© *The Journal of Risk and Insurance*, 2007, Vol. 74, No. 4, 895-906

Book Reviews

The Calculus of Retirement Income: Financial Models for Pension Annuities and Life Insurance, by Moshe A. Milevsky, 2006, Cambridge, England: Cambridge University Press, 321 pages

Reviewer: David F. Babbel, The Wharton School, University of Pennsylvania

Moshe A. Milevsky is one of the most skillful and engaging writers of our profession. His recent book, notwithstanding its rather daunting title—*The Calculus of Retirement Income*—is refreshingly written in a conversational first person voice and surpasses the already high expectations that his previous work has fostered.

Professor Milevsky does a remarkable job at melding several disciplines in his text: financial economics, insurance, mathematics, and actuarial science. The text serves as a bridge for actuaries into the world of finance, and for financial economists into the world of pensions, insurance, and actuarial science. This user-friendly bridge is not only a notable achievement but an essential step toward joining the insights achievable through the various disciplines and equipping researchers to make further advances in the fields of pensions and retirement income.

Milevsky introduces and develops, from the financial economics perspective, the basic actuarial models that underlie the pricing and risk analysis of life-contingent annuities, pensions, and life insurance. He handles the mathematical and actuarial models so deftly that a person with a basic understanding of calculus and finance would find his text quite accessible. Moreover, he sprinkles his analysis with practical and witty examples (such as folk wisdom from the famous Brazilian playboy, the late Jorge Guinle) that demonstrate to the reader their relevance and analytical power of the tools presented.

He introduces the challenges of rational wealth decumulation at retirement through the metaphor of a drunk gambler he once witnessed who exhibited a rather curious roulette wagering and waitress tipping strategy. Here is not the place to reveal the parallels he draws between the drunk gambler and the retiree facing investment and consumption decisions, but suffice it to say that they rivet the reader's attention and impel her to turn the next page. Next, he presents the demographics and economics of what has been termed the central policy concern of our time—providing retirement security for our aging populations.[1]

With this backdrop, he presents the models of actuarial finance that will allow him later to address the challenges of wealth management during retirement. His presentation

[1] See Blitzstein, Mitchell, and Utkus (2006).

begins with modeling of the human life cycle, including a helpful section on real versus nominal planning with inflation. He follows this with an illuminating treatment of models of human mortality, including the modeling of joint lifetimes and a technical note on the incomplete gamma function in Excel. This chapter is followed by two basic chapters on finance, which would be less helpful for financial economists than for actuaries except that they introduce a consistent nomenclature and approach for financial assets that correspond to those used for life-contingent liabilities. Moreover, he accomplishes this in an elegant and succinct fashion. These chapters introduce concepts such as duration and convexity, Brownian motion, simulation, and space-time diversification in a tight and cohesive fashion.

The next three chapters contain models of pension life annuities, life insurance, and a comparison of defined benefit and defined contribution pensions. These chapters are interlaced with thoughtful questions and intriguing subjects, with headings such as "The Wrong Way to Value Pension Annuities," "How Is the Annuity's Income Taxed?" "How Much Life Insurance Do You Need?" and "A Free (Last) Supper." The questions and intrigues that he poses are always addressed by using the tools that have already been presented, validating and solidifying their use to the reader.

The book concludes with two sections on which Milevsky has authored numerous studies. The first presents wealth management applications and implications, including an examination of longevity insurance, self-annuitization, advanced-life delayed annuities, and a demonstration of who incurs mortality risk and investment rate risk. The second presents advanced topics such as variable annuities, guaranteed minimum accumulation, withdrawal and death benefits, and the utility of annuitization, including the real option to defer annuitization.

The scope and treatment of financial models for pension annuities and life insurance are well conceived; however, it is inevitable that some readers will infer too much from the stylized examples given on self-annuitization and advanced-life delayed annuities, both of which hinge on the underlying distribution assumed. This is worrisome, especially when the inferences are reported by the press and shape the opinions of policymakers. Therefore, we have a suggestion for the next edition of his volume. It reflects the direction in which the finance profession is heading and would serve as a valuable segue into that literature.

