QED

Precautionary saving and portfolio allocation: DP by GMM

Marc-Andre Letendre   Gregor Smith
McMaster University   Queen's University

Precautionary saving and portfolio allocation: DP by GMM

Marc-André Letendre and Gregor W. Smith*

Abstract

There is much research on consumption-savings problems with risky labor income and a constant interest rate and also on portfolio allocation with risky returns but non-stochastic labor income. Less is known quantitatively about the interaction between the two forms of risk. Under CRRA utility, undiversifiable income risk should be reflected in both savings rates and portfolio allocations. To quantify these effects in a model of consumption and portfolio choice, we adopt a semi-parametric projection method for solving dynamic programmes, based on generalized method of moments estimation of the parameters of approximate decision rules. We find that background income risk does affect optimal portfolios but that this effect may be difficult to detect empirically.

Keywords: portfolio theory, precautionary saving

JEL classification: D91, G11, C63

1. Introduction

Recently, there has been considerable research into the effect of background risk on risk-taking behavior. It is natural to think that additional background risk will increase local risk aversion. For example, one might expect that households which are relatively exposed to uninsurable labor income risk will hold relatively conservative investment portfolios.

The effect of background risk depends on preferences and the form that risk takes. Pratt and Zeckhauser (1987), Kimball (1993), and Gollier and Pratt (1996) discussed this effect by analyzing behavior when an independent, zero-mean background risk is added to a decision problem. Kimball characterized the necessary and sufficient conditions for this addition to reduce risk-taking as 'standard risk aversion' while Gollier and Pratt characterized them under more general circumstances as 'risk vulnerability.' von Neumann-Morgenstern utility functions with constant relative risk aversion (CRRA) satisfy these conditions. That finding is of interest because CRRA utility is widely used in macroeconomics and finance because of its homogeneity properties.

This paper studies the classic Merton-Samuelson-Hakansson problem of saving and portfolio allocation. In this model there are no frictions so the portfolio choices depend solely on preferences and returns. As in utility-based portfolio models there is a range of risky investments available. As in studies of precautionary saving, there is uninsurable labor income risk. We describe a numerical method – dynamic programming (DP) by the generalized method of moments (GMM) – for measuring the effect on portfolios of adding independent, background income risk. Measuring this effect is of interest as a test of models of saving and investment. Uninsurable income risk differs across households and perhaps also across countries, and savings rates and portfolio shares may reflect this. Some confidence in these predictions is needed in order to predict the welfare effects of insurance schemes and international financial integration, for example.

While measuring the scale of this effect is interesting, there also is a second reason for interest in numerical examples. Adding an independent, zero-mean background risk may not be the most natural way to think of increased background risk. One alternative is a

change with first-degree stochastic dominance, for example in the form of an additional noise which is never positive. A second alternative is a change with second-degree stochastic dominance, for example in the form of a mean-preserving spread. Eeckhoudt, Gollier, and Schlesinger (1996) showed that, with these characterizations of additional background risk, CRRA is not sufficient for an increase in local risk aversion. In other words, added income risk will not necessarily reduce portfolio risk-taking. Thus, numerical examples also can be used (though they may or may not be necessary, pending further theory) to find the sign, and not just the scale, of the effect of income risk on portfolios.

Several applied studies have examined background risk. Elmendorf and Kimball (1991) used a two-period model to predict the effect on portfolios of income taxes which reduce income risk. Guiso, Jappelli, and Terlizzese (1996) studied data from a survey of Italian households. They found that the larger the household's uninsurable income risk the less its demand for risky assets. In subsequent research, Guiso and Jappelli (1998) also found that purchases of health and casualty insurance were higher for households with greater self-reported uncertainty in earnings. They interpreted these results as consistent with decreasing prudence, but they did not relate the scale of the effect to a utility function. Haliassos and Bertaut (1995) examined the U.S. *Survey of Consumer Finances* and found some evidence that income risk reduced the likelihood that a household held stocks.

Recently Bertaut and Haliassos (1997), Heaton and Lucas (1997), Koo (1999), and Viceira (1999) also have studied dynamic portfolio problems with income risk, although with different environments and methods than we adopt. Our economic findings are complementary to theirs. In addition, we develop a solution method which allows statistical inference about the effect of background risk, without parametrizing income and returns processes. We demonstrate this method in simulations and also with historical data.

Section 2 outlines the budgeting problem which we study, and briefly reviews related research on precautionary saving and on portfolio allocation. Section 3 describes the solution method. Section 4 provides numerical results under a range of information sets. We find results which are consistent with theory in that (a) increasing income risk reduces consumption, and (b) increasing return risk reduces consumption (for relative risk aversion

2

greater than one) and tilts portfolios away from risky assets. Moreover, (c) increasing income risk tilts portfolios away from risky assets. However, this effect of background risk on the portfolio is smaller than the effect on consumption and may be difficult to detect statistically. Section 5 contains Monte Carlo evidence and a simple empirical application using historical returns. Section 6 lists conclusions and some possible extensions.

2. Decision-theoretic problem and analytical results

We begin with the economic problem, in which an agent has wealth a_t and savings in period t of $a_t - c_t$, where c_t is consumption. A proportion ω_t of savings is held in an asset with gross return R_1 and a proportion $1 - \omega_t$ is held in an asset with gross return R_{2t}. In our examples R_1 will be a constant, riskless return.

The agent also receives a stochastic income stream $\{y_t\}_{t=0}^{\infty}$ which follows an exogenous process:

$$y_t = y_{pt} y_{nt}, \tag{1}$$

where y_{pt} is a permanent component and y_{nt} is a transitory component. The permanent component follows a logarithmic random walk with drift:

$$\ln y_{pt} = g + \ln y_{pt-1} + \epsilon_{pt}, \tag{2}$$

with $\epsilon_{pt} \sim NID(0, \sigma_p^2)$. The transitory component is

$$\ln y_{nt} = -\frac{\sigma_n^2}{2} + \epsilon_{nt}, \tag{3}$$

where $\epsilon_{nt} \sim NID(0, \sigma_n^2)$ and is independent of ϵ_{pt}. Using the properties of the log-normal density, $E[y_n] = 1$ and $Var[y_n] = \exp(\sigma_n^2) - 1$.

