The Role of International Investment in a Privatized Social Security System

Marianne Baxter and Robert G. King
University of Virginia and NBER

December 1998

1 Introduction

This paper asks what role international financial assets might play in a privatized Social Security system. Answering this question requires that we begin by indentifying the implicit and explicit public policy goals of the Social Security system, as follows: (1) provision of income insurance; (2) ensuring that individuals save for retirement; and (3) redistribution of income within and between generations.[1] In this paper, we will be primarily concerned with the savings issue and the related question of optimal portfolio choice. Specifically, we ask what role international assets might play in a well-designed lifetime savings plan. We will not address questions of redistribution, or even the important question of how to resolve the problem of Social Security financial solvency.

At this point, it is not clear what type of Social Security reform will ultimately be undertaken. However, one unifying element of many proposals involves the potential for retirement funds to be invested in risky assets such as equities. Under some proposals, these risky assets would be invested by the government on the individual's behalf. Under other proposals, these assets would be managed directly by the individual in a 'Private Retirement Account' (PRA). In each case, however, the idea is that including risky assets could improve welfare of retirees because risky assets on average yield higher returns.

What is special about international assets? It is well known that individuals typically do not diversify their holdings of risky assets internationally to the extent

[1] Miron and Weil (1997) discuss the ways in which Social Security has changed from its inception in the 1930s to the present. They make the point that improvements in health and reductions in mortality risk mean that Social Security has evolved from a scheme that mainly provided old-age insurance to a scheme that is essentially a forced-saving plan. Sargent (1998) argues that Social Security was first conceived as a remedy for too much saving, while current policymakers are more concerned about too little saving.

that standard portfolio theory would predict–U.S. investors hold more than 90% of their risky assets in the form of U.S. securities. A large number of studies have documented the improvement in the risk-return profile that could be gained through international diversification, yet home bias in asset holdings persists. Many believe that home bias is due to a kind of investment xenophobia. In any case, there is no consistent theory that explains home bias as a rational choice by fully-informed individuals.[2] Thus, the challenge to overcoming home bias in retirement portfolios is to show that international investments can yield important welfare gains.

The goal of this paper is to explore whether it is likely that there are important gains to international diversification in a reformed, or 'modernized,' Social Security system. There are two main ways in which international assets, or any asset for that matter, may be important in a life-cycle savings plan. First, the asset may provide useful diversification benefits–including the asset in a portfolio may improve the risk-return tradeoff facing the investor. Second, the asset may be useful for hedging risks associated with nontraded assets. This paper focuses on two important, nontraded asset. The first is human capital–the largest source of wealth for all non-retired individuals. For the most part, people cannot trade claims to future labor income. Because there are random fluctuations in labor income, human capital is risky. Individuals would like to hold a portfolio of traded financial assets that acts to hedge the risk associated with nontraded human capital.

The current Social Security system also represents an important nontraded risky asset. Individuals contribute to Social Security during their working life and receive payouts during retirement according to specific rules. Social Security represents an asset that individuals *must* purchase, and this asset cannot be sold or traded. It is a risky asset for individuals who are still working, as we make clear in Section 3 below. To the extent that elements of the current Social Security system are retained in a reformed system, individuals will want to hedge this Social Security risk.

The paper is structured as follows. Section 2 provides a quick overview of Social Security, focusing on the nature of its cash flows. We then ask what privately-provided assets most closely resemble Social Security. This is useful both because it helps us understand the nature of Social Security as a financial claim and because it indicates what products might be desirable to individuals as part of a reformed Social Security system. Section 3 addresses the diversification issue, by exploring the consequences of allowing individuals to add risky assets to their retirement portfolios. Using data on U.S. and foreign asset returns, we evaluate the extent to which individuals could improve their risk/return tradeoff by investing in foreign securities.

Section 4 takes up the question of whether financial assets could be used to hedge the risk of nontraded human capital. We provide evidence on the behavior of labor income, in the aggregate as well as across industries and across educational groups. We then explore whether returns on financial assets are significant predictors of labor

[2] See French and Poterba (199x) for a rationalization of home bias that depends on small errors in individuals' subjective probability distributions for home and foreign returns.

income growth – a necessary condition for financial assets to be useful as hedges for nontraded human capital. We do find significant predictability for many categories of labor income. We also provide an overview of two approaches to computing 'returns' on human capital and provide some new evidence on human capital returns.

Section 5 takes a closer look at the optimal structure of an individual's retirement portfolio. We develop an optimizing model that incorporates subsistence consumption, and ask how individuals' optimal portfolios depend on (i) their level of risk aversion; and (ii) their expected lifetime. We work first with a model in which the lifetime is certain; we then extend the model to a situation with uncertain lifetimes. In each case, our results indicate an important role for foreign assets in the retirement portfolio.

2 Social Security: An overview

For most Americans, Social Security plays an important role in the retirement portfolio. Table 1, based on Gustman, et al. (1997) and Gustman and Steinmeier (1998) describes asset holdings for 51-61-year-olds as compiled in the Health and Retirement Survey. This table shows the value of Social Security claims for these individuals, broken down by lifetime earnings percentiles. For the group as a whole, the value of Social Security represented 31% of other assets. Social Security value represents a smaller fraction–about 12%–of other assets for individuals in the highest and lowest earnings groups. At the high end, Social Security value is 54% of other assets for individuals in the 10-25% bracket (these individuals have average lifetime earnings of $392,781, or about $13,000 per year, assuming the individual worked from age 25 to age 65).

Social Security represents 40% of all income to the aged population, as reported by the Social Security Administration (page 6, "Facts and Figures about Social Security 1998"), and represents more than 50% of income for two-thirds of "beneficiary units (couples or nonmarried persons). Social Security was the only sources of income for 18%.

Under each of the Social Security reform proposals currently under review in the U.S. Congress, an element of the traditional Social Security System will be kept intact. Viewed in finance terms, Social Security represents a nontraded financial asset – we say it is nontraded because a worker who has paid into the Social Security and thus accumulated retirement benefits through the system cannot sell or trade his claim on Social Security. In this section we explore the risk and return characteristics of this nontraded financial asset. We begin by reviewing the salient characteristics of the Social Security claim and then compare this claim to tradable financial assets available in the private market.

2.1 Social Security cash flows

A covered individual pays Social Security tax equal to 7.65% of his income up to a cap of \$68,400; this contribution is matched by employers. Upon retirement, the individual becomes eligible for payments from the Social Security system.

It will be useful to develop notation for the different types of cash flows embedded in the Social Security asset. An individual begins his working life at date t, making contributions Z_{t+j} to Social Security (Z_{t+j} includes employer contributions). These contributions continue until the individual retires in period $t+R$ (that is: the worker's last period of employment is period $t+R$). From retirement at date $t+R+1$ until the individual's death at date $t+T$, the Social Security system pays benefits B_{t+R+j} to the individual. (We are ignoring survivor benefits and death benefits.)

Working years: Covered workers are required to contribute to Social Security an amount equal to 7.65% of their earnings up to a cap, or ceiling amount denoted CAP_{t+j}. Currently, the cap is equal to \$65,000. Individuals' Social Security contributions, or payments, are matched by employers. Thus the total payment is

$$Z_{t+j} = 0.15 * \max(W_{t+j}, CAP_{t+j}),\ j = 1, 2, ..., R.$$

The retirement point and beyond: When an individual retires at date $t+R$, Social Security benefits are computed as follows. A worker's earnings during the pre-retirement years are adjusted to reflect the increase in the general level of wages during the particular working year in which the wages were earned and the retirement date. The idea behind this adjustment, or indexation, is to ensure that a retired worker's benefits reflect the rise in wages that occurred over his working life. This adjustment is carried out through the computation of the individual's "Average Indexed Monthly Earnings," or AIME. The AIME is constructed as follows. Let $\overline{W}_{t+j}$ denote the national average wage index for year $t+j$, and let $\overline{W}_{t+R}$ denote the national average wage index for the retirement year, $t+R$.[3] The worker's indexed earnings $W^{indexed}_{t+j}$ are computed as

$$W^{indexed}_{t+j} = W_{t+j}\left(\frac{\overline{W}_{t+R}}{\overline{W}_{t+j}}\right),\ j = 1, 2, ..., R \quad (1)$$

The worker's average indexed monthly earnings is then the average over the indi-

[3] The "national average wage index" computed by Social Security is actually a measure of average income for covered workers. As such, it is comparable to a measure of average per capita labor income, not an average wage rate.