On page 88 of his text, Milevsky speaks of the random annualized growth rate of the investment portfolio and states: "There is, of course, a multitude of statistical distributions that we could select to describe the annualized growth rate g, and you might be surprised to learn that there are many different distributions that have been proposed for g over the last century of scholarly writing, during which thousands of research papers have been written on this topic. Forecasting the evolution of g over time is something of a holy grail in the field of financial economics. Although I could probably write an entire book on models and calibration of g, I will take the easy path and assume that g satisfies the most ubiquitous of all statistical quantities: the normal distribution."

His choice is understandable, but a subsequent volume or a revised edition of this volume would be improved by also showing how the economic implications would be altered were another distribution to be used that better captures the extreme events,

which are increasingly commonplace in the arena of investments. The g-and-h transformation of the normal, introduced by Tukey (1977), is a natural candidate for modeling extreme events because it subsumes most of the other distributions in common use and accommodates any possible combination of skewness and kurtosis that lies within the possible range.[2] Moreover, it is easy to use with standard spreadsheet programs.

Martinez and Iglewicz (1984) have shown that the g-and-h distribution covers most of the Pearsonian family of distributions up to an adequate approximation and can also generate a variety of other types of distributions. They show how more than 12 distributions can be generated from the g-and-h distribution by choosing the values of A, B, g, and h.[3] The normal and lognormal distributions are two very important special cases of the g-and-h distribution. This property of the g-and-h distribution makes it a very general leptokurtic distribution for modeling asset returns. Badrinath and Chatterjee (1988, 1991) demonstrated that the g-and-h is excellent in modeling the returns on common stock, while Dutta and Babbel (2005) achieved similar results for interest rates and derivative fixed income assets.[4]

When security returns are fit to a g-and-h distribution, the implications for lifetime income security are much worse than when a normal distribution is assumed, owing to the reverse dollar cost averaging nature of decumulation. These different distributions change the implications of the retirement income analysis, and would be a valuable addition for practitioners and policymakers as well, while providing researchers with a foray into making contributions to the economic literature on retirement income. And because most other common distributions are either subsets or effectively subsumed by g-and-h, it would serve as a useful point of departure for future studies.

In summary, *The Calculus of Retirement Income* offers an enjoyable trip to a reader interested in deepening her understanding of this vital topic. Throughout, it provides well-chosen supplemental references on each of its major topics. It has the feel of being led personally by an experienced guide on a safari of discovery. It is well worth the trip!

REFERENCES

Badrinath, S. G., and S. Chatterjee, 1988, On Measuring Skewness and Elongation in Common Stock Return Distributions: The Case of the Market Index, *Journal of Business*, 61(4): 451-472.

[2] Note that "g" in the g-and-h distribution is not to be confused with Milevsky's g, the annualized growth rate of the investment portfolio; rather, it is a parameter related to skewness.

[3] In addition to normal and lognormal distributions, they showed how to generate uniform, student-t, exponential, double exponential, Cauchy, beta, gamma, Weibull, chi-square, and logistic distributions. We can also get Burr III as a combination of the beta and gamma distributions.

[4] An alternative to the g-and-h distribution would be the GB2 distribution analyzed by Bookstaber and McDonald (1987). While it is more cumbersome and less stable than g-and-h, it does accommodate a sizable subset of all feasible combinations on the skewness-kurtosis plane that are mapped by g-and-h and has often been used for modeling security returns.

Badrinath, S. G., and S. Chatterjee, 1991, A Data-Analytic Look at Skewness and Elongation in Common-Stock-Return Distributions, *Journal of Business & Economic Statistics*, 9(2): 223-233.

Blitzstein, D., O. S. Mitchell, and S. P. Utkus, eds., 2006, *Restructuring Retirement Risks* (Oxford: Oxford University Press).

Bookstaber, R. M., and J. B. McDonald, 1987, A General Distribution for Describing Security Price Returns, *Journal of Business*, 60(3): 401-424.