This process restricts labor income to be positive, and allows us to study both temporary and permanent shocks. In addition, this description has been used in studies of saving by Carroll (1997) and Carroll and Samwick (1997) who carefully discussed its calibration to U.S. data. We follow their calibration, to allow comparison with the literature on precautionary saving. The covariance between R_2 and y is zero, so there is no hedging problem and the stream of stochastic income cannot be replicated by trading in the available securities.

The agent is infinitely-lived, with discount factor β and constant relative risk aversion (CRRA) subutility. She chooses a path for consumption $\{c_t\}_{t=0}^\infty$ and portfolio weights $\{\omega_t\}_{t=0}^\infty$ to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\alpha}}{1-\alpha}, \tag{4}$$

subject to the evolution of wealth,

$$a_{t+1} = [\omega_t R_1 + (1-\omega_t)R_{2t+1}] \cdot (a_t - c_t) + y_{t+1} \tag{5}$$

with a_0 given. The vector of exogenous state variables $\{y_t, R_{2t}\}$ is Markov. If there is no labor income and no persistence in returns then the state vector reduces to a_t.

The case without labor income was solved by Samuelson (1969) and Merton (1969). The optimal consumption policy involves consuming a constant fraction of wealth:

$$c_t = \lambda_0 a_t. \tag{6}$$

The marginal propensity to consume, λ_0, is decreasing in β, as more patient planners save more, decreasing in the variance of R_2 when $\alpha > 1$, and decreasing in risk aversion α. With log utility ($\alpha = 1$), the marginal propensity to consume out of wealth is $\lambda_0 = 1 - \beta$. Furthermore, the share of savings invested in the safe asset is of the form:

$$\omega_t = \delta_0, \tag{7}$$

a constant which is independent of a_t and β if $y = 0$.

One of our goals is to compare the effect of income risk on the portfolio with its effect on the savings decision. Analytical solutions for multi-period precautionary savings problems are available only in the special case of constant absolute risk aversion (CARA), as described by Caballero (1990), or used by Kimball and Mankiw (1989). Svensson and Werner (1993) showed that an analytical solution to the joint problem of precautionary saving and portfolio allocation requires CARA utility. But decreasing absolute risk aversion (DARA) in the Arrow-Pratt sense is necessary for standard risk aversion or risk vulnerability, so adopting CARA here would rule out an effect of background income risk on the portfolio.

The case with constant relative risk aversion is of interest here because this utility function displays standard risk aversion. Furthermore, CRRA utility is widely used in macroeconomics where its homogeneity properties are consistent with balanced growth. Isoelastic utility rules out negative consumption. It also implies that high values of wealth are associated with flatter consumption paths, a feature which matches life-cycle evidence (see Blanchard and Mankiw (1988)). Recent studies of precautionary saving under CRRA utility include those of Zeldes (1989), Hubbard, Skinner, and Zeldes (1995), Normandin (1994), and Carroll and Samwick (1997). While these studies examine empirical issues beyond the scope of this paper, each involves a non-stochastic return on a single asset.

It is well-known that rate-of-return risk also may affect saving. Levhari and Srinivasan (1969) showed that, with CRRA utility, increased interest-rate variance leads to increased consumption with relative risk aversion less than one and increased saving with relative risk aversion greater than one. Analytical results are available for the portfolio problem with CRRA utility but these results cannot accommodate undiversifiable labor income risk (see Adler and Detemple (1988), Bodie, Merton, and Samuelson (1992), and Stulz (1984)).

An intermediate case, with constant labor income, was solved by Hakansson (1970) and Merton (1971). Denote the present value of labor income, calculated at the riskless interest rate, by Py. The optimal consumption and investment policies now are of the form:

$$
\begin{aligned}
c_t &= \lambda_0(a_t + Py) \\
\omega_t &= \delta_0 + \delta_1\Big(\frac{Py}{a_t}\Big).
\end{aligned}
\tag{8}
$$

The share of savings invested in the safe asset now is a linear function of the ratio of income to wealth. Merton (1969) and Hakansson (1970) also showed that the transversality condition is satisfied and that wealth remains non-negative in their problems. In turn, Dybvig and Huang (1988) showed that nonnegative wealth precludes arbitrage and rules out doubling strategies, for example. Their result also applies when labor income is stochastic, so when we construct wealth series we check for nonnegative wealth directly rather than imposing a boundary (transversality) condition in the solution algorithm.

Several recent studies have examined the joint precautionary saving/portfolio alloca-

tion problem. Bertaut and Haliassos (1997) studied a model with finite lives, in which a lifetime is divided into three twenty-year periods. Their model includes a portfolio problem and risky income which is calibrated to the U.S. *Survey of Consumer Finances* for 1992. They solved the model numerically under CRRA utility, and found that portfolio risk-taking increases when the variance of persistent income shocks falls. They also found that the variance of temporary income shocks has little effect on the portfolio. Like this study, their research focused on the model without frictions such as short-sales constraints or transactions costs. However, we study the infinite-horizon model, and present a solution method which does not require us to parametrize the income and return processes. The infinite-horizon case is widely used in portfolio and asset-pricing theory, and also has the advantage that analytical solutions are known for special cases. Those cases provide a test of the solution method and initial conditions for other cases.

Koo (1999), Viceira (1999), and Heaton and Lucas (1997) all analyzed dynamic portfolio problems with partly unhedgeable income risk. Koo examined the effect of liquidity constraints, while Viceira examined the effect of a retirement horizon and also studied the hedging problem. These two studies reach one of the same conclusions that we do: income risk has a relatively small effect on optimal portfolios. Heaton and Lucas considered an environment in which there are transaction costs and short-sale constraints. In their examples the portfolio is entirely invested in stocks, a stance which is unaffected by an increase in income risk.

