



1-1-2000

Income Shocks, Asset Returns, and Portfolio Choice

Steven J. Davis
University of Chicago

Paul Willen
University of Chicago

Follow this and additional works at: https://repository.upenn.edu/prc_papers

 Part of the [Economics Commons](#)

Davis, Steven J. and Willen, Paul, "Income Shocks, Asset Returns, and Portfolio Choice" (2000). *Wharton Pension Research Council Working Papers*. 482.
https://repository.upenn.edu/prc_papers/482

The published version of this Working Paper may be found in the 2002 publication: *Innovations in Retirement Financing*.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/prc_papers/482
For more information, please contact repository@pobox.upenn.edu.

Income Shocks, Asset Returns, and Portfolio Choice

Disciplines

Economics

Comments

The published version of this Working Paper may be found in the 2002 publication: *Innovations in Retirement Financing*.

Innovations in Retirement Financing

Edited by Olivia S. Mitchell, Zvi Bodie, P. Brett Hammond, and Stephen Zeldes

Pension Research Council
The Wharton School of the University of Pennsylvania

PENN

University of Pennsylvania Press

Philadelphia

Copyright © 2002 The Pension Research Council of The Wharton School of the University of Pennsylvania

All rights reserved

Printed in the United States of America on acid-free paper

10 9 8 7 6 5 4 3 2 1

Published by

University of Pennsylvania Press

Philadelphia, Pennsylvania 19104-4011

Library of Congress Cataloging-in-Publication Data

Innovations in retirement financing / edited by Olivia S. Mitchell . . . [et al.].

p. cm.

Includes bibliographical references (p.) and index.

"Pension Research Council publications"

ISBN 0-8122-3641-6 (alk. paper)

1. Retirement income — planning. 2. Finance, personal. I. Mitchell, Olivia S. II. Wharton School. Pension Research Council.

HG179.I4866 2002

332.024'01 — dc21

2001053386

Chapter 2

Income Shocks, Asset Returns, and Portfolio Choice

Steven J. Davis and Paul Willen

Other chapters in this book stress the importance of portfolio allocation decisions in managing the financial risks associated with asset accumulation and decumulation for retirement purposes (Leibowitz et al., this volume; Bernheim et al., this volume). In this chapter, we develop and apply a simple graphical approach to portfolio selection that accounts for correlation between asset returns and an investor's labor income.¹ Our approach easily handles the realistic case in which income shocks are partly, but not fully, hedgeable.²

We first show how the properties of labor income shocks and their correlation with asset returns affect portfolio choice. Next, we estimate the properties of shocks to the occupation-level components of individual income and investigate their correlations with aggregate equity and bond returns, selected industry-level equity returns, and the returns on portfolios formed on firm size and book-to-market equity values. We then use the theoretical framework and empirical results to calculate optimal portfolio allocations over the life cycle for workers in selected occupations.

Our analysis captures several important factors that influence portfolio choice over the life cycle: the drawdown of human capital as a worker ages, the impact of labor income innovations on the present value of lifetime resources, the increase in an investor's effective risk aversion as income smoothing ability declines with age, and systematic life cycle variation in the correlation between labor income shocks and asset returns. Each of these factors affects an investor's optimal level of risky asset holdings, as we show below.

According to the two-fund separation principle of traditional mean-variance portfolio analysis, all investors should hold risky financial assets in the same proportions, with only the level of holdings differing across people. We show why and how this simple prescription breaks down when an investor has a risky income stream (from work or business ownership) that is

correlated with asset returns. We quantify this breakdown and several contributory factors, and we show that even moderate correlations between income shocks and asset returns can drive large differences between optimal portfolio shares and the shares implied by a more traditional approach that ignores risky labor income.

Portfolio Choice with Risky Labor Income

If an investor can only borrow and lend but cannot invest in risky assets, her consumption is limited by the sum of her initial risk-free asset holdings (e.g., government bonds) and the present discounted value of her current plus future labor income. Now suppose she can also invest in a risky asset that offers an expected rate of return greater than the risk-free interest rate. If she borrows a dollar and invests it in the risky asset, then her expected future income increases by the difference between the expected return on the risky asset and the amount she has to pay back on the loan—in other words, by the excess return on the risky asset. The same point holds if she finances the investment in the risky asset by drawing down her initial position in the risk-free asset. Either way, this increase in expected income comes at a cost, because the riskiness of future consumption also rises.

This tradeoff between higher risk and higher return is well known in finance and economics. Based on various characteristics such as age, wealth, and risk aversion, investors may be more or less willing to take on risk in order to increase their expected level of consumption.

Age, Wealth, Risk Aversion, and Portfolio Choice

Let PVL_R stand for the expected present value of lifetime earnings from working plus the present value of lifetime excess returns on risky assets. Investing in risky assets increases PVL_R and thus the expected amount an investor can consume over her life cycle. In this sense, investing in risky assets increases wealth. The solid line in Figure 1 shows this wealth measure as a function of investment in a risky asset with an excess return of 8 percentage points.

But every dollar of investment in risky assets also increases the risk of her consumption, as mentioned above. The dashed line in the top panel of Figure 1 shows the cost of increased consumption variability as a function of investment in a risky asset with a standard deviation of 15 percent. We measure the cost as the amount of wealth that the investor would forgo to eliminate the consumption variability caused by the additional risky asset holdings. Suppose, for example, that you invest \$10,000 in risky assets, and this increases the variability of your future consumption by 10 percent. To measure the cost of the added risk, we ask how much additional income you would require to compensate for the 10 percent increase in consumption variability. Note that the slope of the cost curve increases with investment—

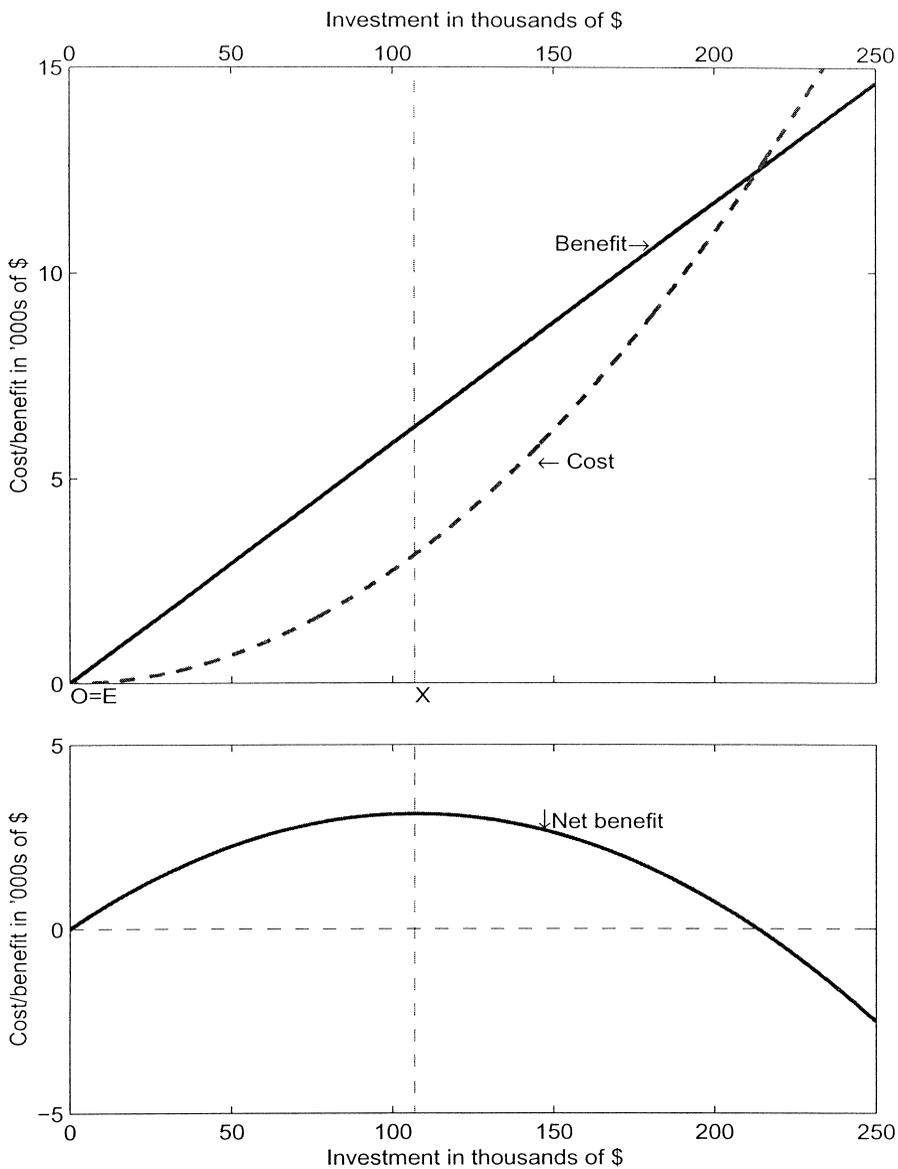


Figure 1. Portfolio choice.

in contrast to the slope of the benefit curve, which is constant. The more of a particular risk an investor takes on, the less willing she becomes to take on yet more of that same risk.³

Given this tradeoff between risk and return, how should an investor choose the optimal level of the risky asset? To answer this question, observe that the net benefit of the risky asset equals the difference between the solid and dotted lines — that is, between the benefit of higher expected consumption and the cost of higher risk. The bottom panel of Figure 1 shows the net difference between cost and benefit as a function of the amount invested. To maximize utility, an investor chooses the amount that maximizes the net benefit, point *X* in Figure 1. We call the distance from point *O* to point *X* an investor's "desired exposure" to the risky asset.