Dutta, K., and D. F. Babbel, 2005, Extracting Probabilistic Information from the Price of Interest Rate Options: Tests of Distributional Assumptions, *Journal of Business*, 78(3): 841-870.

Martinez, J., and B. Iglewicz, 1984, Some Properties of the Tukey g and h Family Distributions, *Communications in Statistics—Theory and Methods*, 13(3): 353-369.

Tukey, J. W., 1977, *Exploratory Data Analysis* (Reading, MA: Addison-Wesley).

Measuring Risk Aversion, by Donald J. Meyer and Jack Meyer, 2006, NOW Publishers, Inc., 101 pages

Reviewer: Joseph G. Eisenhauer, Professor and Chair of Economics, Wright State University, Dayton, Ohio

The empirical literature on risk aversion contains a vast array of estimates undertaken in various contexts, with different types of data, different populations, different models and assumptions regarding utility, and different estimation techniques. The contexts for studying risk have included agricultural production, insurance purchases, financial investments, gambling, and intertemporal consumption, among others; the decision makers in these studies have included college students, working adults, retirees, citizens of less developed countries, and those in developed economies; and cross-sections, time series, laboratory experiments, and survey responses have all been used as data. As a consequence, the results have been quite diverse. Some studies have estimated relative risk aversion to be less than 1 (Hansen and Singleton, 1982), others obtain values that are generally between 1 and 10 (Szpiro, 1986a,1986b; Eisenhauer and Ventura, 2003), still others conclude that relative risk aversion must be well in excess of 10 in order to explain observed behavior (Mehra and Prescott, 1985), and outliers have been reported in the hundreds (Halek and Eisenhauer, 2001). Conflicting results can also be found concerning the slope of the risk aversion function—that is, whether risk aversion increases, remains constant, or decreases as the argument of the utility function rises. Blake (1996) and Levy (1994), for example, found evidence of decreasing relative risk aversion (DRRA), while Eisenhauer and Ventura (2003) reported increasing relative risk aversion (IRRA).

A number of observers have been disconcerted by this lack of consistency across studies. Mehra and Prescott (1985) and Kocherlakota (1996), finding that stock market investments among households have historically been far too limited to be consistent with widely reported relative risk aversion values below 10, refer to the discrepancy as the equity premium puzzle. Similarly, Gollier (2001, pp. 424–425) has remarked, "It is quite surprising and disappointing for me that almost 40 years after the establishment of the concept of risk aversion by Pratt and Arrow, our profession has not been able to

attain a consensus about the measurement of risk aversion." Indeed, "this confusion is viewed by some as so severe that the existing body of evidence is dismissed in favor of introspection" (Meyer and Meyer, 2005, p. 244).

In *Measuring Risk Aversion*, Donald J. Meyer and Jack Meyer point out a previously unrecognized reason for some of the discrepancy: differences in outcome variables, or the arguments of the utility function. Economic models have variously defined utility over wealth, income, consumption, rates of return, profit, payoffs from gambling, and other variables. Moreover, even when two studies appear to apply the same outcome variable, they often define it in different ways—wealth, for example, may be measured more or less broadly by the inclusion or exclusion of components such as human capital. With that as a background, the book's central theme is twofold: neither the magnitude of the risk aversion coefficient nor the slope of the risk aversion function can be compared directly across studies when outcome variables differ, but if there exists a known functional relationship between the outcome variables, then the corresponding risk aversion functions and coefficients can be made comparable. The authors conclude that when existing estimates of risk aversion are converted to a common basis, much of the reported discrepancy disappears.