When labor income is stochastic, there is no analytical solution for the optimal consumption and portfolio behavior. However, Viceira (1999) derived the optimal policies in the log-linearized version of the model. They are of the form:

$$c_t = \lambda_0 a_t^{\lambda_1} y_t^{\lambda_2}$$
$$\omega_t = \delta_0. \tag{9}$$

We adopt this approximation, but find the log-linear decision-rule coefficients $\{\lambda_0, \lambda_1, \lambda_2, \delta_0\}$ numerically. While more flexible shape-preserving approximations could be studied, this one is parsimonious, which should aid identification.

6

This functional form nests the analytical solution (6)-(7) for the case with zero labor income. Perturbing the environment around analytically tractable cases is a standard first step in numerical analysis, as described for example by Judd (1997). In addition, this form also draws on recent theoretical results characterizing the optimal decision rules. Carroll and Kimball (1996) proved that consumption is a concave function of wealth when labor income is risky, and the parameter λ_1 is included to measure this effect. Duffie, Fleming, Soner, and Zariphopolou (1997) and Koo (1998) studied a continuous-time version of the problem, in which y follows a geometric Brownian motion, the continuous-time version of the process used here. Koo (1999) studied the discrete-time version of the same problem, and the existence, uniqueness, and characterization of optimal policies. Wealth is non-negative in all our numerical examples, a feature which is equivalent to the liquidity constraint in Koo's study. Thus his results on the existence of optimal policies apply here directly. These studies also show that the decision rule coefficients can be written as functions of the ratio of income to wealth. For example, if a becomes large relative to y then the solution approaches the Merton case. In the environment studied here, we were unable to identify effects of the income-wealth ratio on the optimal portfolio, and so we report results for the simple portfolio rule (9).

There also is significant research which generalizes the description of returns in the multi-period portfolio problem. Analytical expressions for portfolio weights are available only with log utility or a constant investment opportunities set. A variety of studies examine discrete distributions for returns and construct portfolios numerically. Campbell and Viceira (1999) provided some more general results in environments with persistence in returns, using an analytical approximation to describe optimal portfolios. Kim and Omberg (1996) described behavior when risk premiums are persistent. Brandt (1999) estimated portfolio shares from conditional Euler equations, and thus avoided parametric decision rules by conditioning the investment strategy on the current state. He smoothed across states to estimate the optimal rules given empirical return processes.

Our method also can accommodate various preferences and time-varying conditional returns, though returns are *iid* in our simulations. We estimate behavior from Euler

equations, but we include undiversifiable income risk. We wish to study the effect of bearing uninsurable income risk on willingness to bear rate-of-return risk. In addition, we wish to learn whether standard tools of statistical inference can detect the effect of background risk. The next section outlines a simple numerical method for answering these questions.

3. DP by GMM

Several special cases of the budgeting problem without income risk can be solved analytically, for example by the guess-and-verify method. Sargent (1987, pages 31-35) gave two discrete-time examples. Our algorithm is the natural extension of the guess-and-verify method to cases which must be approximated. We guess a form (such as equations (9)) for the approximate decision rules, substitute that guess into the Euler equations, and estimate the decision-rule coefficients by GMM, whether in simulated or historical data. This procedure inherits several appealing features of GMM. For example, it allows us to attach standard errors to decision-rule coefficients, and it does not require complete knowledge of the laws of motion of income and returns. The method also does not impose certainty equivalence, and so it is appropriate for problems of precautionary saving and portfolio allocation. It is very simple to program, requiring only an econometric software package. Smith and Zin (1997) applied the method to historical U.S. business cycles, and so avoided parametric modeling of technology shocks.

Euler-equation orthogonality conditions are used directly as a design criterion, and so this is a projection method as described by Judd (1992), which here is made semi-parametric. We seek a vector of decision rules which give consumption c_t and the investment share of the riskless asset ω_t as functions of the state. The algorithm begins with the Euler equations:

$$E_t \beta R_{it+1} c_{t+1}^{-\alpha} - c_t^{-\alpha} = 0 \tag{10}$$

for $i = 1, 2$.

We next approximate the decision rules. To explain the algorithm we use a simple,

linear guess:

$$c_t = \lambda_0 a_t$$
$$\omega_t = \delta_0$$

$\qquad(11)$

where λ_0 and δ_0 are scalars. Combining the guess with the law of motion for wealth (5) gives:

$$a_{t+1} = (1 - \lambda_0)a_t \cdot [\delta_0 R_1 + (1 - \delta_0)R_{2t+1}] + y_{t+1}. \qquad(12)$$

Using this result and our guess we then replace c_t and c_{t+1} in the Euler equations (10) to define the approximate Euler equation residuals:

$$\eta_{it+1} = \beta R_{it+1} \left[\lambda_0 \Big((1 - \lambda_0)a_t \cdot [\delta_0 R_1 + (1 - \delta_0)R_{2t+1}] + y_{t+1} \Big) \right]^{-\alpha} - \left[\lambda_0 a_t \right]^{-\alpha} \qquad(13)$$

for $i = 1, 2$. Equation (10) holds at an optimum, so the approximation is selected with the same principle:

$$E[\eta_{it+1}|z_t] = 0, \qquad(14)$$

where z_t is a vector of instruments. Equations (14) define an operator from the decision-rule function space to the Euler equation residuals. We seek a zero of this operator, which is an example of computing an implicit function (see Judd, 1998 chapter 6).

In guessing the form for the decision rules, we assume that agents use the correct state vector. For example, income is Markov and returns have no persistence in the simulations in this paper, and the lag length in the guess reflects this information. In applications, longer lag lengths could be included to test for longer memory in the state. However, agents do not know the parametric law of motion for income or returns. For that reason, the information set includes labor income y_t but not its decomposition into permanent and temporary components.

The expectations in equations (14) are calculated in three different ways in the next section. First, we simulate sequences of returns and incomes, R_{2t} and y_t, from a normal random number generator, and use GMM to find the decision-rule parameters which most closely satisfy the Euler equations. The expectations are replaced by simulated sample means. The residual function consists of the fitted Euler equation errors. Finally, we project the residual function on $z_t = \iota$, a vector of ones. Hence we are requiring that

9

the average error be zero in each Euler equation. The algorithm then minimizes the sum of two squared sample means, so that the two decision-rule coefficients are found by just-identified GMM. The GMM criterion function is minimized using the Nelder-Mead simplex routine from Press *et al* (1992). Notice that, if the information set includes a_t and y_t, then the conditional expectations can be calculated simply by using the densities of the exogenous state variables R_{2t+1} and y_{t+1}, because no endogenous variables dated $t+1$ appear in equations (13). Such conditional calculations save the time involved in looping over equation (5) to construct $\{a_t\}$.