vidual's working life of $W_{t+j}^{indexed}$:[4]

$$AIME = \frac{1}{12}\left[\frac{1}{R+1}\sum_{j=0}^{R} W_{t+j}\left(\frac{\overline{W}_{t+R}}{\overline{W}_{t+j}}\right)\right]. \tag{2}$$

The individual's Social Security benefit, B_{t+R+j}, depends on the individual's AIME. Specifically, the individual's "Primary Insurance Amount," PIA, equals 90% of the first \$477 of AIME plus 32% of AIME over \$477 through \$2,875 plus 15% of AIME over \$2875. The fact that the PIA is a concave function of the AIME can be interpreted as implicit taxation, or redistribution, from retirees who had high incomes over their lifetime toward retirees who had lower incomes. Let PIA_{t+R+1} denote the PIA for a worker who retires in period $t+R$ and begins receiving benefits in period $(t+R+1)$, expressed on an annual basis. In subsequent periods, the PIA will be adjusted for inflation, so that benefits are given by:

$$B_{t+R+j} = PIA_{t+R+1}\left(\frac{P_{t+R+1+j}}{P_{t+R+1}}\right),\ j = 1, 2, ..., (T-t-R-1),$$

where $P_{t+R+1+j}$ is the CPI price index.

2.2 Private-market assets similar to Social Security

Under some proposed reforms, Social Security will be largely or completely privatized. For this reason, it is important to understand what private-market financial instruments could be used to replace a government-run Social Security system. Further, if Social Security remains as a government-run program, but is "modernized" in the sense of offering new products, it is useful to understand the similarities and differences between Social Security and similar, privately-provided products.

To begin, Social Security benefits are best viewed as the flow from a real annuity. We described in the prior section how this real annuity is 'purchased' through Social Security contributions during the individual's working life. We also discussed the factors determining the level of Social Security benefits after retirement. In this section, we describe private-market annuity products, highlighting the similarities and differences with the Social Security annuity. Annuities are technically insurance contracts, and are marketed by the major insurance companies. Increasingly, however, annuity products are being offered by insurance company affiliates of the major mutual fund companies, such as Vanguard and Fidelity.[5]

[4] The AIME calculation excludes some low-wage years and years of non-employment–a feature not captured by our formula.

[5] OASDI Administrative costs as a percent of contributions for FY 1997 = 0.8%. Compare Vanguard costs.

2.2.1 Fixed annuities

A fixed annuity specifies a sequence of known, fixed cash flows that will be received by the annuitant. Generally, the cash flows will continue (once begun) for the lifetime of the annuity owner; alternatively, the time period may be shorter and would be specified in advance of the first distribution. Private annuity markets provide fixed annuities with a variety of survivorship characteristics. Some annuities will provide benefits just for the lifetime of the annuitant; others provide survivorship benefits of various kinds; see Table 2. A "fixed" annuity is not a "real" annuity, because it provides a guaranteed stream of fixed, nominal cash payments. Social Security, by contrast, indexes the cash payments to the CPI inflation rate.

2.2.2 Variable annuities

A variable annuity is a flexible-premium annuity, in the sense that the annuitant can make investments whenever he wishes during the accumulation phase. The private annuity market currently provides a broad range of investment options for variable annuities, including Money Market funds, a variety of bond funds, and several stock funds, including international portfolios. During the accumulation phase the individual invests in one or more of these funds and is typically allowed to reallocate wealth across the funds with few restrictions.

The "distribution phase" of the annuity begins when the annuitant starts to take distributions, or cash flows, from the annuity. The annuity is called a deferred annuity if a period of time elapses between the accumulation and distribution phases; the annuity is called an immediate annuity if the distribution phase begins immediately after the annuity is purchased.[6]

Owners of variable annuities have many options for how the accumulated assets are paid out. The total accumulated value can be paid out in a single lump sum, or 'withdrawals' can be made at irregular intervals. Alternatively, and of the most interest for our purposes, the owner of the variable annuity can choose to annuitize his investment. The decision to annuitize is irreversible. A sample of the choices for annuitization are given in Table 2, which describes annuitization options available to investors in Vanguard's Variable Annuity products. This Table shows a wide range of options that vary according to whether the payouts are fixed dollar amounts

[6] Mitchell, Poterba and Warshawsky (1997) analyze the 'money's worth' of single premium, immediate annuities since it is clear that "the life annuity represents the predominant decumulation method," (page 5), something cannot be known with certainty for other types of annuities. According Mitchell, Poterba and Warshawsky, (1997, page 5) "flexible premium payments accounted for \$46.6 billion in 1995, compared with \$52.5 billion of single premium annuity payments. Single Premium Immediate Annuities, however, accounted for only \$6.2 billion of premium payments in 1995, while Single Premium Deferred Annuities accounted for \$46.3 billion. The small volume of SPIA purchases suggests that the recent growth of the aggregate annuity market has not resolved the long-standing puzzle, discussed for example in Friedman and Warshawsky (1990), of why individuals do not choose to annuitize their wealth."

that vary each period, or whether they vary according to the portfolio chosen by the investor. Further, the payout options differ in the length of time that payments are made and the extent to which payments may be transferred to a beneficiary in the event of the annuitant's death.

2.2.3 Social Security as a variable annuity

Using the terminology of the private annuity markets, Social Security is a variable annuity. Specifically, it is a flexible premium, deferred annuity, since the 'purchases' or contributions to Social Security vary over time and because the payouts are deferred to the retirement period. The private annuity from Table 2 that is closest to the Social Security annuity is the "Joint and Last Survivor Annuity" with fixed payments. This annuity yields fixed payments as long as either of the two annuitants is alive, which is similar to the Social Security provision to continue to provide benefits to surviving spouses.

Let's take a closer look at the similarities and differences between Social Security and private annuities. During the accumulation phase, the rate of return on the private annuity is the rate of return on the portfolio that the investor has chosen. Further, the investor in a private annuity may have the right to reallocate invested funds to alternative portfolios without cost. By contrast, the contributor to Social Security does not choose the investment vehicle for his contributions–Social Security contributions are ostensibly invested in U.S. Treasury securities. However, the rate of return on the Social Security 'asset' during the accumulation phase is *not* the rate of return on U.S. Treasury securities, as will be discussed more fully in Section 3. There, we argue that the rate of return is related to the rate of change of the U.S. average wage index.

Private annuities and Social Security also differ during the distribution phase (the retirement period). First, fixed-payment annuities offered by private firms are fixed in nominal terms, while Social Security payments are indexed to inflation, and therefore are fixed in real terms. Second, distributions from private annuities depend proportionately on the value of the annuity investment at the annuitization point. By contrast, the formulas used by the Social Security Administration to compute the AIME and the PIA, combined with rules about the family maximum, mean that distributions from Social Security are not simple linear functions of the value of contributions. Third, private annuities that have survivorship benefits do not condition these benefits on the income or wealth of the survivor. Social Security does have conditions of this sort: a surviving spouse who has worked at a high-income job may receive no survivorship benefits. This is a redistributive "estate tax" component built into the current Social Security system. Finally, private annuities are legal contracts specifying cash flows (or rules for determining cash flows), and cannot be changed.[7] By contrast, Social Security distributions can be changed by

[7] The Vanguard variable annuities have "separate account" status, which means that the assets

an act of Congress, adding an important element of uncertainty for participants in the Social Security System.

3 Social Security as part of a retirement portfolio

Many proposals for Social Security reform allow for individual choice among assets in which to invest retirement funds. Some proposals keep important elements of Social Security in its current form while allowing additional funds to be invested in assets of the individual's choosing. Other proposals contemplate allowing the Social Security Trust Fund to invest in a broader menu of traded financial assets. Our goal in this section is to evaluate the risk-return tradeoffs facing investors under the current Social Security system and also under systems that may be implemented in the future. We pay particular attention to the importance of the diversification benefits provided by including international assets in the investment portfolio.