After a brief overview in chapter 1, chapter 2 introduces the analysis. Arrow–Pratt wealth (W) is defined to include only those assets whose quantities can be freely adjusted—that is, wealth, which is fungible. Meyer and Meyer explain that the coefficients of absolute and relative risk aversion developed by Pratt (1964) and Arrow (1965) were originally based on this specific concept of wealth, but because W is difficult to measure in practice, most subsequent analyses have used other forms of wealth or nonwealth variables as arguments of the decision maker's utility function. The substitution of other outcome variables for Arrow–Pratt wealth can alter the size and shape of the risk aversion function. To illustrate, they compare relative risk aversion computations across outcome variables using three examples: wealth versus return on investment, taxable wealth versus nontaxable wealth, and wealth versus consumption. As a generalization, the authors show that if Y is a function of X and the utility of X is equal to the utility of Y so that $u(X) = v(Y(X))$, then by simply taking first and second derivatives, the relative risk aversion for X, denoted $R_u(X)$, and relative risk aversion for Y, denoted $R_v(Y)$, can be expressed in terms of one another in the following predictable way:

$$R_u(X) = \frac{-Xu''(X)}{u'(X)} = \frac{-Yv''(Y)}{v'(Y)}\left[\frac{Y'X}{Y}\right] - \frac{Y''X}{Y'} = R_v(Y)\left[\frac{Y'X}{Y}\right] - \frac{Y''X}{Y'}. \qquad (1)$$

Equation (1) states the basic result in its most general form. Unless Y is strictly proportional to X (so that $Y'X = Y$ and $Y'' = 0$), $R_u(X)$ and $R_v(Y)$ will differ in both magnitude and rate of change, yet if the function relating Y to X is known, $R_v(Y)$ can be converted into $R_u(X)$. In principle, this should allow the coefficient of relative risk aversion taken from a study in which utility is defined over consumption (or any other outcome variable) to be rescaled and compared directly to the coefficients of relative risk aversion from unrelated studies that use income or wealth as the argument of the (indirect) utility function. The trick is to find the functional relationship between the two original outcome variables. To that end, much of the subsequent analysis focuses on the special linear case, $Y = a + bX$, for which Equation (1) reduces to

$$R_u(X) = R_v(Y)[bX/Y]. \tag{1'}$$

If $a > 0$, then $b > 1$ implies that Y is more inclusive than X, and $0 < b < 1$ implies that Y is less risky than X. Thus, the book's central theorem states that for $a > 0$ and $b > 0$, "the relative risk aversion measure associated with an outcome variable is larger when the outcome variable is measured more inclusively or is less risky" (p. 29). In addition, differentiating (1′) with respect to X gives the slopes of the risk aversion functions,

$$R'_u(X) = \left[\frac{ab}{Y^2}\right] R_v(Y) + \left[\frac{b^2X}{Y}\right] R'_v(Y), \tag{2}$$

so that dividing (2) by (1′) and multiplying by X yields the following elasticities:

$$\frac{R'_u(X)X}{R_u(X)} = \frac{a}{Y} + \left[\frac{bX}{Y}\right]\left[\frac{R'_v(Y)Y}{R_v(Y)}\right]. \tag{3}$$

These results imply that "the slope and elasticity of the relative risk aversion measure for the more inclusive and less risky outcome variables tend to be smaller than those for the less inclusive and riskier outcome variables" (p. 30).

The theorem provides an intriguing and useful insight for cases in which outcome variables can be linked in this manner. Of course, not all outcome variables that might be of interest necessarily have either linear or positive relationships, especially when the relationships are behavioral. Consumption may not be linear in income, for example, and simple labor supply models suggest that an increase in exogenous wealth tends to reduce an individual's optimal hours of work, reducing earned income for any given hourly wage. Thus, the assumption of a positive linear relationship between outcome variables will not always hold, as the authors recognize.

The next four chapters apply the theorem to specific cases: utility defined over various forms of wealth, consumption, profit, and gambling payoffs, respectively. In doing so, each of these chapters provides a short survey of both the theoretical and empirical literatures.