Marcet (1988) and Smith (1991) pioneered the use of econometric methods to solve dynamic macroeconomic models. Like those authors, we parametrize the decision rules to incorporate the theoretical restrictions (such as concavity) that are available in this problem. Although simulation is typically less accurate and slower than numerical integration, we adopt it because we wish to study a method which can be used in historical data. There we cannot choose data points, and we may wish to avoid assuming that returns are drawn from a specific probability density function. Hence, although the pseudo-random numbers are normally distributed, in this case the algorithm is not told that. Thus the solution method is a projection method with one hand tied behind its back, or is semi-parametric.

Second, we again estimate the decision-rule parameters by GMM, but this time we recursively construct the endogenous state variable a_t. This step slows the computations but provides additional instruments such as lagged consumption or wealth, and so allows over-identification. This information set simulates the instruments z_t which are available in applications, and so serves as a precursor to the application in section 5.

Third, we next give the algorithm information about the densities of income and returns. Instead of replacing the expectations (14) by their simulated sample means, we evaluate the integrals by numerical integration. Since R_2 and the income shocks ϵ_p and ϵ_n are Gaussian, Gauss-Hermite quadrature is the natural choice as an integration method. The integration uses standard nodes and weights as tabulated by Krylov (1962) for example. This third exercise serves as an accuracy check, for it gives population measures of the effect of income risk on the portfolio, given parameter values, the assumed

normality, and the approximation to the decision rules.

4. Information and numerical results

The environment in the numerical examples begins with the standard portfolio problem without labor income. The solution to this problem involves a linear consumption function (6) and a constant portfolio share (7). This environment provides a behavioral benchmark, an accuracy test of the solution method, and starting values for later experiments.

The mean returns are taken from Table 8.1 of Campbell, Lo, and MacKinlay (1997), which is based on the real return on six-month commercial paper (R_1) and the log real return on the S&P 500 index (R_2) in long spans of annual data. The variance of the risky return is greater than in their table, so as to give a positive portfolio share for the riskless asset. The alternative to adopting this unrealistically high value of σ_{R_2} would be to find an unrealistically large, negative value for δ_0, implying that households borrow to invest in stocks. This is simply the quantity dual of the equity premium puzzle.

Table 1 gives the estimated consumption and portfolio rules. The first two columns give α, the coefficient of relative risk aversion, and σ_{R_2}, the standard deviation of the risky return. The middle columns give estimates of the marginal propensity to consume out of wealth, $\hat{\lambda}_0$, and of the share of assets invested in the safe asset, $\hat{\delta}_0$, along with their standard errors. Both decision-rule coefficients are quoted as percentages. The final columns give the J-test of the overidentifying restrictions, with its p-value.

The top panel of Table 1 conditions on the current value of wealth, so that with two Euler equations the two parameters are just identified. The middle panel loops to construct wealth endogenously. The resulting additional instruments, in the form of lagged Euler-equation residuals, allow a test of overidentification. The bottom panel gives the results of Gauss-Hermite quadrature. The sample size is $10,000$. Seeds are held constant across parametrizations. At $T = 10,000$, sampling variability remains in the estimates, and GMM standard errors are given in brackets.

Each aspect of Table 1 conforms to theory. For a given risk aversion, the share of the

portfolio invested in the riskless asset rises as the variance of the risky return rises. For a given structure of returns, the marginal propensity to consume falls and the share of wealth held in the riskless asset rises as risk aversion rises. With log utility, the marginal propensity to consume is independent of the structure of returns and equal to $1-\beta$. When α exceeds one, however, an increase in the variance of the risky return reduces the marginal propensity to consume, even though the portfolio share allocated to that investment also falls. In the top panel, the algorithm also yielded the same coefficients whatever the constant value of wealth.

Table 1 provides two accuracy tests in addition to conforming with analytical results. In the central panel of Table 1, the instruments consist of a constant term and a lag of each Euler equation residual. With six moment conditions the two decision-rule coefficients are overidentified. The last column contains the J-test statistic, which is asymptotically distributed as $\chi^2(4)$. In each case, the J-test does not reject at conventional significance levels. The standard errors do not change from the first to the second panel because in this environment the decision-rule form is known analytically, and the exogenous state variables have no persistence. Thus the added instruments do not add to the precision of the estimates. This result will not hold when DP by GMM is used with approximate decision rules.

The second accuracy test is provided by the comparison with numerical integration, in the bottom panel. Like other simulation methods, DP by GMM is inefficient as an experimental method. But the comparison of panels in Table 1 shows that it is quite accurate, for the estimates are virtually identical to those found by Gauss-Hermite quadrature.

Next, Table 2 introduces risky labor income, following the geometric random walk (1)-(3) with temporary and permanent shocks. The parameter values for the income process are the same as those of Viceira (1999) and very similar to those of Carroll (1997), who discussed their calibration to U.S. panel data from the PSID. The drift term g is set so that income growth averages 3 percent per year. Permanent shocks to labor income have a standard deviation, σ_p, of 15 percent per year, while temporary shocks have a standard deviation, σ_n, of 10 percent per year. Baseline parameters are given at the top of the table,

and the associated decision rules are given in the first row. Later rows vary the parameters relative to the baseline list. Decision-rule coefficients again are quoted as percentages.

The solution method is DP by GMM with endogenous wealth, as in the central panel of Table 1. The idea is that the top panel of Table 1 gives the investigator too little information – by fixing the value of wealth and conditioning on current income – while the bottom panel gives too much, in the form of knowledge of the distribution of returns.