Investors differ by age, wealth level, risk aversion, occupation, region, and many other characteristics. How do these things affect portfolio choice? They do not affect the benefit of risky asset investment — a dollar investment in a risky asset increases PVL_R by the same amount for all people who face the same risk-free interest rate. They typically do affect the cost to an investor of taking on the type of risk implied by investments in the risky asset. To develop this point, we first discuss how age, wealth, and risk aversion affect the cost of risky asset holdings. We then consider the role of labor income uncertainty, especially as it relates to an investor's occupation.

Risk aversion. Some people find risk highly unpleasant, so that they are willing to forego a relatively large amount of wealth to eliminate a given amount of risk. Hence, the cost curve is higher and steeper for investors with greater risk aversion (see Figure 2). As a consequence, desired exposure is lower, other things equal.

Age. Younger investors have more time to smooth income or wealth shocks, so the cost of income variability is lower for younger persons. For example, a dollar shock to wealth for an investor who only expects to live for another year results in a dollar shock to consumption. By contrast, a dollar shock to wealth for an investor who expects to live for another forty years results in a very small shock to consumption. Thus for old people the cost curve (the dashed line) is higher and "desired exposure" is lower. A shorter planning horizon, because of more advanced age or other reasons, affects portfolio choice in much the same way as greater risk aversion.

Wealth. The more wealth you have, the less overall effect a dollar shock to your wealth has on your consumption. If you lose \$50,000 on the stock market and you have a net worth of \$100 million, that is not so bad. But if you lose \$50,000 on the stock market and you have a net worth of \$50,000, that is a disaster. Thus the less wealth you have, the higher and more steeply sloped is the cost curve and the lower your desired exposure. Increased wealth affects portfolio choice in a similar way to decreased risk aversion.

Figure 2 shows portfolio choice for two investors, one with high risk

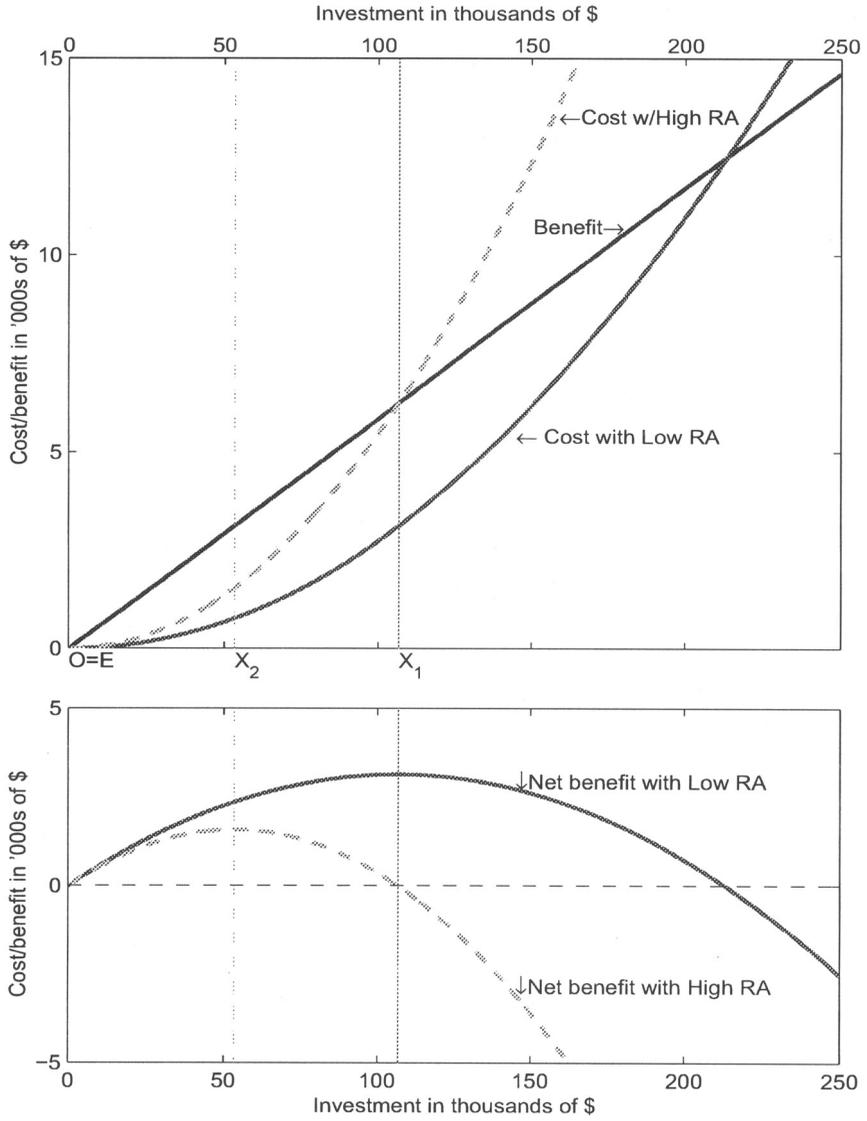


Figure 2. Portfolio choice with high and low risk aversion (RA).

aversion and the other with low risk aversion. The picture would look the same if we compared high age and low age or low wealth and high wealth.

Risky Labor Income

The focus of this chapter is on risky labor income, particularly labor income risk tied to a worker's occupation. How does the presence of risky labor income affect portfolio choice? Unlike risk aversion, age, and wealth, risky labor income does not change the shape of the cost curve. Rather, risky labor income changes its location.

Consider the following example. Suppose an investor works for Ford. Suppose she knows that the price of Honda stock is negatively correlated with her labor income. When things go well for Honda — sales of Odyssey minivans increase, for example — Ford sales drop and bonuses shrink. When things go poorly for Honda — the original Odyssey minivan was something of a dud — Ford sales increase and bonuses increase. Suppose she invests some money in Honda. When things go poorly at Ford (and well at Honda), her wealth and consumption fall by less than they would without the investment in Honda. When things go well at Ford (and badly at Honda), her wealth and consumption increase by less than they would without the investment in Honda. What has happened here? For the Ford employee, an investment in the risky Honda asset actually *reduces* the variability of her consumption! Figure 3 illustrates this graphically. The benefit of investing in a risky asset is still the same as before, but now the cost curve initially slopes down as she invests.

In the preceding example, our investor can increase her wealth and lower her risk at the same time. This seems like a free lunch, and it is, but only up to a point. Since her exposure to Ford risk is fixed (determined by her employment situation) the effects of added exposure to Honda risk eventually swamp the reduction in Ford risk as she adopts larger positions in the risky Honda asset. Her net benefit is maximized at point *X*. To sum up, negative correlation between returns and labor income shifts the cost curve down and to the right.

One can make a related argument for an investor who has labor income that is positively correlated with the returns on a risky asset. For example, a Honda employee may find that, by taking a short position in Honda stock, she can reduce the variability of her total income. But for any long position the cost is higher than it would be if she had no risky labor income at all, or risky labor income that was uncorrelated with the returns on Honda stock. Figure 4 illustrates this situation graphically. By taking a short position — that is, moving to the left of point *O*, — our investor sees her risk fall. But in contrast to the case of negative correlation in Figure 3, now her benefit falls as well. And for any positive investment the cost curve will be higher than it would be in the absence of labor income risk.