3.1 Social Security risk and return

This section takes a first look at the risk and return characteristics of Social Security during the accumulation phase. A worker who contributes \$1 to Social Security will have that \$1 'marked up,' according to the AIME formula (1), by the rate of change of the Social Security wage indexing series for each period between the contribution date and the retirement date. That is: \$1 contributed at date t will grow as follows:

$$
\begin{gathered}
\text{\$1 at date t grows to:} \\
\$1\left(\frac{\overline{W}_{t+1}}{\overline{W}_t}\right) \text{ at date t+1} \\
\$1\left(\frac{\overline{W}_{t+1}}{\overline{W}_t}\right)\left(\frac{\overline{W}_{t+2}}{\overline{W}_{t+1}}\right) \text{ at date t+2} \\
\$1\left(\frac{\overline{W}_{t+1}}{\overline{W}_t}\right)\left(\frac{\overline{W}_{t+2}}{\overline{W}_{t+1}}\right)...\left(\frac{\overline{W}_{t+R}}{\overline{W}_{t+R-1}}\right) \text{ at retirement date, t+R.} \qquad (3)
\end{gathered}
$$

Equation (3) shows that the nominal return during the accumulation phase to \$1 'invested' in Social Security is the growth rate of the wage indexing series, since this is the amount by which the \$1 is 'marked up' each period up to the retirement date.[8]

purchased by the annuitant are not available to the insurance company's creditors in the case of bankruptcy.

[8] This definition of the return to Social Security over the accumulation phase ignores elements of risk associated with the possibility of future changes in the rules for 'marking up' contributions. It also ignores the "bend points" in the computation of the PIA as a function of the AIME. For workers in higher-wage groups, the mean Social Security return and the volatility of the return will be lower than the one we report here.

Figure 1: Growth rates of real wage index, per capita labor income, and GNP

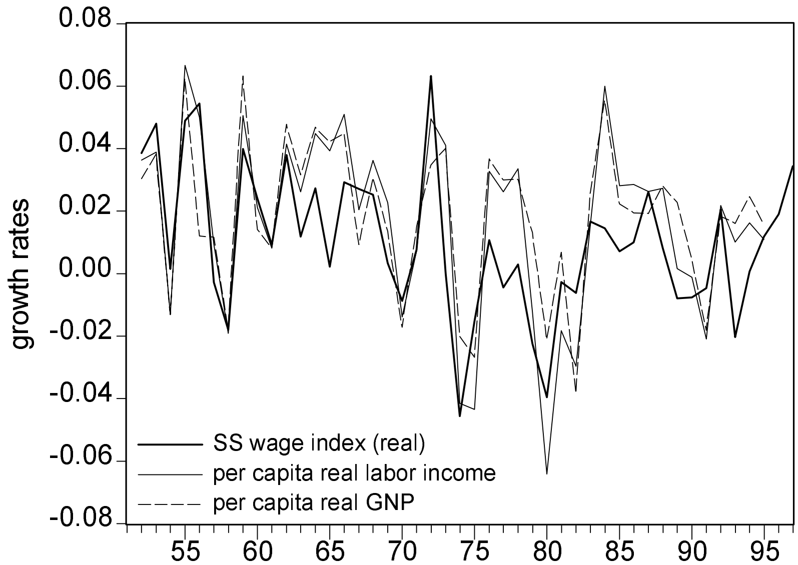

Since many proposals for Social Security reform call for alternative investment vehicles for Social Security contributions, it is of interest to compare the risk and return characteristics of investment in the present Social Security system to some potential, alternative rules governing Social Security investments. To begin, we studied how closely the growth in the wage indexing series corresponds to a measure of per capita income growth and per capita GNP growth. The growth rates of these variables, in real terms, is shown in Figure 1, and summary statistics are given in Table 3. The wage index and per capita labor income are deflated using the CPI deflator; GNP is deflated using the GNP deflator.

The real wage indexing series grew at only 0.95% per year for the 1951-1996 period, while real per capita income grew at 1.65% per year and per capita real GNP grew at 1.85% per year. The wage indexing series is about as volatile as GNP growth, while per capita labor income is somewhat more volatile. Although per capital real income growth is highly correlated (0.90) with per capita real GNP growth, the growth rate of the wage index has a correlation of only 0.84 with income growth and only 0.64 with GNP growth. This is potentially of interest because alternative rules for 'marking up' Social Security contributions might use the growth rate of per capita labor income or per capita GNP. Either of these rules would have meant higher 'returns' on Social Security over the past 45 years, with little increase in volatility.[9]

[9]We were surprised that the wage indexing series did not bear a closer relationship to these measures of aggregate economic activity. To the extent that "macro markets" of the type proposed

3.2 Risk and return of traded financial assets

An important building block for our analysis of the role of international investment in retirement portfolios is the risk/return structure of traded financial assets. The extent to which diversification and hedging benefits can be obtained by holding foreign assets depends on the correlation between returns on foreign assets and the returns on domestic assets, both traded and nontraded.

We begin by looking at the relationship between U.S. and U.K. equity returns, over the period 1918-1996, as reported in Table 4. The returns are annual and are expressed in real, local-currency units, deflated using consumer price indices.[10] Over this period, U.K. returns had higher real returns and higher volatility compared with the U.S.. The correlation between equity returns in the U.S. and the U.K. was 0.52.

If the end of the Bretton Woods system and the recent increase in openness of financial markets means that the structure of asset returns has changed, it may be useful to focus on a shorter, more recent time period. Table 5 gives information on equity and bond returns for a variety of portfolios, using quarterly data from 1970:1-1998:1. Returns are real, U.S. dollar returns, deflated using the GNP deflator. The equity returns are total returns, as reported by Morgan Stanley Capital International. The bond returns are total returns on long-term government securities, published by the IMF. The "Europe" bond portfolio is an equally-weighted portfolio of the U.K., France, Italy, Germany, and Switzerland. The "Asia" bond portfolio is and equally-weighted portfolio of Japan and Australia. The broad impression from Table 5 is that equity and bond returns were positively correlated across countries during this period. In particular, the correlation of U.S. equity returns with other equity returns is reasonably high: 0.67 in the case of U.S.-Europe and 0.43 in the case of U.S.-Far East. The correlation of U.S. equity returns with bond returns is lower, suggesting a potentially important role for U.S. and foreign bonds in a diversified portfolio.

To get an idea of the extent of the benefits from diversification among traded assets, Figure 2 shows the "efficient frontier" for combinations of U.S. and foreign assets. The inner (dashed) line is the frontier for U.S. assets only – combinations of U.S. equities and U.S. bonds. The points with a star inside a circle are U.S. assets (bonds and equities); the open circles are foreign assets.[11] The outer (solid) line is the frontier for U.S. and foreign assets combined. This graph indicates that diversifying internationally could be beneficial in terms of improving the risk-return tradeoff for a U.S. investor.

---

by Shiller (19xx) could allow individuals to hedge aggregate risk, the current Social Security wage indexing scheme appears to introduce idiosyncratic risk into the Social Security asset that could not be hedged with a claim on aggregate GNP.

[10] Data are from Campbell (1998).

[11] See Table 5 for the statistics that will help identify particular points – we will figure out how to label the points in the next draft of the paper!

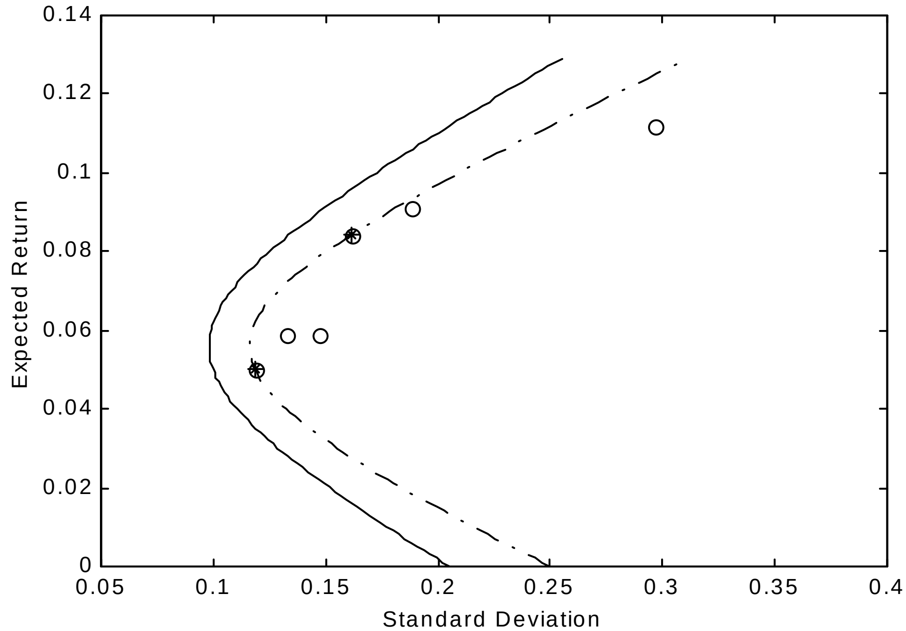

Figure 2: Portfolios of traded assets

3.3 Expanding the menu of Social Security assets

This section investigates the benefits of broadening the menu of assets that investors may use for their retirement savings, either as part of add-on PRAs of the type proposed by Feldstein and Samwick (1997) or as part of a 'modernized' Social Security scheme run by the government. In this section we will be focusing only on the diversification benefits obtained through a broader menu of assets; the potential for hedging benefits will be considered in the next section.