Recursively constructing wealth again introduces instruments which now can add to precision, given the persistence in the exogenous state variable y. Again the J-test of the overidentifying restrictions serves as an accuracy test. None of the test statistics leads to a rejection at the 5 percent level, except in the final row where income is deterministic. In that case, we know that the functional form (9) for the consumption rule is wrong: Hakansson (1970) and Merton (1971) showed that consumption is linear in wealth (as in equation (8)) in the portfolio problem with a deterministic income stream. This line in Table 2 shows that the GMM diagnostics have some power, in that they reject when the approximation is poor, while Table 1 showed that they did not reject when the functional form was correct. These findings add to our confidence in the results elsewhere in Table 2. Thus the functional form and numerical optimization seem to allow a study of the two-risks problem. As we perturbed the environment by adding income risk, we also tried to identify effects of the income-wealth ratio on the optimal portfolio but were not successful. Thus we report results for the simple linear rules in equation (9).

The coefficients in the consumption rule once again satisfy theory. The consumption function is shifted down by increases in risk aversion (rows 2 and 3). With log utility, the consumption function is insignificantly affected by changes in the variance of returns (rows 4 and 5). Decreases in income risk, whether in the temporary (rows 6 and 7) or permanent (rows 8 and 9) component of income, lead to significant decreases in consumption tilting ($\hat{\lambda}_0$ rises and $\hat{\lambda}_1$ falls), consistent with precautionary saving. Finally, $\hat{\lambda}_1$ is less than 100 percent throughout the table and its standard error generally is small. Thus consumption is a concave function of wealth, consistent with the result of Carroll and Kimball (1996).

The portfolio behavior also is consistent with theory. The share of wealth held in

the safe asset rises as risk aversion rises (rows 2 and 3). These changes are large and significant, given the standard errors. Increases in the variance of the risky return also lead to increases in the share held in the safe asset (rows 4 and 5), but these changes are smaller and have large standard errors. Finally, decreases in income risk (rows 6 to 10) lead to riskier investment behavior, as one would predict from risk vulnerability or standard risk aversion. However, these effects on the portfolio are difficult to measure precisely and sometimes are insignificant. The effect of a given change in income risk is easier to detect in the consumption rule than in the portfolio decision.

These large-sample results suggest that models of precautionary saving may mis-predict if they ignore the endogenous effect of income risk on the portfolio rate of return. Similarly, the predicted effects on the savings rate and on the portfolio strategy of a change in rate-of-return risk may be sensitive to the scale of income risk. While the effects of income risk on the portfolio can be difficult to measure precisely – because the portfolio share has a large estimated standard error – the same can be said of the traditional effects of rate-of-return risk.

5. Monte Carlo evidence and historical application

We next present a simple application, which uses GMM to estimate optimal consumption rules and portfolio shares. The application shows that the method can be implemented easily with historical returns.

We consider annual U.S. returns data from 1934 to 1998, so that $T = 65$. This frequency is chosen because the parameters of the income process are calibrated from annual PSID data. For comparability with section 4, R_1 is set at a constant value of 1.02. The $\{R_{2t}\}$ series is defined as R_1 plus the value-weighted return on the NYSE and AMEX including dividends (from CRSP, label vwretd) minus the three-month T-bill rate (from FRED at the Federal Reserve Bank of St.Louis, label tb3ms). Inflation cancels with this definition of R_2, so this can be viewed as R_1 plus a real excess return. This series is converted from monthly to annual frequency as a geometric average. It has a mean of 1.095, a standard deviation of 0.182, and a first-order autocorrelation of 0.13. This return

thus has minimal persistence at annual frequency, which again mimics the environment in section 4. For that reason, R_{2t} is not included in the state vector.

The historical return series is combined with a simulated income series, as if the PSID had begun in 1934. This approach is taken because there is no annual series on labor income for this time period. In addition, the use of a simulated series ensures independence from returns, so that there is no hedging problem. Estimates are averaged over 1000 replications to minimize simulation sampling error.

Table 3 contains the results. The top panel examines the sampling variability in the estimated decision rules using Monte Carlo methods. Here we use the baseline parametrization of table 2 (with both income and returns simulated), but now with 1000 draws over 65 observations. The first goal of the simulations is to see whether standard tools of inference are reliable. We thus calculate the average standard error across replications and compare it to the Monte Carlo, cross-replication standard error. The second goal is to see what economic inferences can be drawn at these sample sizes. The point estimates of the decision-rule coefficients are consistent with theory, in that consumption is concave in wealth. The average value of the J-test statistic is 5.66, which does not lead to a rejection of the over-identifying restrictions at the 5 percent significance level.

The two sets of brackets below the average coefficient estimates contain, first, the average GMM standard error and, second, the Monte Carlo, cross-replication standard error. The Monte Carlo standard errors are less than half the asymptotic GMM ones, so that using the latter would underestimate the precision. This discrepancy between the two methods of calculating standard errors naturally declines with the sample size, but it remains large even at larger sample sizes such as 200 or 500. We draw the lesson that repeated simulation should be used to find standard errors in applications that involve some simulated state variables.

In the simulation in the top panel of Table 3, even the Monte Carlo standard errors are large relative to the coefficients, so that none of the coefficients is statistically significant. However, the decision rules are estimated with much more precision in the historical data. The next panel of Table 3 summarizes results using the historical sequence of 65 equity

premiums, with a constant riskless interest rate. Using the Monte Carlo standard errors for inference, estimation of the baseline model yields two significant coefficients, $\hat{\lambda}_1$ and $\hat{\delta}_0$.

At the baseline parameter values, with $\alpha = 1$, the optimal policy is to borrow heavily at the riskless rate so as to invest 217 percent of wealth in the stock market (with a standard error of 18 percent). This large short position in the riskless asset arises because the application uses historical equity returns, which are less variable than the equity returns in the simulations.

The next two rows of Table 3 show that one can detect the effects of changes in α, the coefficient of relative risk aversion. For example, higher values of risk aversion lead to portfolios which are tilted significantly towards the safe asset, as one would expect. When $\alpha = 5$, 166 percent of wealth is invested in equities, with a standard error of 16 percent. However, the average J-test rejects the over-identifying restrictions for large values of α. Although the tendency of the J-test to over-reject in small samples is well-known this tendency probably cannot explain such large test statistics.