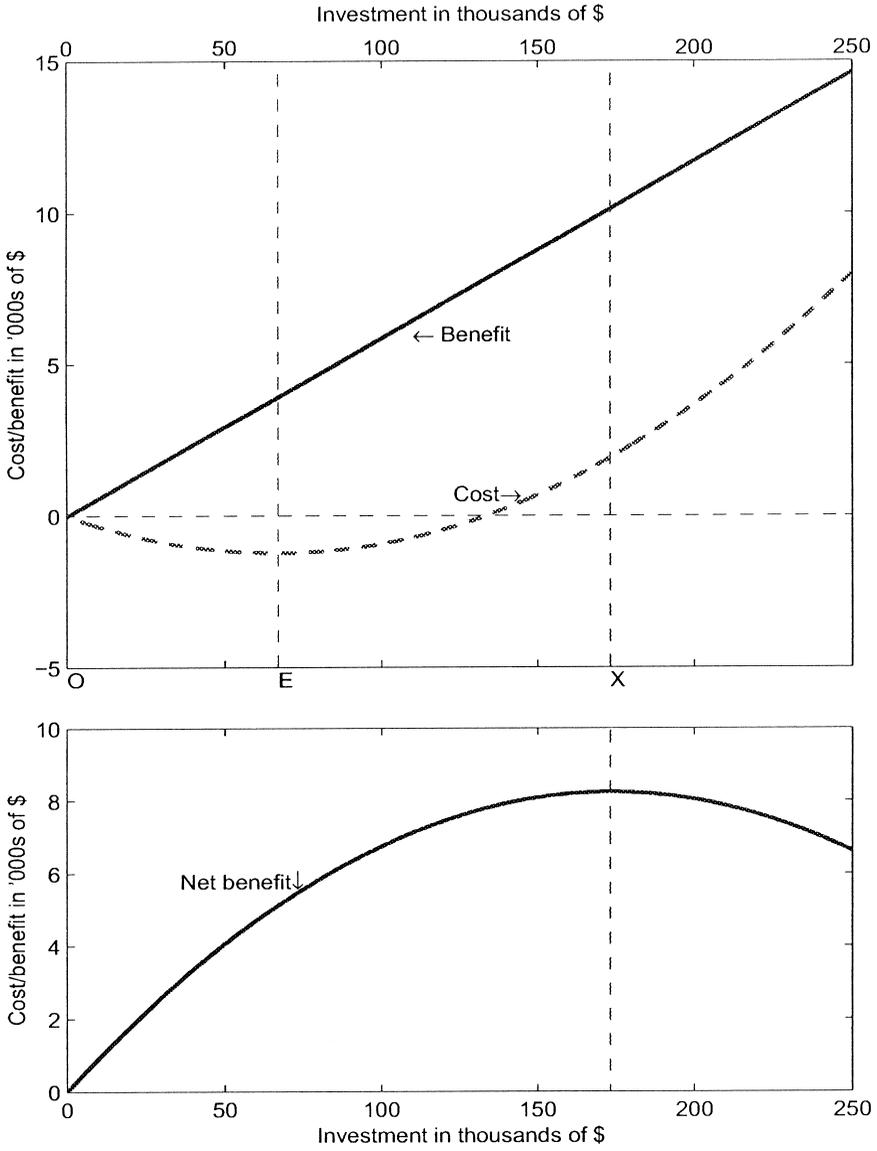


Figure 3. Portfolio choice with negative correlation between asset returns and labor income.

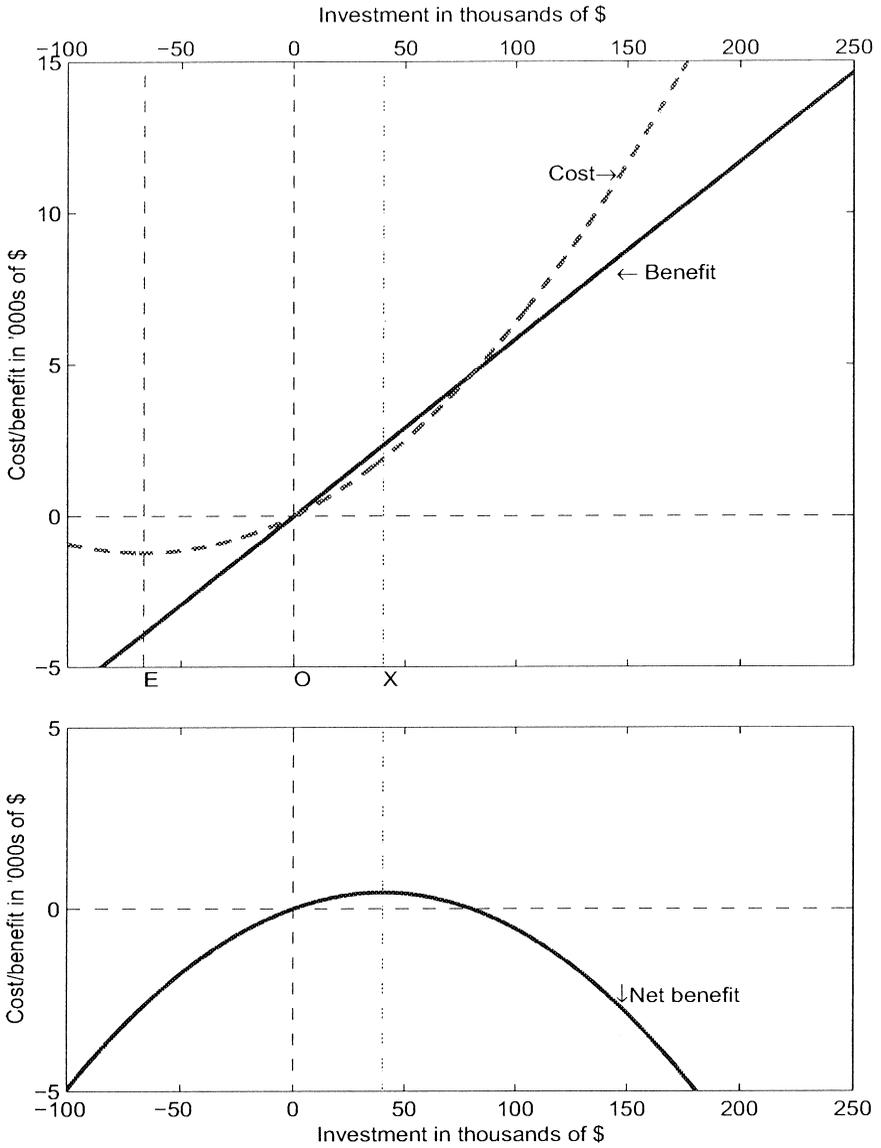


Figure 4. Portfolio choice with positive correlation between asset returns and labor income.

TABLE 1. Summary of Portfolio Choice Decision

<i>Distance</i>		<i>Name</i>	<i>Description</i>
<i>From</i>	<i>To</i>		
<i>E</i>	<i>X</i>	Desired exposure	Quantity invested in the risky asset in the absence of any labor income risk
<i>O</i>	<i>E</i>	Endowed exposure	Quantity invested to minimize variability of consumption
<i>O</i>	<i>X</i>	Optimal portfolio choice	Desired Exposure minus Endowed Exposure

Why is this so? Labor income for a Honda employee is (by assumption) positively correlated with returns on Honda stock, so our investor already has exposure to the stock, even if she owns none of it. Hence, she reaches the point of maximum net benefit from Honda stock much sooner than she would if she had no risky labor income or worked for Ford. To sum up, positive correlation between returns and labor income shifts the cost curve down and to the left.

How can we relate this to our earlier discussion of age, wealth, and so forth? The key point here is that age, wealth, and risk aversion affect the shape of the cost curve. In contrast, the risk characteristics of labor income move the whole curve while preserving its shape. The invariance of the shape of the cost curve to the risk characteristics of labor income provides a convenient decomposition for portfolio choice analysis. Since the shape of the cost curve doesn't change with occupation, the distance from the optimal portfolio choice point X to the minimum cost point E is also unaffected. We call this distance the desired exposure. It is equivalent to the optimal investment level for an investor whose minimum cost point coincides with zero investment in the risky asset. We call the distance from zero investment O to the minimum variance point E "endowed exposure." Endowed exposure reflects the risk characteristics of the investor's labor income, and it also measures the level of disinvestment required to minimize risk.

The optimal portfolio choice is simply desired exposure minus endowed exposure (the distance from O to X in the figures). The decomposition makes it possible to capture all the effects of occupation in one number, "endowed exposure," and to capture the effects of age, risk aversion, and wealth in another number, "desired exposure."⁴ Table 1 summarizes the relationship between desired exposure, endowed exposure, and optimal portfolio choice.

Previous research on portfolio choice has focused on the determinants of desired exposure. We focus on endowed exposure, which is ignored in most analyses of portfolio choice.

Measuring Endowed Exposure

The magnitude of an investor's endowed exposure to a risky asset depends on three things:

Correlation. As we have discussed already, the higher the correlation, the higher is endowed exposure.

Variability of income. The higher the variability of income, all else equal, the higher is endowed exposure. Why? Return to our example. If you work for Ford and your compensation is well insulated from the ups and downs of Ford's fortunes, then the effect of a shock to Ford on your income will be very small. Even if the correlation between Ford stock returns and your income is high, the potential for risk reduction is low, because you simply don't have very much risk to begin with.

Persistence. An investor's lifetime consumption possibilities, as we already noted, depend on the present value of her lifetime resources. A shock to labor income obviously affects current labor income, but current labor income is typically only a small part of the present value of lifetime resources. In general, a shock to labor income today also conveys information about future labor income, and thus may have a large impact on lifetime resources. Consider a Ford employee who gets a reduced bonus because of poor sales one year. From a life cycle perspective, the reduced income this year is not so important. And if she expects the company to rebound quickly, this shock will have little impact on her PVLR or her consumption. We call this a shock of *low persistence*. But if the reduced bonus presages major long-term cutbacks at the company, her lifetime resources may decline sharply—we call this a shock of *high persistence*. The more persistent the shock, the more the shock affects PVLR and consumption, and the higher the endowed exposure to a given shock.