Figure 3 shows the risk-return tradeoff that can be achieved by combining the traditional Social Security 'asset' described earlier with other traded financial assets. This figure should be interpreted as describing the risk-return tradeoff facing a worker in the accumulation phase (i.e., still saving for retirement). The statistics used to compute this tradeoff are given in Table 5.[12] Social Security is the asset in the lower-left corner of the figure, with average return of 0.24% per year and standard deviation of 2.16%. As before, the dashed line represents the tradeoff for investors who confine their portfolios to U.S. assets: in this case, Social Security plus U.S. bonds and equities. The solid line represents the tradeoff for investors who invest in combinations of Social Security, U.S. assets, and foreign assets.

[12] We experimented with using a longer time period, 1951-1996, with only two traded securities: U.S. and U.K. equity returns (these are the only returns that we have for this longer time period). The results were qualitatively very similar to those reported in the text.

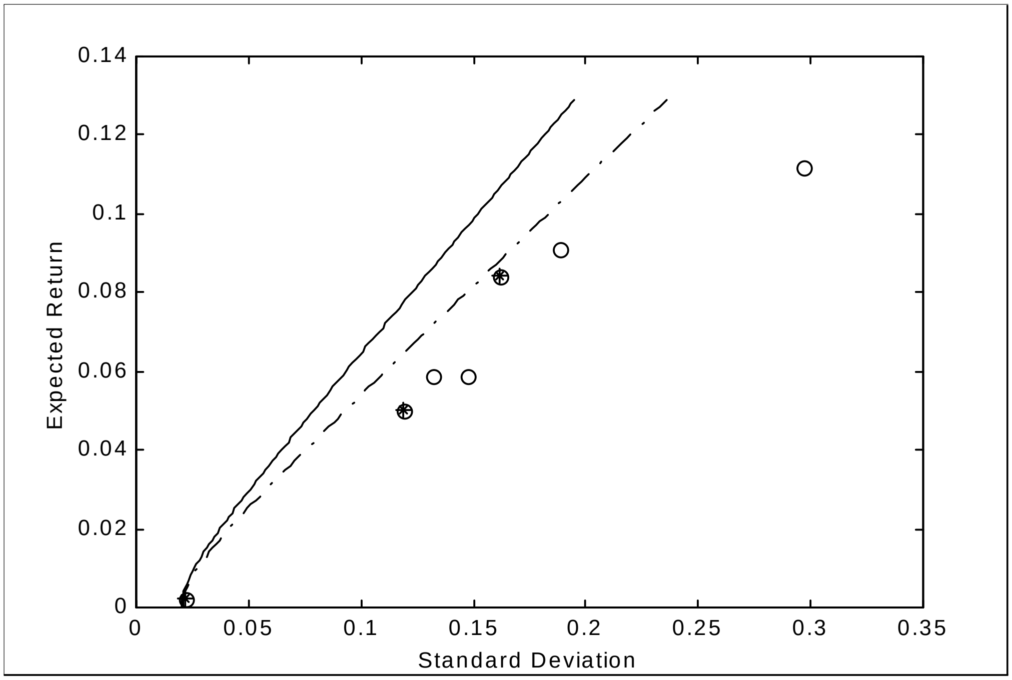

Figure 3: Portfolios combining Social Security with other assets

How useful are international assets to a retirement investor, through the diversification channel stressed here? The answer depends critically on several factors: (i) the tolerance of the investor for the higher levels of risk entailed in international investments; (ii) the risk-return profile of the investor's current portfolio; and (iii) the constraints on the investor due to the rules of particular retirement schemes.

Consider, for example, an investor who is invested 100% in the current Social Security system. Now imagine that this individual can undertake a small amount of investment in financial assets of his or her choice, either through PRAs or through some other scheme. How much does it matter whether this investor makes use of (or has access to) international investments? To be more concrete, suppose the investor wants to earn an expected return of 3.00% per year, rather than 0.25% per year (the return from being 100% in the current Social Security scheme). If the investor confines himself to Social Security plus U.S. assets, the portfolio shares necessary to achieve this return are as follows: 58.8% in Social Security, 23.1% in U.S. equities, and 18.1% in U.S. bonds (see Table 6 for more detail). The standard deviation of this portfolio is 5.8% per year. Adding international investments reduces the standard deviation of the portfolio only slightly, to 5.1% per year (the portfolio fractions are as follows: Social Security, 58.8%, U.S. equities, 17.9%, European equities, -0.4%, Far Eastern equities, 0.3%, U.S. bonds, 0.6%, European bonds, 17.9%, and Asian bonds, 1.2%).

When the individual's starting point is being fully invested in the current Social

Security scheme and the investor is contemplating only 'small' increases in the risk and return of his portfolio, there are large gains to including stocks or long-term bonds in the portfolio. This situation would apply to investors who are very risk averse, or are effectively risk averse because of their position in their life cycle or income level. It would also apply to investors who are constrained by the rules of a modified Social Security scheme that permitted only small fractions of one's investment portfolio to be invested in assets of the individual's choice. The marginal gains to international diversification are small, however, compared with investing all of the risky component of one's portfolio in domestic assets.

The benefits to international diversification are more noticeable when the circumstances or preferences permit individuals to take on higher levels of risk in their retirement portfolios. For example, a portfolio with expected return of 10% has standard deviation of 14.71% when only U.S. assets are considered, versus 12.27% when international assets are part of the portfolio.[13]

4 Human capital as a nontraded asset

Human capital is the single most important asset for most individuals. Since labor income represents about two-thirds of GNP, human capital represents roughly two-thirds of aggregate national wealth. There will be variations across individuals: young people who have had little opportunity to save and who have most of their working lives ahead of them will have a higher fraction of total wealth in the form of human capital; retirees will have a relatively small ratio of human capital to total wealth.

Table 7 gives an idea of the importance of human capital relative net worth for different types of individuals.[14] We computed the value of human capital by assuming a real growth rate of 2% per year for labor income over the individual's remaining working life, and discounted future labor income at the real rate r. If human capital is viewed as riskless, or as having only idiosyncratic risk, then it would be appropriate to use the risk-free interest rate, say $r = 0.02$, to discount future labor income. If, alternatively, human capital is about as risky as risky financial assets, then it would be more appropriate to use a discount like $r = 0.08$, approximately the average long-run return on equity. No matter which method is used to discount future labor income, Table 7 shows that human capital is the most important single asset held by non-retired individuals. It is especially important for low-income and young individuals.

[13] Portfolio shares in the U.S.–only portfolio are as follows: -12.8% in Social Security (or, nearly equivalently, borrowing T-bills); 69.1% in U.S. equities, and 43.7% in U.S. bonds. Portfolio shares in the international portfolio are -12% Social Security, 54.8% U.S. equities, -16.4% European equities, 12.0% Far Eastern equities, 9.7% U.S. bonds, 48.9% European bonds, and 2.3% Asian bonds.

[14] Friend and Blume (1975) also treat human capital as a nontraded, risky asset, and use data from an older Survey of Consumer Finances to compute an estimate of human capital as a fraction of net worth.

Even for individuals that are near retirement, in the 55-65 year age bracket, human capital is still the largest component of wealth.

Human capital is a risky asset because of uncertainty in labor income which represents the flow of 'dividends' from this asset. Part of the volatility of an individual's labor income may be common across all individuals, which can be characterized as 'aggregate' risk. The remaining component of income volatility is individual-specific, 'idiosyncratic' risk. This section looks at the characteristics of labor income volatility and the returns to human capital. Because human capital is a nontraded asset, it is necessary to compute 'synthetic returns' on this asset by combining observable data on labor income and other variables with an asset pricing model. This section reviews prior studies that have used alternative methods for computing these synthetic returns and provides some new results. Next, we investigate how the nontradability of human capital affects portfolio composition. Intuitively, traded financial assets with returns that are highly correlated with human capital returns can be used to hedge human capital risk. We evaluate the extent to which domestic vs. international assets may be useful as hedging tools.