Chapter 3 is divided into three sections. It begins by reviewing the seminal work of Pratt (1964) and Arrow (1965), and summarizing subsequent theoretical papers. The second section analyzes how adding a component to Arrow–Pratt (A–P) wealth alters the coefficient of relative risk aversion, and the third section reviews empirical research. The authors conclude that *theoretical* work endorses both $R_u(W) < 1$ and IRRA for A–P wealth. Of course, *empirical* studies often find relative risk aversion to be well in excess of 1, and many find it to be decreasing in wealth. Meyer and Meyer offer a novel approach to reconciling the empirical evidence with theory. In particular, they demonstrate that for any component X of a wealth measure broader than W, "Whether X is fixed or random and independently distributed, as long as X is positive, the more comprehensive measure of wealth has the larger measure of relative risk aversion and tends to be more negatively sloped and has a smaller elasticity. When X is perfectly correlated with W, what is important to know is whether this added component is more or less risky than W. If less risky, the magnitude of relative risk

aversion is increased and the slope and elasticity decreased. The opposite occurs when the added component is riskier than W" (p. 50). Thus, empirical evidence that relative risk aversion exceeds unity and/or decreases as wealth rises can be attributed to the inclusion of wealth components that are fixed, distributed independently of W, or correlated with W but less risky. On the other hand, Morin and Suarez (1983) found that excluding housing from wealth generated DRRA, but including it produced IRRA; such results are also consistent with the Meyer and Meyer analysis if housing is considered to be a riskier asset than A–P wealth.[1]

The chapter could have been further strengthened by some quantification of how much impact the use of broader wealth measures has on the magnitude of relative risk aversion. There are two issues here: the extent to which various empirical estimates can be reconciled with each other, and the extent to which they can be reconciled with economic theory. The book does not undertake such an analysis, though the authors do refer the reader to their earlier article on the same topic, Meyer and Meyer (2005), which uses Friend and Blume (1975) data to illustrate the theorem. A statistical investigation using other estimates of risk aversion taken from the empirical literature would make an interesting topic for future research.

Because the focus is on $R_u(W)$, one of the chapter's few comments regarding absolute risk aversion is that "in . . . theoretical work where A–P wealth W is the outcome variable . . . the assumption of decreasing absolute risk [aversion] for W has very strong support, and increasing absolute risk aversion is quite universally rejected" (p. 43). This interpretation of the literature is perhaps a bit too strong; Eichner (2000), for example, has shown that utility functions exhibiting increasing absolute risk aversion (IARA) can generate plausible results in both insurance problems and redistributive taxation problems. Moreover, there is an empirical literature that finds evidence of IARA, at least over some intervals, in a variety of settings.[2] Whether these findings would translate into IARA under A–P wealth is another question.

Chapters 4 and 5, covering consumption and profit, respectively, can be read together, particularly inasmuch as the authors observe, "The outcome variable referred to as profit or net income is more similar to consumption than wealth" (p. 80). Indeed, chapter 4 proposes a consumption function of the form $C = a + bW$, while chapter 5 uses $C = Y + bW$, with Y now defined as income. The assumption of linearity allows the authors to apply their basic theorem to cases involving utility defined over consumption, and to conclude, "Small and relatively constant values for relative risk aversion for wealth are consistent with high and severely declining measures of relative risk aversion for consumption" (p. 71). Similarly, "relative risk aversion for net income should always be less than that for consumption" (p. 82). The two chapters demonstrate that these propositions are consonant with much of the empirical literature, though the authors have also been careful to acknowledge that the use of

[1] Additional evidence could be cited as well. Siegel and Hoban (1982), for example, also found DRRA for narrowly defined wealth and IRRA for a broader wealth measure that included nonmarketable and contingent (i.e., uncertain) assets.

[2] This research encompasses such areas as investment behavior (Jackwerth, 2000), agricultural production and storage (Abdulkadri et al., 2003), insurance purchasing (Eisenhauer, 1997; Eisenhauer and Halek, 1998), and gambling (Hamid et al., 1996).

the linear consumption function and the magnitudes of its parameters are subject to debate.[3]