Persistence also depends on the retirement horizon. If a worker is only one year from retirement, a shock to income can only affect current income—so all shocks are of low persistence. Because of this horizon effect, labor income shocks effectively become less persistent as an investor ages.

When we actually go to the data below we measure persistence as the shock to the present value of lifetime resources of a dollar shock to labor income this year. If a shock to labor income is highly persistent, then future labor income will be significantly affected and PVLR will change considerably—often by as much 20 times the shock to current income. By contrast, if a shock to labor income is of low persistence, then future labor income won't be greatly affected and PVLR will change by only a small amount.

To measure endowed exposure, therefore, we need to measure three things: the variability of income shocks, the persistence of those shocks, and the correlation of labor income shocks with stock returns or other risky assets under consideration.

Multiple Risky Assets

So far, we have considered an investor investing in only one risky financial asset. In reality, investors have tens of thousands of risky financial assets to choose from, from stocks in tiny start-up companies to mutual funds that try to mimic the entire universe of stocks in the U.S. How can we analyze these financial decisions? In the absence of labor income risk, this decision is actually remarkably simple. Consider two investors who have different levels of risk aversion and no labor income risk. Using methods described above, we can calculate their optimal portfolios for each asset separately. For each asset, the cost of risk curve will be higher for the more risk averse investor. Conveniently, the difference in costs across agents will be the same for each asset — which means that the relative investment in any two stocks will be the same — the ratio of shares held of Ford and Honda stock will be the same for all investors. This means that all investors will hold portfolios with identical weights on individual stocks and mutual funds. This drastic simplification of the portfolio problem is known as the principle of *two-fund separation*, because it says that a single mutual fund with the correct weights will enable all investors to implement optimal financial plans using just the mutual fund and a riskless bond.

When we add occupational or other sources of labor income risk, the principle of two-fund separation breaks down. Why? The key to two-fund separation is that risk aversion, age, and wealth affect demand for all assets by the same amount — they change the shape of the cost curve for every asset in the same way. But labor income risk changes the *location* of the cost curve for each investor in a potentially different way. To return to our example, for a Ford employee, labor income risk moves her Ford cost curve to the left and her Honda cost curve to the right — whereas for a Honda employee, labor income risk moves her Ford cost curve to the right and her Honda cost curve to the left. The ratio of holdings of Honda and Ford stock will no longer, in general, be the same for different investors.

Limitations of Our Approach

Our approach makes many assumptions about life cycle portfolio choice. It allows for unlimited borrowing of the riskless asset and unlimited short-sales of the risky asset. Short-sale restrictions on risky assets can be treated without great difficulty, but a proper treatment of borrowing constraints for the riskless asset complicates the analysis greatly in a many-period setting. Our approach also requires certain assumptions about the utility function and the time series processes for labor income and asset returns, which allow us to consider each period separately. In other words, investors do not need to worry about the effects of their current choices on their choice sets in future

periods. (A mathematical treatment of these issues appears in Davis and Willen 2000b.)

Occupation-Level Earnings Innovations

The chief empirical requirements of our approach are measures of (1) correlation between labor income shocks and asset returns, (2) variability of income shocks, and (3) persistence of income shocks. It is interesting that asset returns have received substantial attention from financial researchers, but only a handful of scholars have investigated their correlation with labor or proprietary business income. One such study, by Campbell et al. (1999), considers the correlation between aggregate equity returns and the permanent component of household income for three education groups. A second, by Davis and Willen (2000a), uses a synthetic panel to create demographic groups defined by sex, educational attainment, and birth cohort. Although they use rather different empirical designs, both of these studies find that the correlation between labor income shocks and equity returns rises with education. Heaton and Lucas (2000) emphasize a positive correlation between equity returns and the income of self-employed persons.⁵

Whereas prior studies have relied on panel data sets or synthetic panels constructed from repeated cross sections as the basis for analysis, here we pursue a somewhat different empirical approach. In particular, we rely on the repeated cross-section structure of the Current Population Survey to extract mean occupation-level income shocks, while controlling for a host of observable worker characteristics. We then focus our empirical investigation on the properties of the occupation-level shocks and their correlation with asset returns.

Data Sources and Definitions

The Current Population Survey (CPS) randomly samples about 60,000 U.S. households every month. Among other items, the survey inquires about labor earnings, employment status, hours worked, educational attainment, occupation, and demographic characteristics for each household member. The Annual Demographic Files in the March CPS contain individual data on these items for the previous calendar year. Using the CPS March files, we estimate occupation-level components of individual annual earnings from 1967 to 1994.

To compute annual earnings, we use CPS data on wage and salary workers in the private and public sectors who were 23 to 59 years old in the earnings year. Excluded from the earnings calculations are unincorporated self-employed persons, though we do include self-employment and farm income for persons who were mainly wage and salary workers. The sample is

TABLE 2. Occupational Classifications and Summary Statistics

<i>Occupational Description</i>	<i>1980 Standard Occupational Classification</i>	<i>Sample Period</i>	<i>Mean Cell Count</i>	<i>Minimum Cell Count</i>	<i>Average Earnings in 1982\$</i>
Accountants and Auditors	23	1967–94	542	327	24,881
Electrical Engineers	55	1967–94	246	150	33,923
Registered Nurses	95	1967–94	704	392	17,823
Teachers, Elementary	156	1967–94	842	679	18,325
Teachers, Secondary	157	1967–94	733	487	20,886
Janitors and Cleaners	453	1967–94	805	336	11,846
Auto Mechanics	505	1967–94	389	306	17,675
Electricians	575	1967–94	325	267	23,646
Plumbers	585	1967–94	220	168	22,437
Truck Drivers	804, 805	1967–94	1079	744	18,665

Source: Authors' tabulations from the Annual Demographic Files of the March CPS; see text.

Note: Average earnings is the simple mean from 1967 to 1994 of the unweighted mean annual earnings among persons who satisfy the selection criteria.

limited to those persons who worked at least 500 hours during the year, and people who were not students and not in the military.⁶ In addition to these individual-level selection criteria, we also limit attention to 57 detailed occupational classifications that can be tracked from 1967 (or 1970) to 1994. We further restrict our analysis to 10 detailed occupations with large numbers of individual-level observations in each year. These ten occupations differ widely in terms of educational requirements and annual labor income.⁷ These occupations and associated summary statistics on cell counts and average annual earnings in 1982 dollars appear in Table 2.⁸

The Occupation-Level Component of Earnings Innovations

To extract the occupation-level component of individual earnings shocks, we first fit ordinary least squares earnings regressions to the individual-level data. Separate models are fitted for each occupation after pooling the data over all available years. Explanatory variables include sex, educational attainment, age interacted with sex, and a full set of occupation-specific year effects. We estimate regressions using annual earnings as the dependent variable.⁹ The specification allows the age-earnings profile to vary freely across occupations (and sex), but not to shift over time. Effectively, we treat the occupation's average age-earnings profile over the 1967–94 period, adjusted for sex and education, as predictable variation in a worker's earnings. As implied by the occupation-level earnings specifications described below, we also treat the average occupational earnings growth from 1967 to 1994 (conditional on worker characteristics) as part of expected earnings growth.

To characterize the stochastic properties of the occupation-level compo-

ment of individual earnings shocks, we fit autoregressive moving average (ARMA) models to the first-differenced values of the occupation-year effects. We denote the occupation-year effects estimated in the first-stage earnings regressions as $\epsilon_{i,t}$, $t = 1967, 1968, \dots, 1994$. Then we fit second-order moving average processes, following MaCurdy (1982), who uses panel data on individuals, and Davis and Willen (2000a), who use synthetic panel data for demographic groups:

$$(1) \quad \Delta \epsilon_{i,t} = \alpha + \eta_t + \psi_1 \eta_{t-1} + \psi_2 \eta_{t-2} .$$

Here η_t denotes the time- t innovation to the occupation-level component of individual earnings shocks. These innovations and their correlation with asset returns are the main focus of the empirical investigation and the applied portfolio analysis below.¹⁰

Magnitude and Persistence of Earnings Innovations

The standard deviation of η_t in equation (1) quantifies the magnitude of innovations to the occupation-level component of individual earnings. As described earlier, we measure persistence as the shock to PVLR of a dollar shock to income. We refer to this persistence measure as the present value multiplier (PVM). The magnitude of the PVM depends on the persistence of η (a function of ψ_1 and ψ_2), the risk-free rate of interest, and the number of years remaining until retirement. By combining these elements, we can easily calculate the magnitude of a typical shock to PVLR at a given age. The magnitude of this shock declines with age, because fewer years remain until retirement.

Table 3 reports the present value multipliers on the occupation-level earnings shocks at ages 30 and 50, assuming a real discount rate of 2.5 percent per year and retirement after age 59. To illustrate the calculation of the shock to PVLR implied by an occupation-level income innovation, consider the example of the accountants and auditors occupation at age 30. According to Table 3, the standard deviation of innovations to the occupation-level component of earnings is \$1,080, or 4.3 percent of annual earnings. At age 30, the present value multiplier on this innovation is 20.0, so that the implied impact on PVLR amounts to $1080 (20.0) = \$21,600$. This figure is 87 percent of the average annual earnings for accountants and auditors reported in Table 2. These calculations show that occupation-level earnings innovations are of modest size, but the implied effects on the present value of lifetime earnings are not.