4.1 Labor income volatility

As an initial measure of the risk associated with human capital, we can look at the volatility of labor income and its relationship to the business cycle. If labor income followed a simple random walk, then the return to human capital is simply the growth rate of labor income. The data do not support the random-walk model for human capital, but we will begin by looking at raw volatility of labor income growth.

4.1.1 Aggregate labor income

We begin by looking at a measure of aggregate labor income.[15] As we saw in Figure 1 above, the growth rates of real GNP and real labor income move together quite closely. In prior work, Baxter and Jermann (1997) studied data from the U.S. and three other OECD countries, and found that one cannot reject the hypothesis that aggregate labor income is cointegrated with GNP with a cointegrating vector of [1 -1] (implying that labor's share is stationary). Because aggregate labor income moves so closely with aggregate GNP, a claim to aggregate labor income in the U.S. is almost the same thing as a claim on aggregate GNP. Although the "macro markets" proposed by Shiller (19xx) do not (yet) exist, it is useful to note that a market in GNP could provide a useful hedging vehicle for the risk associated with aggregate labor income.

[15] Our measure of aggregate labor income is computed as the sum of compensation of employees (Citibase code GCOMP) and proprietor's income (GPROJ), deflated using the GNP deflator (GD).

4.1.2 Labor income at the industry level

This section looks at the properties of labor income at the industry level. We are interested in learning how the volatility of labor income and its relationship to the business cycle differs across industries. We gathered data on unemployment rates, u_{jt}, hourly wage rates, w_{jt}, and hours worked per employed person, h_{jt}, for nine manufacturing industries indexed by $j = 1, 2, ..., 9$. We wanted a measure of the earnings that would be received in an industry in a particular period, taking into account the possibility of being unemployed. Thus, we computed a measure of the "unemployment-corrected" total earnings in industry j as follows:

$$earnings_{jt} = (1 - u_{jt})w_{jt}h_{jt}.$$

Table 8 shows how the mean growth rate of real earnings and its volatility differ across industries; the growth rate of real GNP is included for comparison.[16] Panel A shows that there is a great deal of cross-industry variability in the behavior of unemployment-corrected labor income growth. In each of these manufacturing industries, the mean growth rate of wages is lower than the mean growth rate of GNP. Only two industries: services and wholesale trade–had wage growth that was less volatile than GNP growth. Most industries had wage growth that was substantially more volatile than GNP growth – more than twice as volatile for construction and mining.

To get an idea of the cyclic sensitivity of wages in these industries, we ran a regression of wage growth on GNP growth, and have reported the coefficient on GNP growth together with the R^2 of this regression.[17] The "beta" coefficient gives an idea of the cyclic sensitivity of wage growth to GNP growth; we intend the terminology to suggest an individual stock's "beta" with respect to a market index. If there were a market claim on GNP, this "beta" would help determine the loading on GNP in the hedge portfolio for that category of labor income. The R^2 of the regression indicates how much of the volatility of wage growth in a particular industry represents "aggregate risk," where we take aggregate risk to mean the risk associated with aggregate movements in GNP or changes in aggregate wages, which we have seen are nearly the same thing.

Table 8 shows that the manufacturing industries differ widely in their sensitivity to GNP changes, with estimated "betas" ranging from a low of 0.22 for services to coefficients close to 1.0 for construction, durables manufacturing, and mining. The R^2s for these regressions are rather low, indicating that most of the volatility in labor

[16] We have used the GNP deflator for wages instead of the CPI because we are interested in comparing the behavior of real earnings with real GNP. Conceptual problems with the CPI deflator, combined with the fact that the ratio of the CPI to the GNP deflator exhibits a significant upward trend over our sample, with at least one significant "jump" around 1980, led us to use the same deflator for all variables.

[17] For most industries, adding further lags of GNP growth did not increase the adjusted R^2 of the regression.

income growth in these industries is idiosyncratic. An implication is that human capital risk for workers in these industries may not be hedged very well with a claim on GNP.

4.1.3 Labor income by educational group

It is likely that the risk/return characteristics of human capital differ across income groups. Campbell, Cocco, Gomes, and Maenhout (1998) have constructed information on income by educational group that will allow us to investigate this hypothesis. They specify a model of income for individual i at date t of the form:[18]

$$\log Y_{it} = f(t, Z_{it}) + \nu_{it} + \nu_{it}$$

where $f(t, Z_{it})$ depends on age and other individual characteristics Z_{it}, ν_{it} follows a random walk with innovation u_{it}, and ε_{it} is a temporary shock. Letting $\hat{f}(t, Z_{it})$ denote the fitted value of $f(t, Z_{it})$, CCGM define

$$\log Y^*_{it} = \log Y_{it} - \hat{f}(t, Z_{it})$$

which effectively purges income of fluctuations due to age or other identifiable demographic characteristics.

CCGM generously provided annual data from 1971-1992 on the average value of Y^*_{jt} for individuals within educational group j. There are three such groups: (i) never finished high school; (ii) finished high school; (iii) additional education beyond the high school level. Table 9 shows the standard deviation of the growth rates of the Y^*_{it} (ΔY^*_{it}) and their sensitivity to GNP growth. The standard deviation of ΔY^*_{it} for the group that never finished high school (group 1) is nearly twice the standard deviation of GNP growth; whereas the highest educational group has ΔY^*_{it} about as volatile as GNP growth. In terms of correlations, ΔY^*_{it} is highly correlated with GNP growth for the lowest educational group, with a correlation of 0.83; for the middle education group the correlation is 0.71, but is only 0.33 for the group with the highest education level. The last line in the table shows the "beta" in a regression on GNP growth similar to the one run for the industry groups above, together with the R^2s for these regressions. The lowest educational group has a large value of "beta" at 1.63, while the highest educational group has a "beta" of only 0.32. Further, the R^2 for the lowest educational group is 0.68, but is only 0.11 for the group with the most education. These results suggest that there are important differences across educational groups in the extent of aggregate vs. idiosyncratic risk in labor income, and thus the extent to which claims on aggregate output could be used to hedge this risk. Specifically, the lowest educational group seems to be the most exposed to aggregate risk, while the highest educational group seems to be the least exposed.

[18] The following uses the notation of CCGM, not the notation of the present paper.

4.2 Hedging human capital risk

This section explores issues associated with hedging human capital risk. While we have documented that labor income is volatile, implying that human capital is risky, it remains to be seen whether existing financial assets could be useful in hedging this risk. Financial assets will be useful as hedges against human capital risk if they have significant predictive power for current and future labor income growth. This sub-section explores whether such predictive power can be found in the data.[19]

Our results on the predictability of labor income growth using financial asset returns are summarized in Table 10. We consider a variety of measures of labor income and several different possibilities for traded financial assets. In Panel A, we consider aggregate, per capita labor income as our measure of the flow of 'dividends' from human capital; the traded financial assets that we consider are U.S. and U.K. equities. Our data is annual, covering the period 1951-1998. We ran a regression of the form:

$$\Delta \log y_t = c + \beta \Delta \log y_{t-1} + \Gamma_{US}(L) r_t^{US} + \Gamma_{UK}(L) r_t^{UK} + \varepsilon_t$$

where we included lags 0 and 1 in the polynomials $\Gamma_{US}(L)$ and $\Gamma_{UK}(L)$. We found that U.S. and U.K. equities each had significant predictive power for aggregate, per capital labor income growth. This suggests an important role for both U.S. and U.K. equities in hedging human capital risk, where the measure of human capital is aggregate labor income. For an individual, this result could be interpreted as implying that U.S. and U.K. equities could be useful in hedging the *aggregate* component of his human capital risk.

The last line of Panel A runs a similar regression for the "Social Security asset"– the rate of growth of the wage indexing series. Here, we find that U.S. equities but not U.K. equities are useful predictors for the growth rate of the wage indexing series.

Panel B of Table 10 looks at predictability of aggregate labor income (not per capita) when the traded assets include U.S. equities, European equities, and Far East equities. The sample is quarterly, 1969:4-1998:1. Again, we find an important role for U.S. equities in hedging aggregate human capital risk. European equities are not significant predictors of aggregate labor income growth, but Far East equities are strongly significant. At least over this time period, then, U.S. and Far East (mainly Japanese) equities would have provided useful hedges for aggregate human capital.