Although the focus of the book is on differences in outcome variables, a careful reading reveals that another important reason for the discrepancies in the literature is the variation in modeling assumptions. The authors point out, for example, that relative risk aversion for consumption is often reported to be extremely large, but that such estimates result from the imposition of utility functions exhibiting constant relative risk aversion; when relative risk aversion is permitted to decrease in consumption, the magnitude of the estimates can be reduced. Some studies therefore adopt more flexible functional forms for utility, so that the slope of relative risk aversion is not predetermined. However, even these functions are sometimes not neutral; as Meyer and Meyer note, "a weakness in [the] expo-power utility function . . . biases the findings toward increasing relative risk aversion" (p. 77). Still other studies impose no specific functional form on utility at all, but instead construct a Taylor series expansion of utility (or expected utility) in order to obtain a reduced-form expression for relative risk aversion, and then insert data to estimate the magnitude of risk aversion and determine whether it is increasing, constant, or decreasing in its argument. Even the results from that approach are sensitive to the point at which the approximation is made and the number of terms used in the expansion (Eisenhauer, 2003). It seems likely that a substantial portion of the dissonance in reported results may be attributable to such differences in models, aside from differences in outcome variables.

Additional explanations are also possible. In chapter 6, Meyer and Meyer discuss risk aversion for other outcome variables—gambling payoffs in particular—while providing a nice review of some recent studies involving experiments and television game shows. They observe that in many studies, aversion to risk is negligible or nonexistent when the payoff is low, and it increases with the magnitude of the payoff. This is consistent with the Meyer and Meyer theorem to the extent that gambling payoffs are less inclusive than A–P wealth, and a larger expected value of the payoff (for a given W) brings relative risk aversion for the payoff closer to $R_u(W)$. There is, however, a second and subtler issue that deserves attention: a mean-preserving increase in the variance of the payoff also affects risk aversion, and as the interval between potential gains and losses expands, differential calculus becomes a less precise tool for measuring changes in utility. This raises the question of whether the conventional measures of risk aversion are appropriate for both large and small risks (Eisenhauer, 2006).

A final and perhaps obvious explanation for some of the empirical discrepancy is simply that human beings differ from one another in their tolerances for risk. Meyer and Meyer mention briefly that some researchers have investigated demographic differences in risk aversion, though a more extensive discussion of this point may be warranted. Since preferences frequently vary across demographic groups, a study based exclusively on one group is likely to obtain different results from a study based on

[3] Indeed, in the chapter 5 version, for any positive value of b, the individual is assumed to spend all of his income and some of his wealth on consumption in each period. In some cases it might be preferable to assume that $C = aY + bW$ with both $0 < a < 1$ and $0 < b < 1$, making it possible for the individual to either accumulate or deplete wealth, depending upon whether consumption is less than or greater than income.

another. As Meyer and Meyer note, for example, some studies have relied exclusively on game show participants or convenience samples of volunteer college students as experimental subjects, raising the question of whether their attitudes and behavior can be reliably generalized to the larger population. Indeed, given the wide diversity of human personalities, it is not clear that one should expect much similarity in risk preferences across studies that investigate the behavior of different populations.

Even so, Meyer and Meyer make a convincing argument that the underlying similarity is at least somewhat greater than it would appear at first glance. On the whole, the book is nicely written, and clear enough to assign to advanced undergraduates as a supplementary text. Its central theme is certainly correct, and in retrospect, should have been apparent to researchers in the field from the start. The real mystery, it seems, is not why so many different empirical results have been obtained regarding risk aversion, but why it took us so long as a profession to recognize precisely how differences in outcome variables affect the measurement of risk aversion.

REFERENCES

Abdulkadri, A., M. R. Langemeier, and A. M. Featherstone, 2003, Estimating Risk Aversion Coefficients for Dryland Wheat, Irrigated Corn, and Dairy Producers in Kansas, *Applied Economics*, 35: 825-834.

Arrow, K. J., 1965, *Aspects of the Theory of Risk Bearing* (Helsinki: Yrjo Jahnsson).

Blake, D., 1996, Efficiency, Risk Aversion, and Portfolio Insurance: An Analysis of Financial Asset Portfolios Held by Investors in the United Kingdom, *Economic Journal*, 106: 1175-1192.

Eichner, T., 2000, A Note on Indifference Curves in the (μ, σ)-Space, *OR Spektrum*, 22: 491-499.