Occupations differ in terms of magnitude and persistence of occupation-level earnings innovations. The standard deviation of the occupation-level innovations ranges from 2.9 to 6.9 percent of annual earnings. Plumbers have the most volatile occupation-level earnings component in both dollar

TABLE 3. Selected Statistics for the Occupational Component of Individual Earnings, 1968–1994

<i>Occupational Description</i>	<i>Standard Deviation of Labor Income Shocks</i>	<i>Standard Deviation as % of Income</i>	<i>Present Value Multiplier at:</i>	
			<i>Age 30</i>	<i>Age 50</i>
Accountants and Auditors	1080	4.34	20.0	8.3
Electrical Engineers	1283	3.78	6.8	3.4
Registered Nurses	446	2.50	40.2	15.9
Elementary School Teachers	525	2.86	27.2	11.0
Secondary School Teachers	637	3.05	22.5	9.4
Janitors and Cleaners	583	4.92	13.3	5.8
Auto Mechanics	714	4.04	18.9	8.0
Electricians	951	4.02	13.2	6.1
Plumbers	1453	6.48	12.8	5.7
Truck Drivers	790	4.23	18.5	8.0

Source: Authors' calculations, CPS data.

Notes: For each occupation, a second-order moving average process is fit to the occupational component of individual annual earnings in 1982 dollars. The moving average process is estimated by (conditional) nonlinear least squares. See Davis and Willen (2000b) for an explanation of how the occupational component of individual earnings is identified. Present value multipliers computed using a real discount rate of 2.5 percent per year and assuming retirement after age 59.

and percentage terms, while registered nurses and elementary school teachers have the least volatile. Likewise, the present value multiplier at age 30 is 6.8 for plumbers and 40.2 for registered nurses. These two occupations are outliers in terms of persistence. For the other occupations, the present value multipliers at age 30 range from 13 to 27. The last two columns in Table 3 show how the present value multiplier declines between ages 30 and 50, given the assumptions about discounting and retirement. The age-50 multipliers are fairly sensitive to alternative assumptions about retirement age, but the basic point is not. As workers near retirement, earnings innovations have smaller and smaller effects on lifetime resources.

Correlation Between Occupation-Level Income Innovations and Asset Returns

To investigate the correlation between occupation-level earnings innovations and aggregate equity returns, we next regress η_t from equation (1) on the realized market rate of return during period t . Recall that we can use the slope coefficient in an ordinary least squares (OLS) regression of y on x to generate an estimate of the correlation of x with y . Hence, we can use standard regression methods to quantify the correlation between income shocks and equity returns and to test whether the relationship is statistically significant.

Correlation with Aggregate Equity Returns

We find little evidence that occupation-level income innovations and aggregate equity returns are linearly related in annual data over the period 1968 to 1994. None of the ten occupations considered evinces a statistically significant relationship between income innovations and returns on the value-weighted market portfolio (in regressions not detailed here).¹¹ As a check, we also considered the returns on several other broad-based equity indexes: the S&P 500, the New York Stock Exchange, the Wilshire 5000, and a value-weighted composite of the New York Stock Exchange, American Stock Exchange and NASDAQ. For each measure, we see the same pattern of little or no evidence for a relationship between occupation-level income innovations and contemporaneous aggregate equity returns.

This result is puzzling from the vantage point of standard economic theories of growth, fluctuations, and asset pricing. Equilibrium models that obey standard asset-pricing relationships and embed a conventional specification of the aggregate production technology imply a high positive correlation between aggregate equity returns and shocks to labor income.¹² While we note this puzzle here, it is not necessary to resolve it to pursue the remainder of our agenda.

Other Asset Return Measures

We also investigate the correlation between occupation-level income innovations and the returns on long-term government bonds and other assets. Bond returns are significantly correlated with income innovations for a few occupations. In most cases, bonds account for a greater fraction of occupation-level income innovations when the returns are measured in nominal terms.

We examine two additional types of assets which might be highly correlated with labor income shocks. First, we sought to construct industry equity portfolios that respond sensitively to labor income shocks in particular occupations. For example, demand shocks in the construction sector induce a positive correlation between equity returns in Construction industries (SICs 15, 16, 17) and occupation-level income innovations for Electrical Engineers, Electricians, and Plumbers. More generally, industry-level demand shocks and factor-neutral technology shocks impart a positive correlation between returns on industry equity and occupation-level income innovations.

However, prior reasoning alone cannot determine the sign, let alone the magnitude, of the correlation between industry equity returns and labor income innovations for industry workers. For example, labor-saving technological improvements in construction activity might be good for shareholders but bad for the earnings of Electricians and Plumbers. As another example, the deregulation of the trucking industry during the 1970s and early 1980s was bad news for many truck drivers (Rose, 1987) but good news for many trucking firms (Keeler, 1989). The basic point is that factor-biased technology shifts (construction example) and rent shifting between owners and workers (trucking example) impart a negative correlation between industry-level equity returns and occupation-level income innovations.

Clearly the relationship between industry-level equity portfolios and labor income shocks is very much an empirical issue. Furthermore, if the mix of underlying shocks and economic response mechanisms changes over time, the correlation between industry-level equity returns and occupation-level income innovations is likely to change. The weight of this concern is also largely an empirical issue. No single study can definitively settle these empirical issues, so our results in this regard are best viewed as one installment in a broader empirical inquiry.

We construct industry portfolios using firm-level equity returns and market values in the Center for Research in Security Prices (CRSP) database. For each occupation (except Janitors and Cleaners) we identify one or more industries that account for a large fraction of the occupation's employment. In some cases, we had to omit natural SIC counterparts for particular occupations, because CRSP contains no firm-level observations during part of the sample period.¹³ In the end, the SIC industry groups listed in Table 4 are

TABLE 4. Asset Return Measures, Definitions, and Summary Statistics

<i>Variable Name</i>	<i>Description</i>	<i>Mean Annual Return (%) 1968–94</i>	<i>Standard Deviation of Annual Return</i>	<i>Occupation Match</i>
SMB	Fama-French Size Portfolio, Small-Big	0.2	15.5	All
HML	Fama-French Book-to-Market Portfolio, Value — Growth Stocks	5.9	12.9	All
Bonds	Nominal Return on 10-Year Constant Maturity U.S. Government Bonds	8.5	10.1	All
Autos	Real Return on SIC 371 (Auto Mfg.)	6.4	25.0	Auto Mechanics
Elmach	Real Return on SIC 36 (Electrical Machinery Manufacturing)	5.8	21.4	Electrical Engineers
Build	Real Return on SICs 15, 16, 17 (Construction)	3.2	27.8	Elec. Eng., Electricians, Plumbers
Freight	Real Return on SIC 42 and 472, ex. 4725 (Freight Transport by Road)	6.4	27.8	Truck Drivers
Technical	Real Return on SICs 871 and 7336 (Engineering, Architectural and Technical Services)	8.1	31.9	Electrical Engineers
Education	Real Return on SICs 82, ex. 823, and 833 (Education Services)	6.4	37.1	Elementary and Secondary Teachers
Health	Real Return on SIC 80 (Medical, Dental and Health Services)	12.8	37.1	Registered Nurses
Utility	Real Return on SICs 46 and 49, ex. 495 (Electricity, Gas, Steam, Water Works)	5.4	15.8	Electrical Engineers, Electricians, Plumbers
Finance	Real Return on SICs 62, 67 (Investment Banking, Securities Markets, Exchanges)	7.9	19.8	Accountants and Auditors

Sources: Returns data for the SMB and HML portfolios taken from (web.mit.edu/kfrench/www/data.library.html); see also Fama and French (1993) for construction of these portfolios. Returns data on Bonds from Center for Research in Security Prices; all industry-level return series constructed from value-weighted portfolios of firm-level equity returns in the Center for Research in Security Prices database; see Davis and Willen (2000a).

Notes: Nominal returns for the industry-level measures converted to real returns using the GDP deflator for personal consumption expenditures. Data points missing for Health in 1968 and for Technical in 1987 and 1988. Last column lists the occupation for which the returns measure is used as a regressor.

targeted for further analysis. We construct value-weighted industry returns using firms in the CRSP data and update the firm-level weights annually. The rightmost column in Table 4 shows the occupations to which we match each industry-level return measure.