Panel C of Table 10 looks at the predictability of labor income growth in 9 industries, using the same quarterly dataset for financial variables. We find that the U.S. equity return is a significant predictor of labor income growth in 6 of the 9 industries, while European and Far East equities are never significant. The strongest results are for those industries known for their cyclic sensitivity: construction and

[19] Davis and Willen (1998) also address the question of the benefits of using financial asset to hedge human capital.

the manufacturing industries. Our results imply that U.S. equities could provide a useful hedge for human capital risk in several industries.

Finally, Panel D looks at the predictability of labor income growth for the educational groups described above. We find no equity return that is a significant predictor of labor income growth, although including the equity returns in the regressions does increase the fit of the regressions substantially.

Overall, our findings suggest that returns on financial assets, notably U.S. equities, are significant predictors of labor income growth. These results suggest an important role for U.S. equities in hedging the risk of nontraded human capital. Specifically, an individual whose labor income growth is positively correlated with U.S. equity returns would hedge by holding a reduced fraction of his portfolio of traded financial assets in the form of U.S. equities. As a consequence, the individuals holdings of foreign assets (or risk-free assets) would necessarily be increased.

4.3 Computing human capital returns: Literature review

Since human capital is a nontraded asset, measuring human capital returns requires an asset-pricing model that links observable cash flows from human capital to returns on this asset. This section summarizes alternative approaches that have been used in the literature.

4.3.1 The approach of Campbell

Campbell (1996) constructs a multifactor model of asset returns in which risky human capital plays an important role. By assuming that the conditional expected return on human capital is equal to the conditional expected return on financial wealth, Campbell derives the following expression for the return on human capital:

$$r_{y,t+1} - E_t r_{y,t+1} = (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta y_{t+1+j} - (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j r_{a,t+1+j} \tag{4}$$

where $r_{y,t+1}$ is the return on human capital, Δy_{t+1+j} is the growth rate of labor income, $r_{a,t+1+j}$ is the return on financial assets, and ρ is a discount factor. Thus, human capital returns are high when there are upward revisions in expected future labor income growth, but are low when expected future returns on financial assets are high because the given stream of labor income is discounted at a higher rate. Campbell finds that there is important predictability of future labor income growth using current and lagged labor income growth and, in monthly data at least, financial market variables also help predict labor income growth. Campbell also finds that most of the variability in the return on human capital is accounted for by the variance of the 'discount rate news' term, $(E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j r_{a,t+1+j}$. Specifically, the variance of the human capital return in monthly data is reported as 20.8 (percent per month), while the variance of the labor income news is only 1.256. Thus, the correlation between human capital returns and returns on financial assets is necessarily quite high: the correlation is 0.94 for the monthly data, and is 0.54 in annual data.[20]

4.3.2 The approach of Baxter and Jermann

Baxter and Jermann (1997) use a somewhat different approach, working with a version of (4) in which expected returns on financial assets are constant. The starting point for the Baxter-Jermann work is the observation that labor's share of GNP appears to be stationary, so that labor income and capital income are cointegrated. They specify a VECM for labor income growth and capital income growth, as follows:

$$\Delta d_{L,t+1} = \delta_L + \beta_L(L)\Delta d_{Lt} + \gamma_L(L)\Delta d_{Kt} + \alpha_L(d_{Lt} - d_{Kt}) + \varepsilon_{L,t+1} \tag{5}$$

$$\Delta d_{K,t+1} = \delta_K + \beta_K(L)\Delta d_{Lt} + \gamma_K(L)\Delta d_{Kt} + \alpha_K(d_{Lt} - d_{Kt}) + \varepsilon_{K,t+1}\ . \tag{6}$$

where d_{Lt} denotes the log of labor income, d_{Kt} denotes the log of capital income, $\Delta d_{L,t+1} \equiv d_{L,t+1} - d_{Lt}$, $\Delta d_{K,t+1} \equiv d_{K,t+1} - d_{Kt}$, and $\beta_L(L)$, $\beta_K(L)$, $\gamma_L(L)$, $\gamma_K(L)$ are polynomials in the lag operator, L. Returns to labor and capital were then computed as follows:

$$r^L_{t,t+1} - E(r^L_{t,t+1}) = (E_{t+1} - E_t)\left(\sum_{j=0}^{\infty} \rho^j \Delta d_{L,t+1+j}\right) \tag{7}$$

$$r^K_{t,t+1} - E(r^K_{t,t+1}) = (E_{t+1} - E_t)\left(\sum_{j=0}^{\infty} \rho^j \Delta d_{K,t+1+j}\right)\ . \tag{8}$$

Because labor and capital income share a common stochastic trend, and because returns on an asset are dominated by revisions in the expected trend component, labor and capital returns were found to be very highly correlated, with the correlation exceeding 0.92 for the U.S., Japan, and Germany, and a correlation of 0.78 for the U.K. If the return on physical capital is best proxied by the return on the unlevered stock market, as argued by Black (1987), then these results suggest a high correlation between human capital returns and returns on the domestic stock market.

4.3.3 Returns on human capital: Some new evidence

This section uses a version of the approach of Shiller (19xx) and Baxter and Jermann (1997) to compute U.S. aggregate human capital returns where the 'dividend' flow is

[20] In earlier work, Fama and Schwert (1977) study human capital returns in a version of (4) that assumed that labor income growth was unforecastable, and where there were assumed to be no variation in expected future discount rates. These alternative assumptions explain why Fama and Schwert found little relationship between human capital returns and returns on financial assets.

aggregate U.S. labor income. This approach is distinct from Campbell's because it assumes that the discount rate is constant. Thus all variation in returns is due to changes in expected future dividend flows.

We want to explore how useful domestic and foreign assets may be in hedging human capital risk. Intuitively, traded financial assets are useful as hedges for human capital risk if they have returns that are highly correlated with human capital returns. Further, traded financial assets are important determinants of the returns on human capital to the extent that returns on financial assets help predict future growth in labor income. Thus, we estimated a process for income of the form:

$$\Delta \log y_t = c + \Gamma_Y(L)\Delta \log y_{t-1} + \Gamma_{US}(L) r_t^{US} + \Gamma_{foreign}(L) r_t^{foreign} + \varepsilon_t \tag{9}$$

and computed returns on human capital using the formula

$$r_{t,t+1}^Y - E(r_{t,t+1}^Y) = (E_{t+1} - E_t)\left(\sum_{j=0}^{\infty} \rho^j \Delta y_{t+1+j}\right). \tag{10}$$

Table 11 contains information on the return characteristics of various measures of labor income. The structure of Table 11 parallels that of Table 10 in terms of the measure of labor income considered and the menu of financial assets. In each case, we report results for several choices of lag length in the polynomials $\Gamma_j(L)$, since the choice of lag length significantly affects our results in some cases.

Panel A has results for per capita labor income. The standard deviation of the returns on this measure of human capital is about .04 (4% per year) – about 1/3 as volatile as U.S. equities. Our results indicate that human capital returns are strongly correlated with U.S. equity returns and with U.K. equity returns as well; there is some evidence that the correlation with the U.S. equity return is stronger. The results for the Social Security asset are similar, except that the Social Security returns are uniformly more correlated with U.S. equity returns than with U.K. returns.

Panel B contains results for quarterly aggregate labor income and three equity portfolios. We find that this measure of aggregate human capital returns is positively correlated with U.S. equity returns as well as being positively correlated with European and Far East equity returns. There is no clear evidence on which financial asset is most highly correlated with human capital; the results are sensitive to the lag length in the VAR for labor income growth.

Panel C contains results for human capital by industry. In general, human capital returns at the industry are positively correlated with U.S. and foreign equities. The results are quite sensitive to lag length, however, and it is difficult to draw clear conclusions from this table. It appears that those industries that we identified earlier as 'cyclically sensitive:' construction, and manufacturing of durables and nondurables, are the industries with the highest correlation with U.S. equity returns. However, we also estimate high correlation between human capital returns and U.S. equity returns in the wholesale and retail trade industries. Finally, Panel D contains results for the three educational groups. The results here are quite sensitivity to lag length, and the correlations are in some cases negative. Our interpretation is that the very short sample period (24 years, annual data) is insufficient to estimate (9) with any precision.