Eisenhauer, J. G., 1997, Risk Aversion, Wealth, and the DARA Hypothesis: A New Test, *International Advances in Economic Research*, 3: 46-53.

Eisenhauer, J. G., 2003, Approximation Bias in Estimating Risk Aversion, *Economics Bulletin*, 4: 1-10.

Eisenhauer, J. G., 2006, Risk Aversion and Prudence in the Large, *Research in Economics*, 60: 179-187.

Eisenhauer, J. G., and M. Halek, 1998, Prudence, Risk Aversion, and the Demand for Life Insurance, *Applied Economics Letters*, 6: 239-242.

Eisenhauer, J. G., and L. Ventura, 2003, Survey Measures of Risk Aversion and Prudence, *Applied Economics*, 35: 1477-1484.

Friend, I., and M. E. Blume, 1975, The Demand for Risky Assets, *American Economic Review*, 65: 900-922.

Gollier, C., 2001, *The Economics of Risk and Time* (Cambridge, MA: MIT Press).

Halek, M., and J. G. Eisenhauer, 2001, Demography of Risk Aversion, *Journal of Risk and Insurance*, 68: 1-24.

Hamid, S. S., A. J. Prakash, and M. W. Smyser, 1996, Marginal Risk Aversion in a Betting Market, *Applied Economics*, 28: 371-376.

Hansen, L. P., and K. J. Singleton, 1982, Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models, *Econometrica*, 50: 1269-1286.

Jackwerth, J. C., 2000, Recovering Risk Aversion from Option Prices and Realized Returns, *Review of Financial Studies*, 13: 433-451.

Kocherlakota, N. R., 1996, The Equity Premium: It's Still a Puzzle, *Journal of Economic Literature*, 34: 42-71.

Levy, H. 1994, Absolute and Relative Risk Aversion: An Experimental Study, *Journal of Risk and Uncertainty*, 8: 289-307.

Mehra, R., and E. C. Prescott, 1985, The Equity Premium: A Puzzle, *Journal of Monetary Economics*, 15: 145-161.

Meyer, D. J., and J. Meyer, 2005, Relative Risk Aversion: What Do We Know?, *Journal of Risk and Uncertainty*, 31: 243-262.

Morin, R. A., and A. F. Suarez, 1983, Risk Aversion Revisited, *Journal of Finance*, 38: 1201-1216.

Pratt, J. W., 1964, Risk Aversion in the Small and in the Large, *Econometrica*, 32: 122-136.

Siegel, F. W., and J. P. Hoban, Jr., 1982, Relative Risk Aversion Revisited, *Review of Economics and Statistics*, 64: 481-487.

Szpiro, G. G., 1986a, Measuring Risk Aversion: An Alternative Approach, *Review of Economics and Statistics*, 68: 156-159.

Szpiro, G. G., 1986b, Relative Risk Aversion Around the World, *Economics Letters*, 20: 19-21.

Social Security Programs and Retirement Around the World: Micro-Estimation, edited by Jonathan Gruber and David A. Wise, 2004, Chicago, IL: The University of Chicago Press

Reviewer: J. Tim Query, Associate Professor, Mountain States Insurance Group Chair, College of Business, New Mexico State University

Social Security Programs and Retirement Around the World: Micro-Estimation is the second volume of an ongoing research project studying the relationship between social security and labor supply. In the first volume, Gruber and Wise reveal strong disincentives to continue working at older ages in developed countries. In many social security programs, provisions are revealed that essentially encourage retirement by reducing total compensation for working. The second volume uses microdata to further analyze retirement behavior on a country-by-country basis. The results in this volume provide an important complement to the first volume in that the results leave little doubt that social security incentives are an important factor in the retirement decision-making process.