In a different approach, we consider the correlation between occupation-level income innovations and returns on equity portfolios formed on firm size (market equity value) and the ratio of book-to-market equity value (Fama and French 1993).¹⁴ The Fama-French SMB portfolio pays off the return on a portfolio of firms with small market values minus the return on a portfolio of firms with large market values. The Fama-French HML portfolio pays off the return on a portfolio of “value” stocks with a high ratio of book-to-market equity minus the return on a portfolio of “growth” stocks with a low ratio of book-to-market equity. The Fama-French portfolios are rebalanced quarterly and adjusted for transactions costs when firms are bought and sold. Prior research shows that size and book-to-market factors account for much of the cross-sectional variation in returns on common stocks (Fama and French 1992, 1993, 1996). Many other asset-pricing studies confirm an important role for these two factors.¹⁵

A question naturally arises as to what types of risk are being priced by size and book-to-market value. In other words, why do small cap stocks earn a higher average return than large cap stocks? And, why do value stocks earn a higher average return than growth stocks? One possibility is that shocks to labor income covary positively with the size and book-to-market factors. If so, then investors who are exposed to labor income risk will demand a return premium to hold small cap and value stocks. This asset-pricing logic suggests that labor income innovations might be correlated with the returns on the size or book-to-market portfolios. Following this logic, we investigate the correlation between occupation-level income innovations and returns on the SMB and HML portfolios.

Correlation with Other Asset Returns

We examine bivariate and multivariate regressions of occupation-level income innovations on returns for bonds, SMB and HML. For several occupations, the regression results show a fairly large negative correlation between income innovations and the SMB return.

The results in Table 5 suggest there is some scope for hedging occupation-level income risk, as suggested by the asset-pricing logic outlined above. However, the pattern of results runs directly counter to our original motivation for investigating the SMB portfolio. Most of the correlations in Table 5 and all the statistically significant ones imply that the relative return on small cap stocks covaries negatively with occupation-level income innovations. Thus investors who are exposed to labor income risk should be willing to hold small cap equities at a return discount relative to large cap equities. In

TABLE 5. Determinants of Endowed Exposure for SMB Portfolio, 1968–94

<i>Occupational Description</i>	<i>Correlation</i>	<i>Variability (std. dev.)</i>	<i>PVM at Age 30</i>
Accountants and Auditors	-0.37	1080	20.0
Electrical Engineers	-0.33	1283	6.8
Registered Nurses	-0.14	446	40.2
Teachers, Elementary	-0.36	525	27.2
Teachers, Secondary	-0.39	637	22.5
Janitors and Cleaners	-0.32	583	13.3
Auto Mechanics	-0.10	714	18.9
Electricians	0.22	951	13.2
Plumbers	-0.28	1453	12.8
Truck Drivers	0.03	790	18.5

Source: Authors' calculations; see text.

Notes: All regressions estimated by ordinary least squares. For regression results, see Davis and Willen (2000b).

fact, the average return on small cap stocks is higher.¹⁶ So, while the findings can be useful for portfolio design purposes, they serve to *heighten* rather than resolve asset-pricing puzzles related to the return premium on small cap stocks.

Life Cycle Portfolio Choice with Risky Labor Income: Some Examples

We now solve the life cycle portfolio problem with risky labor income, drawing on the just presented empirical evidence to characterize the magnitude, persistence, and correlation properties of labor income shocks.

Preliminaries and Two-Fund Separation

Optimal portfolio allocations when asset returns and labor income are uncorrelated appear in Table 6. The table considers three risky assets—market, size, and value portfolios—and uses a real risk-free return of 3.5 percent per year. We do not impose short-sale constraints on risky asset holdings or restrictions on borrowing at the risk-free rate. Since two-fund separation holds under these conditions, every investor has the same risky asset portfolio shares, as shown in the top row. These shares depend on the joint return distribution for the three assets, which we fit to the first two sample moments in the data. The table also displays optimal risky asset holdings at ages 40 and 60 for two occupations under various assumptions about relative risk aversion and expected returns.

We measure risk aversion using the Arrow-Pratt coefficient of relative risk aversion (RRA). To understand RRAs, consider the following experiment.

TABLE 6. Investment in Risky Assets with Zero Correlation between Earnings and Returns: Two-fund Separation

	<i>RRA</i>	<i>Age</i>	<i>Reduction in Returns (%)</i>	<i>SMB</i>	<i>HML</i>	<i>Market</i>	<i>Total</i>
<i>Portfolio Shares</i>				-25	88	37	100
<i>Asset Levels</i>	3	40	0	-257	903	381	1027
Electrical Engineers	5	40	0	-154	542	229	616
	3	60	0	-148	520	220	592
	3	40	50	-129	451	191	514
Secondary School Teachers	3	40	0	-158	556	235	632
	5	40	0	-95	334	141	379
	3	60	0	-91	320	135	364
	3	40	50	-79	278	117	316

Source: Authors' calculations.

Notes: Portfolio shares are percentage of total investment in risky assets. Asset levels in thousands of 1982 dollars. RRA stands for relative risk aversion level.

An investor is given a choice of a fixed sum of money next period or a lottery that pays \$800 with probability 0.5 and \$1200 with probability 0.5. A risk-neutral investor would be indifferent between the actuarial value of the lottery (\$1000) and the lottery. An investor with a RRA of 3 is indifferent between \$940 and the lottery, and an investor with a RRA of 5 is indifferent between \$900 and the lottery.

The results show that an electrical engineer with relative risk aversion of 3 should, according to the theory, hold a \$1.03 million portfolio of risky assets. This consists of a \$257,000 short position in SMB and long positions in HML and the market portfolio. The optimal risky positions are smaller if we consider an otherwise identical investor who is sixty years old, or one who has relative risk aversion of 5. Optimal holdings are also about 40 percent smaller for a secondary school teacher, because her permanent income is about 40 percent smaller. In line with the two-fund separation principle, none of these changes alter the optimal portfolio shares.

In each of these cases, optimal holdings are quite large relative to casual and systematic evidence regarding actual holdings (40-year-old electrical engineers with million dollar equity portfolios are not the norm!). One possible explanation for this gap between theory and evidence is the high returns on U.S. equities over the last century. Some analysts believe that these high returns are unlikely to hold in the future, so the last row in each panel of Table 6 shows optimal allocations for expected returns on risky assets that are only half as large as the corresponding sample means. Investment positions drop by half as well, but the optimal allocations remain quite

large compared to observed holdings for the typical person. This portfolio puzzle has received little attention in previous research because of the strong proclivity to focus on portfolio shares and to disregard theoretical implications for the level of risky asset holdings.¹⁷

We believe that the resolution of this puzzle rests at least partly on the opportunity cost of investor funds. In computing the portfolio allocations in Table 6, investors are allowed to borrow unlimited amounts at the risk-free interest rate. If, instead, investors must borrow at an interest rate that approximates the expected return on risky assets, then the optimal risky asset position is approximately *zero* when asset returns and labor income are uncorrelated. Many (potential) investors face an opportunity cost of funds at least as great as the expected return on equities, so it is not surprising that half or more of all households have little or no holdings of risky financial assets.

Endowed Exposure and the Breakdown of Two-Fund Separation

Nonzero correlations between asset returns and labor income cause two-fund separation to break down in a particular way. To illustrate this point, Table 7 shows optimal allocations for seven occupations when we account for correlation with labor income shocks. Recall from above that optimal holdings in the zero-correlation case, “desired exposure,” depend only on risk aversion, age, wealth, and asset returns. “Endowed exposure” gives the risky asset position implicit in the correlation between asset returns and the worker-investor’s labor income.

The results in Table 5 above demonstrated that most occupational groups have a negative endowed exposure to the SMB portfolio. As we explained above, the endowed exposure reflects the size and persistence of labor income innovations and their correlation with asset returns. Consequently, although income shocks for janitors and cleaners and electrical engineers have almost the same correlation with SMB, the endowed exposure of electrical engineers is much lower because their income shocks are more variable and more persistent.

To calculate an investor’s optimal portfolio, we simply subtract endowed exposure from desired exposure. Since endowed exposure is not proportional to desired exposure, two-fund separation fails. Other things equal, the bigger the endowed exposure the bigger the departure from the two-fund separation principle. Table 8 illustrates this breakdown by showing optimal portfolio shares under different assumptions about risk aversion and excess returns for each occupation that has a non-zero correlation with one or more of the assets. The base case uses sample average excess returns and a relative risk aversion of 3. Given these assumptions, the departures from two-fund separation are modest. For example, the optimal shares for

TABLE 7. Endowed exposure, desired exposure and portfolio holdings

		<i>SBM</i>	<i>HML</i>	<i>Market</i>	<i>Total</i>
Accountants and Auditors	Endowed Exposure	-36	0	0	-36
	Desired Exposure	-189	662	280	753
	Portfolio position	-153	662	280	789
Electrical Engineers	Endowed Exposure	-28	0	0	-28
	Desired Exposure	-257	903	381	1027
	Portfolio position	-229	903	381	1055
Elementary School Teachers	Endowed Exposure	-42	0	0	-42
	Desired Exposure	-139	488	206	555
	Portfolio position	-97	488	206	597
Secondary School Teachers	Endowed Exposure	-52	0	0	-52
	Desired Exposure	-158	556	235	632
	Portfolio position	-106	556	235	684
Janitors and Cleaners	Endowed Exposure	-13	0	0	-13
	Desired Exposure	-90	315	133	359
	Portfolio position	-76	315	133	372
Plumbers	Endowed Exposure	-46	0	0	-46
	Desired Exposure	-170	597	252	679
	Portfolio position	-124	597	252	725
Truck Drivers	Endowed Exposure	-0	16	-0	16
	Desired Exposure	-141	497	210	565
	Portfolio position	-141	481	210	550

Source: Authors' calculations.