5 Allocation of retirement assets

Under Social Security, individuals receive a certain real annuity with spousal survival rights. By contrast, in the private variable annuities market, individuals can receive annuity payments which depend on the returns on risky portfolios. Thus, one direction for Social Security reform is to expand the menu of annuities available to Social Security participants. For example, individuals might choose a certain real annuity, as under the current plan, or they might choose a risky portfolio whose return depended on the domestic or international stock markets. Further, they might choose some of each.

In this section, we explore the allocation of retirement assets at the retirement point, with a specific focus on the benefits from investment in international risky assets. We will work with a basic model designed to highlight the issues. We also discuss the implications of extending the model in various directions. The discussion focuses on a household who is at the retirement age, which we view as exogenously determined and denote as R as above. This household can live until a maximum age T.

5.1 Preferences

We assume that households have time separable preferences and that momentary utility is

$$u(c_t, \underline{c}_t) = \frac{1}{1-\sigma}[(c_t - \underline{c}_t)^{1-\sigma}].$$

According to this expression, individuals derive utility from the deviation of the level of consumption, c_t, from a mandated level of consumption, $\underline{c}_t$. We want have written utility in this form for several reasons. First, it allows for a "subsistence" level of consumption, which is an idea frequently discussed in public policy discussions of Social Security and other transfer programs. Second, it is a simplification of the idea of "habitual" levels of consumption that are built up over the prior work years (in this latter interpretation, we might wish to include additional terms to the utility function to represent the utility benefits from these levels). In either case, the parameter σ governs the aversion to risk that individuals have (higher values of σ correspond to higher relative risk aversion) and intertemporal substitution (higher values of σ correspond to a lower elasticity of intertemporal substitution).

We consider a group of individuals that are at the retirement point, age R, at calendar date t. We assume that individuals have a subjective discount factor β. Further, individuals have a probability π_j of living for j periods after retirement, and that they discount future utility flows accordingly. Thus, at the retirement point expected utility is:

$$E_t\{\sum_{j=0}^{T-R} \pi_j \beta^j u(c_{t+j}, \underline{c}_{t+j})\}.$$

5.2 Retirement portfolios with a certain lifetime

The initial focus of our discussion is on optimal portfolio construction when there is a known lifetime. In this context, we are interested in learning

(a) the determinants of the demand for risky assets as a fraction of wealth;

(b) the fraction of this demand which is for international assets;

(c) the likelihood that short-sales constraints bind; and

(d) the cost of restrictions on the composition of the portfolio.

To address these questions, we use to the model sketched above. Each period, the individual begins with a level of wealth, a, and must choose how much to save. We let μ denote the marginal propensity to consume out of wealth. The individual must also decide how to allocate his portfolio: we call the fractions allocated into different assets x, with a subscript indicating the short-term bonds, the domestic risky asset, or the international risky asset.

With a certain lifetime, the form of the utility function is:

$$E_t\{\sum_{j=0}^{\overline{T}-R} \beta^j u(c_{t+j}, \underline{c}_{t+j})\}$$

We find it convenient to describe asset evolution with the following three constraints:

$$a_{t+1} = [x_{bt}(1 + r_{b,t+1}) + x_{dt}(1 + r_{d,t+1}) + x_{it}(1 + r_{i,t+1})]f_t \tag{11}$$

$$1 = [x_{bt} + x_{dt} + x_{it}] \tag{12}$$

$$c_t + f_t = a_t \tag{13}$$

Equation (11) states that the future wealth level is determined by the portfolio allocation $[x_{bt}(1+r_{b,t+1})+x_{dt}(1+r_{d,t+1})+x_{it}(1+r_{i,t+1})]$ and the extent of wealth not consumed $f_t = (1-\mu_t)w_t$. Equation (12) defines the restriction on the portfolio shares (lower case $x's$) and equation (13) is the constraint on consumption and accumulation. The individual holds a portfolio with the return

$$1 + \widetilde{r}_{t+1} = [x_{bt}(1 + r_{b,t+1}) + x_{dt}(1 + r_{d,t+1}) + x_{it}(1 + r_{i,t+1})].$$

Equivalently, this portfolio return is

$$1+\widetilde{r}_{t+1}=[(1+r_{b,t+1})+x_{dt}(r_{d,t+1}-r_{b,t+1})+x_{it}(r_{i,t+1}-r_{b,t+1})]$$

when we impose the requirement (12) that the portfolio shares sum to one.

5.2.1 The standard problem

The foregoing is a version of the portfolio problem that has been much studied in finance (e.g. Levhari and Srinvasan, Samuelson, Hakansson) with two exceptions. In those analyses, there is sometimes a certain labor income stream (an "endowment" of the bond asset over the remainder of the lifetime) and there is no subsistence level of consumption. The results of that literature are as follows. First, in the absence of labor income and subsistence consumption, the marginal propensity to consume which depends on the length of the remaining lifetime but not on the level of wealth. Second, the portfolio shares are independent of the level of wealth *and* the length of the lifetime.

These properties can be demonstrated by studying the dynamic programming problems for households of ages $R, R{+}1, R{+}2$, and so forth. To write these compactly, it is useful to have a notation for the age of the household, which we call γ_t:

$$V(a_t,\gamma_t)=\max_{c_t,x_t}\{u(c_t)+\beta E_t V(a_{t+1},\gamma_t)\}$$

subject to the above constraints, where the subscript on the value function indicates the age of the individual. Attaching a multiplier λ_t to the third constraint and treating the others as equality constraints to be used as desirable, we have the first-order conditions:

$$\begin{aligned}
c_t &: \ Du(c_t)=\lambda_t \\
f_t &: \ \lambda_t=\beta E_t\{(1+\widetilde{r}_{t+1})D_1V(a_{t+1},\gamma_t)\} \\
x_{dt} &: \ 0=\beta z_t E_t\{(r_{d,t+1}-r_{b,t+1})D_1V(a_{t+1},\gamma_t)\} \\
x_{it} &: \ 0=\beta z_t E_t\{(r_{i,t+1}-r_{b,t+1})D_1V(a_{t+1},\gamma_t)\}
\end{aligned}$$

By the "envelope theorem," $D_1V(a_t,\gamma_t)=\lambda_t$, so that the last conditions can be written as:

$$\begin{aligned}
0 &= E_t\{(r_{d,t+1}-r_{b,t+1})Du(c_{t+1})\} \\
0 &= E_t\{(r_{i,t+1}-r_{b,t+1})Du(c_{t+1})\}
\end{aligned}$$

(We are treating x_{bt} as the residual to be determined by $x_{bt}=[1-x_{dt}-x_{it}]$.) Working with these conditions, we find that if $c_{t+1}=\mu(\gamma_{t+1})a_{t+1}$ and $Du(c_t)=c_t^{-\sigma}$, it follows that:

$$0=E_t\{(r_{d,t+1}-r_{b,t+1})[(1+\widetilde{r}_{t+1})\mu(\gamma_{t+1})f_t]^{-\sigma}\}$$

$$0 = E_t\{(r_{i,t+1} - r_{b,t+1})[(1+\tilde{r}_{t+1})\mu(\gamma_{t+1})f_t]^{-\sigma}\}$$

where $(1+\tilde{r}_{t+1}) = [(1-x_{dt}-x_{it})(1+r_{b,t+1}) + x_{dt}(r_{d,t+1}-r_{b,t+1}) + x_{ft}(r_{i,t+1}-r_{b,t+1})]$. These latter equations are important for two reasons. First, they indicate why the $x's$ are independent of wealth. Mathematically, the terms $[\mu_{R+1}f_t]^{-\sigma}$ factor out of the conditions and have no bearing on the optimal x's. Economically, the individual has constant relative risk aversion so that his attitude toward proportional risks–including those on investing in risky assets—is independent of the scale of his consumption. Second, they provide the basis for computing x_d and x_i: simply search to find the pair of portfolio weights that "solves" these equations.

5.2.2 The marginal propensity to consume out of wealth

Given that the portfolio shares are constant through time, we can then proceed as if there were a single asset with a return $(1+\tilde{r}_{t+1})$. The marginal propensity to come out of wealth solves the equation, $\lambda_t = E_t\{\beta(1+\tilde{r}_{t+1})\lambda_{t+1}\}$, so that[21]

$$\mu(\gamma_t) = \frac{[(E_t\{\beta(1+\tilde{r}_{t+1})^{1-\sigma}\})^{-\frac{1}{\sigma}}]\mu(\gamma_{t+1})}{1+[(E_t\{\beta(1+\tilde{r}_{t+1})^{1-\sigma}\})^{-\frac{1}{\sigma}}]\mu(\gamma_{t+1})}$$

To determine the marginal propensity to consume out of wealth, one thus works backward from $\mu = 1$ to the period (age) of interest. The longer the horizon, the lower the marginal propensity to consume. But, while μ depends on the horizon and on the interaction of preferences and returns, it does not depend on the level of wealth.