In contrast, the effects of changes in the variance of shocks to labor income, whether temporary or permanent, cannot be detected at this sample size. Here changes in the consumption function and in the portfolio share are both insignificant. Drawing historical inferences about the effect of background risk at this sample size would require additional information, for example by studying panel data or by assuming that households know the parameters of the income and returns processes.

This application has a small number of observations because it is set up to match the earlier examples, calibrated to annual data from the PSID. Even so, several features of the decision rules – and of their changes in response to changed parameters – are statistically significant. The same method could be applied to portfolio problems with higher frequency data.

6. Conclusion

Modeling responses to income and investment risk jointly is potentially important for

a variety of economic questions. For example, changes in tax policies or social security which reduce income risk are predicted to reduce savings rates by standard models of precautionary saving. But the effect on growth of a decline in saving may be offset by a portfolio shift towards investments with higher expected returns. Similarly, the welfare gains from international portfolio diversification depend on the response of portfolios, which in turn depends on background income risk.

The method in this paper provides a way to examine the scale of the interaction between risks that has been studied in recent theoretical papers. It is designed to be useful in applied work, for it does not require knowledge of the distributions of income and return shocks, provided these satisfy the weak requirements of GMM estimation. We have studied two risks in an infinite-horizon problem of saving and portfolio choice. So far, our research suggests that the effect of background risk on portfolios is smaller than the effect on savings rates and may be difficult to detect empirically.

The combination of GMM and simple decision rules studied here stands up to several accuracy tests. First, given the form of decision rules, GMM and quadrature (which exploits the distributions of returns and income) yield very similar results. Second, the decision rules pass the J-test in simulations and in a number of the historical applications. Third, the results reproduce analytical results: with log utility and no income risk, the marginal propensity to consume is $1 - \beta$; with log utility the savings rate is independent of rate-of-return risk. Fourth, this combination also is consistent with known, qualitative properties of this decision-theory problem. In addition to the effect of income risk on the portfolio, for example, income risk affects the savings rate, return risk affects the portfolio (and also the savings rate if $\alpha > 1$), and risk aversion affects the savings rate and the portfolio. Also, consumption is a concave function of wealth when both income and rate-of-return risks are present.

Possible extensions to this work include continuing to examine more flexible approximations to the decision rules, allowing for cross-correlation in exogenous state variables, and calculating expected utility. Ultimately we hope to study the extent to which cross-country differences in income risk can account for cross-country differences in savings

rates and portfolios. To do so would require a general equilibrium environment, beyond the decision-theoretic problem studied here.

References

Adler, M., Detemple, J.B., 1988. On the optimal hedge of a nontraded cash position. Journal of Finance 43, 143-153.

Bertaut, C.C., Haliassos, M., 1997. Precautionary portfolio behavior from a life-cycle perspective. Journal of Economic Dynamics and Control 21, 1511-1542.

Blanchard, O.J., Mankiw, N.G., 1988. Consumption: beyond certainty equivalence. American Economic Review (P) 78, 173-177.

Bodie, Z., Merton, R.C., Samuelson, W.F., 1992. Labor supply flexibility and portfolio choice in a life cycle model. Journal of Economic Dynamics and Control 16, 427-449.

Brandt, M.W., 1999. Estimating portfolio and consumption choice: a conditional Euler equations approach. Journal of Finance 54, 1609-1646.

Caballero, R.J., 1990. Consumption puzzles and precautionary savings. Journal of Monetary Economics 25, 113–136.

Campbell, J.Y, Lo, A.W., MacKinlay, A.C., 1997. The econometrics of financial markets. (Princeton University Press).

Campbell, J.Y., Viceira, L.M., 1999. Consumption and portfolio decisions when expected returns are time varying. Quarterly Journal of Economics 114, 433-495.

Carroll, C.D., Kimball, M.S., 1996. On the concavity of the consumption function. Econometrica 64, 981-992.

Carroll, C.D., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. Quarterly Journal of Economics 112, 1-56.

Carroll, C.D., Samwick, A.A., 1997. The nature of precautionary wealth. Journal of Monetary Economics 40, 41-71.

Duffie, D., Fleming, W., Soner, H.M., Zariphopolou, T., 1997. Hedging in incomplete markets with HARA utility. Journal of Economic Dynamics and Control 21, 753-782.

Dybvig, P.H., Huang, C., 1988. Nonnegative wealth, absence of arbitrage, and feasible consumption plans. Review of Financial Studies 1, 377-401.

Eeckhoudt, L., Gollier, C., Schlesinger, H., 1996. Changes in background risk and risk taking behavior. Econometrica 64, 683-689.

Elmendorf, D.W., Kimball, M.S., 1991. Taxation of labor income and the demand for risky assets. NBER Working Paper No. 3904.

Gollier, C., Pratt, J.W., 1996. Risk vulnerability and the tempering effect of background risk. Econometrica 64, 1109-1123.

Guiso, L., Jappelli, T., Terlizzese, D., 1996. Income risk, borrowing constraints, and portfolio choice. American Economic Review 86, 158-172.

Guiso, L., Jappelli, T., 1998. Background uncertainty and the demand for insurance against insurable risks. Geneva Papers on Risk and Insurance Theory 23, 7-27.

Hakansson, N.H., 1970. Optimal investment and consumption strategies under risk for a class of utility functions. Econometrica 38, 587-607.

Haliassos, M., Bertaut, C.C., 1995. Why do so few hold stocks? Economic Journal 105, 1110-1129.

Heaton, J., Lucas, D., 1997. Market frictions, savings behavior, and portfolio choice. Macroeconomic Dynamics 1, 76-101.

Hubbard, R.G., Skinner, J., Zeldes, S.P., 1995. Precautionary saving and social insurance. Journal of Political Economy 103, 360–399.

Judd, K.L., 1992. Projection methods for solving aggregate growth models. Journal of Economic Theory 58, 410-452.

Judd, K.L., 1997. Computational economics and economic theory: substitutes or complements? Journal of Economic Dynamics and Control 21, 907-942.