Notes: Entries show endowed exposure, desired exposure and optimal portfolio position for indicated risky assets in thousands of 1982\$. Calculations assume a 40-year-old investor who has a relative risk aversion of 3.

electrical engineers never differ from the zero-correlation optimum by more than three percentage points. For secondary school teachers, the traditional zero-correlation portfolio understates SMB holdings by nine percentage points.

Because these effects are small, a portfolio manager might be forgiven for ignoring them. However, if one believes that high equity returns in recent decades are an aberration, or that expected returns have declined in recent years, then the effects of correlation on optimal portfolio shares become more important. As an example, the second line for each occupation in Table 8 shows optimal portfolio shares when we set excess returns to one-half their sample averages. Recall that this change has no impact on the optimal shares when two-fund separation holds. In particular, the optimal SMB share is -25 percent under two-fund separation, regardless whether we scale down excess returns. This invariance result fails when we take correlation into account. As an example, the optimal SMB portfolio share for secondary school teachers is +2 percent when excess returns are half their sample values and relative risk aversion is 5. To understand this result,

TABLE 8. Risk Aversion, Excess Returns, and Optimal Portfolio Shares in the Case of Three Risky Assets

	<i>Reduction in Excess Returns (%)</i>	<i>RRA</i>	<i>SMB</i>	<i>HML</i>	<i>Market</i>
Accountants and Auditors	0	3	-19	84	35
	50	5	-8	76	32
	75	5	5	67	28
Electrical Engineers	0	3	-22	86	36
	50	5	-15	81	34
	75	5	-6	75	31
Elementary School Teachers	0	3	-16	82	35
	50	5	0	70	30
	75	5	17	58	25
Secondary School Teachers	0	3	-16	81	34
	50	5	2	69	29
	75	5	19	57	24
Janitors and Cleaners	0	3	-21	85	36
	50	5	-11	85	33
	75	5	-0	70	30
Plumbers	0	3	-17	82	35
	50	5	-2	72	30
	75	5	14	61	26
Truck Drivers	0	3	-26	88	38
	50	5	-28	87	41
	75	5	-31	85	46

Source: Authors' calculations.

Notes: "Percent Reduction In Excess Returns" of 0 means that expected returns on risky assets are set to realized sample values.

A 50 percent reduction means that the excess return (sample mean return minus a risk-free rate of 3.5%) is set to half its sample value, and similarly for a 75 percent reduction. RRA stands for relative risk aversion level. Entries in the last three columns show the percentage of risky financial asset holdings in the indicated asset. All calculations assume an investor 40 years old.

recall that the level of excess returns has no effect on "endowed exposure." So, as we reduce excess returns and hence desired exposure, the relative size of endowed exposure goes up.

Higher risk aversion has the same effect, and for much the same reason. Greater risk aversion lowers desired exposure but does not affect endowed exposure. The last line in each panel of Table 8 shows optimal portfolio shares for the case of high risk aversion and low excess returns. In this case, the optimal portfolio shares sometimes differ substantially from the two-fund separation case. Based on traditional mean-variance analysis, a portfolio advisor would recommend a 25 percent short position in SMB. In contrast, the optimal position for secondary school teachers is a 17 percent long position in a plausible case that accounts for correlation between asset returns and labor income.

Life Cycle Variation in Endowed Exposure

Endowed exposure to occupation-specific assets varies over the life cycle, as illustrated in Table 9. Given an age-invariant correlation between labor income innovations and asset returns, the endowed exposure declines monotonically with age as the worker-investor draws down the present value of future labor income. This result follows immediately when the correlation between labor income innovations and asset returns is age invariant.¹⁸

A final issue involves life cycle variation in the extent of departures from two-fund separation. Other things equal, a declining path of endowed exposure leads to ever smaller departures from two-fund separation as an investor ages. However, income smoothing capacity also declines with age, which creates a countervailing force. In particular, age intensifies the effect of correlation on optimal portfolio shares, as we discussed above. So, for any given level of endowed exposure, the departure from two-fund separation is bigger for an older investor.

Conclusion and Discussion

When labor income and asset returns are correlated, investors are implicitly endowed with certain exposures to risky financial assets. These endowed exposures have important effects on optimal portfolio allocation. The two-fund separation principle that governs optimal portfolio choice in a traditional mean-variance setting breaks down when investors have endowed exposures to risky assets. In simple terms, an investor's optimal portfolio can be calculated as the difference between her desired exposure to risky assets and her endowed exposure. Because investors typically differ in their endowed exposures, they also differ in their optimal portfolio allocations (levels and shares), even when they have the same tolerance for risk and the same beliefs about asset returns.

Our graphical approach to portfolio choice over the life cycle accounts for an investor's endowed and desired exposures. The approach easily handles risky labor income, multiple risky assets, many periods, and several determinants of portfolio choice over the life cycle. As an added virtue, the chief empirical inputs into the framework are easily estimated using simple statistical procedures.

The empirical model relies on repeated cross sections to extract occupation-level components of individual income innovations. Annual data from 1968 to 1994 yield little evidence that occupation-level income innovations are correlated with aggregate equity returns. This finding, along with similar findings in other work, presents something of a puzzle for standard equilibrium models of economic fluctuations, growth, and asset pricing. Given rational asset pricing behavior, frictionless financial markets, and standard specifications of the production technology, dynamic equi-

TABLE 9. Endowed Exposure to Occupation-Specific Assets

	<i>Age</i>						<i>Asset</i>
	30	35	40	45	50	55	
Electrical Engineers	9.5	8.9	8.2	7.4	6.4	5.2	Build
Registered Nurses	-6.7	-6.0	-5.1	-4.1	-2.9	-1.4	Health
Elementary School Teachers	12.2	10.9	9.4	7.7	5.5	3.0	Educ
Secondary School Teachers	17.2	15.5	13.5	11.1	8.3	4.9	Educ
Electricians	14.5	13.3	11.9	10.2	8.2	5.8	Build
Plumbers	21.2	19.5	17.3	14.8	11.8	8.2	Build

Source: Authors' calculations.

Notes: Table entries report the endowed exposure to the indicated asset based on the best-fitting specification in regressions of occupation-level income innovations on SMB, HML, Bonds and the indicated industry-level return measure. See Table 7 in Davis and Willen (2000b) for the underlying regression results.

librium models imply a high correlation between aggregate equity returns and shocks to the present value of labor income. That implication finds little support in our empirical results.

We do find that several other asset return measures are correlated with occupation-level income innovations. The returns on portfolios formed on firm size (market capitalization) are significantly correlated with occupation-level income innovations for about half the occupations we consider. In a few occupations, income innovations are correlated with returns on long-term bonds. In several instances, industry-level equity returns are correlated with the occupation-level income innovations of the workers in those industries. Both a priori reasoning and our empirical results suggest that industry-level equity returns can covary negatively or positively with labor income innovations for industry workers. It follows that the optimal hedge portfolio for occupation-specific and industry-specific components of risky labor income cannot be discerned without intensive empirical study.

Applying the estimated correlations to our portfolio choice framework, we find sizable departures from the two-fund separation principle for plausible assumptions about expected asset returns and investor risk aversion. To the extent that future research more fully uncovers the correlation structure between labor income shocks and asset returns, the gap between optimal portfolio allocations and the uniform portfolio shares implied by the two-fund separation principle will be larger than the ones shown in our examples.

Notes

The authors thank Deborah Lucas, Olivia Mitchell, and Stephen Zeldes for helpful comments. Jeremy Nalewaik provided valuable research assistance. Davis gratefully

acknowledges research support from the University of Chicago Graduate School of Business.

1. This paper draws heavily on Davis and Willen (2000b). We direct the interested reader to that paper for a more thorough mathematical treatment of the issues and the approach developed herein.

2. Bodie, Merton, and Samuelson (1992) derive analytical solutions for portfolio choice in a continuous time finite horizon setting with fully hedgeable labor income risks. Much other work adopts computationally-intensive approaches to the portfolio implications of unhedgeable or partly hedgeable labor income risks. (For example, see Cocco, Gomes, and Maenhout 1999 for analysis in a finite horizon setting, and Heaton and Lucas 1997, Viceira 1998 and Haliassos and Michaelides 1999 in infinite horizon settings.)

3. For the utility specification that underlies our analysis, absolute risk aversion is unaffected by wealth shocks, and an investor's cost of a particular risk is unaffected by uncorrelated risks. However, an increase in the particular risk under consideration reduces the investor's willingness to take on more of that same risk. If the investor has constant relative risk aversion, then an increase in uncorrelated risks also reduces the investor's willingness to take on more of any particular risk.

4. Endowed exposure depends on the number of years to retirement, as we discuss more fully below, but this horizon effect on endowed exposure is distinct from the age effect on desired exposure.