Around the world, a major demographic tide of declining birthrates is pushing nations further away from the promises that they have made to seniors. Additional salient issues driving the need to reform social insurance include more women in the global workplace, rising divorce rates, changing employment patterns, and rising budget deficits. The combination of changing employment life-cycle behavior and increasing average age of populations in various developed countries, coupled with social security stipulations that have contributed to the decline in the labor force participation of older persons, can substantially reduce the productive capacity of the

labor force. While most of the political discussion, at least in the United States, has centered on "preserving" social security benefits, less attention has been given to the incentive/disincentive employment and early retirement effects of various proposals. The Gruber and Wise book addresses this issue using a logical and salient process.

The social security programs in most industrialized countries are financed on a pay-as-you-go basis. Under this arrangement, most countries have accumulated large unfunded liabilities and face looming financial burdens. The penalty on work induces older employees to leave the labor force early and thus magnifies the financial burden caused by population aging.

While the first stage of the Gruber and Wise project produced two major results—(1) social security systems in many countries are laden with strong incentives to leave the labor force at older ages, and (2) a positive relationship between these incentives to retire and the departure of senior workers from the work force—it did not provide a method for estimating the magnitude of plan provision changes on labor force participation.

Social Security Programs and Retirement Around the World uses substantial databases of individuals compiled by research teams in 12 different countries. The data in each country match information on retirement decisions to the retirement incentives inherent in the social security provisions of each country. Across these 12 countries, noteworthy in their characteristically diverse social security programs and labor market institutions, the results consistently support the contention that program incentives are significantly aligned with retirement decisions.

According to the Gruber and Wise, the key advantage of the microestimation approach is that in each country the effects of changes in social security stipulations on retirement decisions can be predicted. Each paper illustrates the country-specific effects of two different reforms. One simulated reform assumes provisions that are common in each of the countries of interest to the study—reducing retirement incentives in some countries and increasing incentives in others. The other illustrative reform simulates delays in the benefit eligibility ages in each of the 12 countries.

The individual countries studied for this volume and the respective authors include Belgium (Dellis, Desmet, Jousten, and Perelman), Canada (Baker, Gruber, and Milligan), Denmark (Bingley, Gupta, and Pedersen), France (Mahieu and Blanchet), Germany (Supan, Schnabel, Kohnz, and Mastrobuoni), Italy (Brugiavini and Peracchi), Japan (Oshio and Oishi), Netherlands (Vos and Kapteyn), Spain (Boldrin, Martin, and Peracchi), Sweden (Palme and Svensson), United Kingdom (Blundell, Meghir, and Smith), and the United States (Coile and Gruber). The main purpose of each paper is to analyze empirically the impact of social security incentives on retirement decisions of older employees. As expected, there are notable differences among the papers due to a variety of factors, including: methodological differences, availability and structure of data, and the divergent approaches undertaken by authors from all over the world. However, country-specific issues are often discussed as part of the country-specific analysis where appropriate.

The first chapter of this book provides a somewhat detailed "executive summary" of the research project. This includes a description of the incentive measures used in the analysis, along with other model specifications. A synopsis of the parameter estimates

obtained across the 12 countries is also provided. The key empirical consistency found in the analyses, the correspondence between retirement decisions and benefit eligibility ages, is also discussed. The concluding section of the first chapter describes the simulations utilized to summarize the implication of the estimated models, with a subsequent discussion of simulation results.

The results of the country analyses reported in this volume serve to confirm the strong relationship of social security program retirement incentives on labor force participation. Perhaps of even greater interest from a public policy perspective is the magnitude of these effects. The 12 countries of interest to this study provide a diverse sampling of social security programs as well as different labor market systems. The results reported by all of the country studies consistently support the contention that program incentives concur significantly with retirement decisions. The magnitude of the estimated effects varies from country to country. However, the effects are considerable in all countries.

This book adds to the discussion of social security reform, and should be included among the literature addressing social security systems in various countries, and more generally the application of labor economics to social insurance. On a practical basis, this book could conceivably be added as supplemental reading to a graduate- or doctoral-level course examining the structural implications of various proposed changes to the social security system. In general, the collection of research papers provides ample support for the hypotheses developed in the first volume. There are a myriad of issues currently under review in a number of countries involving social security reform, and this collection of research studies adds an important perspective to the "work or retire" decision aspect of these provisions.