Judd, K.L., 1998. Numerical methods in economics. (M.I.T. Press).

Kim, T.S., Omberg, E., 1996. Dynamic nonmyopic portfolio behavior. Review of Financial Studies 9, 141-161.

Kimball, M.S., Mankiw, N.G., 1989. Precautionary saving and the timing of taxes. Journal of Political Economy 97, 863-879.

Kimball, M.S., 1993. Standard risk aversion, Econometrica 61, 589-611.

Koo, H.K., 1998. Consumption and portfolio selection with labor income: a continuous time approach. Mathematical Finance 8, 49-65.

Koo, H.K., 1999. Consumption and portfolio selection with labor income: a discrete-time approach. Mathematical Methods of Operations Research 50, 219-243.

Krylov, V.I., 1962. Approximate calculation of integrals. (trans. A.H. Stroud) (Macmillan, New York).

Levhari, D., Srinivasan, T.N., 1969. Optimal savings under uncertainty. Review of Economic Studies 36, 153–163.

Marcet, A., 1988. Solving non-linear models by parametrizing expectations. mimeo, Carnegie Mellon University, Graduate School of Industrial Administration.

Merton, R.C., 1969. Lifetime portfolio selection under uncertainty: The continuous-time case. Review of Economics and Statistics 51, 247-257.

Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous-time model. Journal of Economic Theory 3, 373-413.

Normandin, M., 1994. Precautionary saving: An explanation for excess sensitivity of consumption. Journal of Business and Economic Statistics 12, 205-219.

Pratt, J.W., Zeckhauser, R.J., 1987. Proper risk aversion. Econometrica 55, 143-154.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1992. Numerical recipes in FORTRAN: The art of scientific computing. (Cambridge University Press, New York).

Samuelson, P.A., 1969. Lifetime portfolio selection by dynamic stochastic programming. Review of Economics and Statistics 51, 239-246.

Sargent, T.J., 1987. Dynamic macroeconomic theory. (Harvard University Press).

Smith, A.A. Jr., 1991. Solving stochastic dynamic programming problems using rules of thumb. Discussion Paper 816, Department of Economics, Queen's University.

Smith, G.W., Zin, S.E., 1997. Real business cycle realizations. Carnegie-Rochester Conference Series on Public Policy 47, 243-280.

Stulz, R.M., 1984. Optimal hedging policies. Journal of Financial and Quantitative Analysis 19, 127-140.

Svensson, L.E.O., Werner, I., 1993. Nontraded assets in incomplete markets: Pricing and portfolio choice. European Economic Review 37, 1149-1168.

Viceira, L., 1999. Optimal portfolio choice for long-horizon investors with nontradable labor income. NBER Working Paper No. 7409.

Zeldes, S.P., 1989. Optimal consumption with stochastic income: Deviations from certainty equivalence. Quarterly Journal of Economics 104, 275-298.

TABLE 1: Portfolio Problem

$$\beta = 0.95 \qquad a_1 = 100$$
$$R_1 = 1.02 \qquad R_{2t} = 1.06 + \sigma_{R_2}\epsilon_{Rt} \qquad \epsilon_{Rt} \sim IIN(0,1)$$
$$c_t = \lambda_0 a_t \qquad \omega_t = \delta_0$$

α	σ_{R_2}	$\hat{\lambda}_0$	(se)	$\hat{\delta}_0$	(se)	$J(4)$ (p)

DP by GMM
$$z_t = \{\iota\}$$
$$T = 10,000 \quad a_t = 100$$

α	σ_{R_2}	$\hat{\lambda}_0$	(se)	$\hat{\delta}_0$	(se)	$J(4)$ (p)
1	0.30	5.00	(0.00)	55.11	(3.29)	— —
1	0.45	5.00	(0.00)	79.86	(2.24)	— —
3	0.30	3.18	(0.03)	84.79	(1.15)	— —
3	0.45	3.07	(0.02)	93.24	(0.76)	— —
5	0.30	2.72	(0.02)	90.86	(0.69)	— —
5	0.45	2.64	(0.01)	95.94	(0.46)	— —

DP by GMM
$$z_t = \{\iota \ \ \eta_{1t} \ \ \eta_{2t}\}$$
$$T = 10,000 \qquad \text{Endogenous wealth}$$

α	σ_{R_2}	$\hat{\lambda}_0$	(se)	$\hat{\delta}_0$	(se)	$J(4)$ (p)
1	0.30	5.00	(0.00)	55.09	(3.29)	
1	0.45	5.00	(0.00)	79.83	(2.24)	
3	0.30	3.18	(0.03)	84.79	(1.15)	1.29 (0.86)
3	0.45	3.07	(0.02)	93.23	(0.76)	1.33 (0.86)
5	0.30	2.72	(0.02)	90.86	(0.69)	1.28 (0.86)
5	0.45	2.64	(0.01)	95.93	(0.46)	1.30 (0.86)

Gauss-Hermite Quadrature

α	σ_{R_2}	$\hat{\lambda}_0$	(se)	$\hat{\delta}_0$	(se)	$J(4)$ (p)
1	0.30	5.00		55.49		— —
1	0.45	5.00		80.01		— —
3	0.30	3.17		84.96		— —
3	0.45	3.07		93.24		— —
5	0.30	2.71		90.94		
5	0.45	2.63		95.97		

Notes: λ_0 is the marginal propensity to consume out of wealth and δ_0 is the share of savings invested in the riskless asset; both are quoted as percentages. Brackets contain standard errors.