5. Other studies investigate the issue at a more aggregated level in an international setting. Botazzi, Pesenti, and van Wincoop (1996) consider the covariance of national labor income shocks with financial asset returns, and Baxter and Jermann (1997) consider their covariance with the returns on hypothetical claims to a country's capital stock. Davis, Nalewaik, and Willen (2000) consider the covariance between national output shocks and a variety of domestic and foreign asset returns for 18 industrialized countries.

6. We also exclude persons who report an hourly wage less than 75 percent of the federal minimum. We handle top-coded earnings observations in the same manner as Katz and Murphy (1992).

7. The detailed occupational classification schemes in the CPS underwent major changes over time. Where possible, we constructed a uniform classification scheme from 1967 or 1970 to 1994 based on the occupational descriptions in the CPS documentation and an examination of changes over time in occupational cell counts and mean occupational earnings. We omitted individual-level observations that met any of the following occupation-level selection criteria: (1) the occupational group could not be extended back to 1970 or earlier in a consistent manner; (2) self-employed persons account for a large fraction of occupational employment (examples include physicians, dentists, lawyers, and farmers); (3) the occupational category was vague (examples include "General Office Supervisors" and "Financial Managers"); and (4) the number of individual-level observations in the occupation had a mean annual cell count less than 100 or a minimum annual cell count less than 50. These selection criteria reduced the number of individual-level observations by about one-half. From these 57 occupations, we selected for further analysis 10 occupations with large cell counts and a consistent definition back to 1967.

8. All earnings are expressed in 1982 dollars using the GDP deflator for personal consumption expenditures.

9. A log earnings specification is more commonly used by empirical researchers, but the specification in natural units fits more closely with our underlying theoretical model. In Davis and Willen (2000b), we show that log specifications yield results that are highly similar to specifications in natural units.

10. The empirical approach abstracts from potential selection issues associated with worker mobility across occupational groups, as well as mobility between the employment and not working. As a consequence, our estimates of the stochastic process for the occupation-level component of individual earnings may be incorrect even for infra marginal workers who do not move. A more complicated treatment of these issues requires long panel data sets. Davis and Willen (2000a) construct long times series for synthetic persons defined in terms of sex, birth cohort, and educational attainment; alternatively, one could use a true panel such as the Panel Survey of Income Dynamics. In practice, the true panel approach has serious limitations imposed by the nature and limited size of available surveys.

11. The data were taken from Professor French's web site (web.mit.edu/kfrench/www/data.library.html).

12. By "conventional" we mean a production technology that is approximately Cobb-Douglas over capital and labor. Given a stable Cobb-Douglas technology and a competitive economy, factor income shares are constant over time. Hence, if the same discount rates apply to future capital and labor income, and asset prices reflect fundamentals, the unobserved value of aggregate human capital fluctuates in a manner that is perfectly correlated with the observed value of claims to the aggregate capital stock. Models with these ingredients are standard, but they are hard to reconcile with the emerging body of work that finds how correlations between aggregate equity returns and labor income innovations.

13. For example, SIC 872 (Accounting and Auditing) is a natural industry counterpart for the Accounting and Auditing occupation, but CRSP contains no firm-level observations for SIC 872 during much of the sample.

14. These data are obtained from Professor French's web site (web.mit.edu/kfrench/www/data.library.html).

15. For references to related work see Fama and French (1992, 1993, 1996). Cochrane (2000) reviews the asset-pricing evidence related to size and book-to-market factors and provides references to more recent work.

16. Table 4 shows a very modest return premium on small cap stocks during our sample period. As others have observed, the realized premium on small cap stocks has declined in recent decades. The average annual value of the Fama-French SMB portfolio return was about 8 percentage points from 1964 to 1980 and -4 percentage points from 1981 to 1994.

17. Davis, Nalewaik, and Willen (2000) discuss this portfolio puzzle in connection with the gains to international trade in risky financial assets.

18. This covariance is allowed to vary smoothly with age in Davis and Willen (2000a), but they find only modest life cycle variation for demographic groups defined in terms of sex, education and birth cohort. Given their findings, and since their empirical design is better suited for uncovering age effects of this sort, we imposed an age-invariant covariance structure in this paper.

References

- Ameriks, John and Stephen P. Zeldes. 2000. "How Do Household Portfolio Shares Vary with Age?" Columbia University Working Paper, Columbia University. May 22.
- Baxter, Marianne and Urban J. Jermann. 1997. "The International Diversification Puzzle Is Worse Than You Think." *American Economic Review* 87, 1 (March): 170-80.
- Bernheim, B. Douglas, Lorenzo Forni, Jagadeesh Gokhale, and Laurence J. Kotlikoff. This volume. "An Economic Approach to Setting Retirement Saving Goals."
- Bodie, Zvi, Robert C. Merton, and William F. Samuelson. 1992. "Labor Supply Flex-

- ibility and Portfolio Choice in a Life Cycle Model." *Journal of Economic Dynamics and Control*, 16: 427–449.
- Botazzi, Laura, Paolo Pesenti, and Eric van Wincoop. 1996. "Wages, Profits, and the International Portfolio Puzzle." *European Economic Review* 40, 2: 219–54.
- Campbell, John Y., João F. Gomes, Francisco J. Gomes, and Pascal J. Maenhout. 1999. "Investing Retirement Wealth: A Life-Cycle Model." NBER Working Paper 7029.
- Canner, Niko, N. Gregory Mankiw, and David N. Weil. 1998. "An Asset Allocation Puzzle." *American Economic Review* 87, 1 (March): 181–91.
- Cocco, Joao, Francisco J. Gomes, and Pascal Maenhout. 1999. "Consumption and Portfolio Choice over the Life-Cycle." Harvard University Working Paper.
- Cochrane, John H. 2000. "New Facts in Finance." In *Economic Perspectives*. Chicago: Federal Reserve Bank of Chicago.
- Davis, Steven J. and Paul Willen. 2000a. "Using Financial Assets to Hedge Labor Income Risks: Estimating the Benefits." University of Chicago and Princeton Working Paper.
- . 2000b. "Occupation-Level Income Shocks and Asset Returns: Their Covariance and Implications for Portfolio Choice." NBER Working Paper 7905.
- Davis, Steven, Jeremy Nalewaik, and Paul Willen. 2000. "On the Gains to International Trade in Risky Financial Assets." NBER Working Paper 7796.
- Davis, Steven J., Wendy Edelberg, Jeremy Nalewaik, and Paul Willen. 2000. "The Opportunity Cost of Funds and Cross-Sectional Variation in Risky Asset Holdings." In progress, University of Chicago Graduate School of Business.
- Dréze, Jacques H. and Franco Modigliani. 1972. "Consumption Decisions Under Uncertainty," *Journal of Economic Theory* 5: 308–35.
- Fama, Eugene F. and Kenneth F. French. 1992. "The Cross-Section of Expected Stock Returns." *Journal of Finance* 47, 2 (June): 427–65.
- . 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33: 3–56.
- . 1996. "Multifactor Explanations of Asset-Pricing Anomalies." *Journal of Finance* 51, 1 (March): 55–84.
- Fama, Eugene F. and William Schwert. 1977. "Human Capital and Capital Market Equilibrium." *Journal of Financial Economics* 4, 1 (January): 95–125.
- Haliassos, Michael and Alexander Michaelides. 1999. "Portfolio Choice and Liquidity Constraints." Working Paper, University of Cyprus.
- Heaton, John and Deborah Lucas. 1997. "Market frictions, Savings Behavior and Portfolio Choice." *Macroeconomic Dynamics* 1, 1: 76–101.
- . 2000. "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *Journal of Finance* (June).
- Ingersoll, Jonathan E., Jr. 1987. *Theory of Financial Decision Making*. Totowa, N.J.: Rowman & Littlefield.
- Jagannathan, R. and Narayana R. Kocherlakota. 1996. "Why Should Older People Invest Less in Stocks Than Younger People?" *Federal Reserve Bank of Minneapolis Quarterly Review* 20, 3: 11–23.
- Katz, Lawrence F. and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–87: Supply and Demand Factors." *Quarterly Journal of Economics* 107: 35–78.
- Keeler, Theodore E. 1989. "Deregulation and Scale Economies in the U.S. Trucking Industry." *Journal of Law and Economics* 32: 229–55.
- Leibowitz, Martin L., J. Benson Durham, P. Brett Hammond, and Michael Heller. This volume. "Retirement Planning and the Asset/Salary Ratio."
- MaCurdy, Thomas E. 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," *Journal of Econometrics* 18: 83–114.

- Rose, Nancy L. 1987. "Labor Rent Sharing and Regulation: Evidence from the Trucking Industry." *Journal of Political Economy* 95: 1146–78.
- Samuelson, Paul A. 1969. "Lifetime Portfolio Selection by Dynamic Stochastic Programming." *Review of Economics and Statistics* 51: 239–43.
- Viceira, Luis. 1998. "Optimal Portfolio Choice for Long-Horizon Investors with Non-tradable Labor Income." Harvard University Working Paper, Cambridge, Mass.
- Willen, Paul. 1999. "Welfare, Financial Innovation and Self Insurance in Dynamic Incomplete Market Models," Princeton University Working Paper, Princeton, N.J. May.