TABLE 2: Effects of Income Risk

$$\beta = 0.95 \qquad a_1 = 100 \qquad T = 10,000$$

$$R_1 = 1.02 \qquad R_{2t} = 1.06 + \sigma_{R_2}\epsilon_{Rt} \qquad \epsilon_{Rt} \sim IIN(0,1)$$

$$y_t = y_{pt}y_{nt} \qquad \ln y_{pt} = g + \ln y_{pt-1} + \epsilon_{pt} \qquad \ln y_{nt} = -\frac{\sigma_n^2}{2} + \epsilon_{nt}$$

$$g = \log 1.03 - 0.5(\sigma_p^2 + 2\sigma_n^2) \qquad \epsilon_{pt} \sim IIN(0,\sigma_p^2) \qquad \epsilon_{nt} \sim IIN(0,\sigma_n^2)$$

$$c_t = \lambda_0 a_t^{\lambda_1} y_t^{\lambda_2} \qquad \omega_t = \delta_0$$

Instruments: 1, η_{1t}, η_{2t}

Baseline parameters: $\alpha = 1$, $\sigma_{R_2} = 0.30$, $\sigma_n = 0.1$, $\sigma_p = 0.15$

	$\hat{\lambda}_0$	(se)	$\hat{\lambda}_1$	(se)	$\hat{\lambda}_2$	(se)	$\hat{\delta}_0$	(se)	$J(2)$	(p)
Baseline	139.31	(2.09)	89.31	(0.71)	7.44	(0.58)	-2.24	(6.48)	3.37	(.19)
$\alpha = 3$	123.55	(38.15)	82.63	(8.44)	-2.54	(39.83)	67.60	(2.38)	4.16	(.12)
$\alpha = 5$	83.01	(7.66)	82.88	(2.39)	-33.72	(8.53)	84.95	(4.47)	2.12	(.35)
$\sigma_{R_2} = 0.4$	137.71	(2.99)	89.70	(0.89)	8.11	(0.63)	16.86	(7.48)	3.67	(.16)
$\sigma_{R_2} = 0.5$	134.63	(3.84)	90.39	(1.10)	7.96	(0.82)	25.75	(9.55)	4.64	(.10)
$\sigma_n = 0.05$	149.87	(3.28)	87.26	(0.74)	11.33	(0.55)	-15.43	(5.32)	4.03	(.13)
$\sigma_n = 0.0$	152.85	(5.47)	86.75	(1.07)	12.07	(0.82)	-17.35	(8.76)	5.87	(.05)
$\sigma_p = 0.10$	150.37	(2.76)	87.11	(0.66)	11.28	(0.43)	-12.96	(6.22)	5.48	(.06)
$\sigma_p = 0.05$	156.87	(3.69)	85.83	(0.76)	12.91	(0.55)	-15.28	(5.51)	6.09	(.05)
$\sigma_p,\ \sigma_n = 0.0$	159.76	(16.63)	85.33	(2.97)	13.97	(2.47)	-21.99	(36.41)	22.84	(.00)

Notes: λ_0, λ_1, and λ_2 are the parameters of the consumption function while δ_0 is the share of savings invested in the riskless asset; all are quoted as percentages. Brackets contain standard errors.

TABLE 3: Monte Carlo Evidence and Application

$$\beta = 0.95 \qquad a_1 = 100 \qquad T = 65$$

$$y_t = y_{pt}y_{nt} \qquad \ln y_{pt} = g + \ln y_{pt-1} + \epsilon_{pt} \qquad \ln y_{nt} = -\frac{\sigma_n^2}{2} + \epsilon_{nt}$$

$$g = \log 1.03 - 0.5(\sigma_p^2 + 2\sigma_n^2) \qquad \epsilon_{pt} \sim IIN(0, \sigma_p^2) \qquad \epsilon_{nt} \sim IIN(0, \sigma_n^2)$$

$$c_t = \lambda_0 a_t^{\lambda_1} y_t^{\lambda_2} \qquad \omega_t = \delta_0$$

Instruments: 1, η_{1t}, η_{2t}

Baseline parameters: $\alpha = 1$, $\sigma_n = 0.1$, $\sigma_p = 0.15$

	$\overline{\hat{\lambda}_0}$	$\overline{\hat{\lambda}_1}$	$\overline{\hat{\lambda}_2}$	$\overline{\hat{\delta}_0}$	$\overline{J(2)}\ (p)$
$R_1 = 1.02$	$R_{2t} = 1.06 + \sigma_{R_2}\epsilon_{Rt}$	$\sigma_{R_2} = 0.30$	$\epsilon_{Rt} \sim IIN(0,1)$		
Baseline	45.50	76.31	11.04	34.16	5.66 (.06)
	(182.46)	(187.41)	(57.19)	(173.08)	
	(51.52)	(33.95)	(31.67)	(75.92)	
$R_1 = 1.02$	$R_{2t} = R_1 + VWRET - TBILL$		$(1934 - 1998)$		
Baseline	3.97	100.43	4.29	-117.22	5.10 (.08)
	(35.03)	(83.62)	(182.04)	(119.25)	
	(14.76)	(14.21)	(30.61)	(18.26)	
$\alpha = 3$	149.35	46.65	-0.28	-75.94	17.99 (.00)
	(660.85)	(77.48)	(93.25)	(242.24)	
	(43.17)	(4.03)	(5.13)	(12.60)	
$\alpha = 5$	333.28	30.54	-0.14	-66.48	17.05 (.00)
	(970.40)	(55.28)	(38.15)	(265.88)	
	(80.08)	(3.59)	(4.19)	(16.28)	
$\sigma_n = 0.05$	4.42	100.08	3.54	-116.90	4.17 (.12)
	(36.27)	(93.24)	(251.84)	(130.13)	
	(8.37)	(14.81)	(40.41)	(19.17)	
$\sigma_n = 0.0$	4.60	101.68	0.66	-116.05	4.68 (.10)
	(32.13)	(98.99)	(287.89)	(136.48)	
	(7.93)	(14.33)	(42.66)	(18.44)	
$\sigma_p = 0.10$	3.51	99.65	5.75	-117.90	6.57 (.04)
	(39.66)	(78.41)	(199.76)	(113.72)	
	(27.71)	(13.96)	(34.28)	(18.12)	

$\sigma_p = 0.05$	5.09	99.50	5.18	-117.88	8.60 (.01)
	(24.74)	(73.36)	(209.30)	(108.51)	
	(8.12)	(13.37)	(37.63)	(17.28)	

Notes: The number of replications is 1000. The first standard error in brackets is the cross-replication average of the GMM standard errors. The second standard error is the Monte Carlo one. $\overline{J}$ is the average of the J-test statistics. Brackets contain its asymptotic p-value.