5.2.3 The growth rate of consumption

The setting implies that the expected growth rate of consumption is constant over time and independent of the investment horizon. Since consumption is proportional to wealth, it follows that:

$$\begin{aligned}
\frac{c_{t+1}}{c_t} &= \frac{\mu(\gamma_{t+1})}{\mu(\gamma_t)}\frac{a_{t+1}}{a_t} \\
&= \frac{\mu(\gamma_{t+1})}{\mu(\gamma_t)}(1+\tilde{r}_{t+1})(1-\mu(\gamma_t)) \\
&= \frac{1}{[(E_t\{\beta(1+\tilde{r}_{t+1})^{1-\sigma}\})^{-\frac{1}{\sigma}}]}(1+\tilde{r}_{t+1})
\end{aligned}$$

[21] The steps in this derivation are to substitute in the consumption decision rule, $c_t = \mu_t w_t$, and the equation for future wealth:

$$\begin{aligned}
(\mu_t w_t)^{-\sigma} &= E_t\{\beta(1+\tilde{r}_{t+1})(\mu_{t+1}w_{t+1})^{-\sigma}\} \\
(\mu_t w_t)^{-\sigma} &= E_t\{\beta(1+\tilde{r}_{t+1})^{1-\sigma}(\mu_{t+1}(1-\mu_t)w_t)^{-\sigma}\}
\end{aligned}$$

and then to solve the resulting expression.

Thus, give the portfolio rules, we can determine the rate at which consumption will grow on average and in response to various disturbances.

5.2.4 The value functions

It is easy to establish that the value functions for this retirement problem take the constant relative risk aversion form,[22]

$$v(a,\gamma) = q(\gamma)^{1-\sigma}\left[\frac{1}{1-\sigma}a^{1-\sigma}\right]$$

with the multiplicative coefficient satisfying the recursive relationship,

$$q(\gamma)^{1-\sigma} = [\mu(\gamma)^{1-\sigma} + (1-\mu(\gamma))^{1-\sigma}E\{\beta(1+\widetilde{r})^{1-\sigma}q(\gamma')^{1-\sigma}].$$

5.2.5 Nontraded, risk-free assets or subsistence consumption

One standard extension of this basic model (e.g. Hakansson) has been to been to treat the individual as having a certain endowment of a nontraded assets. Hakansson discusses this as labor income, but in our context it could equivalently be certain transfer payments that the individual expects to receive (e.g., Social Security). The approach is to treat this nontraded certain income stream as a (positive) endowment of bonds. Then, the bond quantities chosen in the foregoing analysis must be "corrected" for the endowment of the bond-like nontraded asset. A constant, or at least nonstochastic, subsistence level can be analyzed in much the same manner, except that there is a negative endowment of bonds from this source: the individual is required to be able to finance subsistence consumption before anything else (otherwise he faces infinitely negative utility).

Taking both of these features together and calling the nontraded asset's income stream $\underline{b}$, the above analysis would then apply (at the retirement date) to a modified wealth measure,

[22] The algebra is as follows. If we conjecture that to be of a form that is of constant relative risk aversion form,

$$v(a,\gamma) = q(\gamma)^{1-\sigma}\{\frac{1}{1-\sigma}a^{1-\sigma}\}$$

then the value recursion with the optimal consumption decision, $c(a,\gamma) = \mu(\gamma)a$ implies that

$$\begin{aligned} v(a,\gamma) &= \{\frac{1}{1-\sigma}[\mu(\gamma)a]^{1-\sigma}\} + \beta E v(a',\gamma') \\ &= \{\frac{1}{1-\sigma}[\mu(\gamma)a]^{1-\sigma}\} + \beta q(\gamma')^{1-\sigma}E\{\frac{1}{1-\sigma}(a')^{1-\sigma}\} \\ &= \{\frac{1}{1-\sigma}[\mu(\gamma)a]^{1-\sigma}\} + \beta q(\gamma')^{1-\sigma}E\{((1+\widetilde{r})(1-\mu(\gamma))a)^{1-\sigma}\} \end{aligned}$$

so that the conjecture is correct if the following coefficient recursion is satisfied.

$$\widehat{a}_t = [a_t + \sum_{j=0}^{T-R} (\frac{1}{1+r_b})^j (\underline{b}_{t+j} - \underline{c}_{t+j})]$$

That is: the individual's modified wealth measure $\widehat{a}_t$ includes measured financial wealth plus the present value of the nontraded asset's income stream less the present value of the mandated consumption level. This modified wealth measure would be allocated into risky assets and bonds proportionately as described previously, so that the individual's net demand for bonds at date t would be

$$\sum_{j=0}^{T-R} (\frac{1}{1+r_b})^j (\underline{c}_{t+j} - \underline{b}_{t+j}) + x_b \widehat{a}_t.$$

Thus, there would be the proportional demand for bonds discussed earlier, $x_b \widehat{a}_t$, plus some additional bond purchases necessary to cover the excess of required (subsistence) consumption over the income from the nontraded asset.

5.3 Results for specific portfolio opportunities

What retirement portfolios would individuals choose if their opportunities were described by the recent risk and return characteristics of traded financial assets? Table 12 shows how individuals would structure their retirement portfolios, given their degree of risk aversion, σ, and the specific "international" portfolio under consideration. When risk aversion is low, with σ=2, individuals want to hold about 150% of their retirement portfolio in the form of risky assets–they would have to borrow at the risk-free rate in order to accomplish this. As the degree of risk aversion rises, the share of the retirement portfolio invested in risky assets falls: when $\sigma = 5$, the share of risky assets is about 60%, falling to about 20% for $\sigma = 15$. The share of risky assets in the retirement portfolio is not very sensitive to the choice of the international portfolio, if we focus on the 1970-1998 period.

The share of the risky portfolio invested in U.S. assets is between 66% and 77%, depending on the particular international portfolio. Our results indicate that the U.S. share in total risky assets is quite insensitive to the level of risk aversion.[23]

5.3.1 Welfare gains from including equities in retirement portfolios

The prior sub-section showed that there was an important role for risky assets in individuals; retirement portfolios, for all but extremely risk-averse individuals. The current Social Security system is set up so that the portfolio held by a retired person is

[23]This result is related to the standard result of the Markowitz model augmented with a risk-free asset: all investors hold the same risky portfolio, but investors with different levels of risk aversion hold different mixes of the risky and risk-free assets.

essentially a real, risk-free bond. It is natural to ask how much better off individuals would be if their retirement portfolios instead contained risky assets. Of course, the answer will depend on the length of the retirement horizon and the individual's degree of risk aversion.

To answer this question, we proceed as follows. Let a^b denote the asset holdings of a retired person who invests all of his assets in the risk-free bond. For a retirement period of length θ, the value function for this individual was shown to be:

$$v(a^b, \theta) = q^b(\theta)^{1-\sigma} \left[\frac{1}{1-\sigma} \left(a^b\right)^{1-\sigma} \right].$$

Now consider an individual who optimally invest in stocks, and has wealth a^s with value function:

$$v(a^s, \theta) = q^s(\theta)^{1-\sigma} \left[\frac{1}{1-\sigma} \left(a^s\right)^{1-\sigma} \right].$$

We set the value functions implied by the two portfolio strategies equal to each other, $v(a^b, \theta) = v(a^s, \theta)$, and then ask what wealth levels a^b and a^s are implied by this equality. It is intuitive that higher wealth levels will be needed to deliver a specific level of utility (as summarized by the value function) when the individual is constrained to hold only the risk-free bond. Specifically, $V^b = V^s$ implies $a^b/a^s = q^s/q^b$. Another way to state this is to note that allowing an individual to move from a bond-only portfolio to his optimal bond/stock portfolio is equivalent to giving the individual an increase in wealth, in percentage terms, equal to $100 * \left[\left(q^s/q^b\right) - 1\right]$.

Table 13 reports the welfare gains from allowing individuals to move from a bond-only portfolio to the optimal portfolio, as measured by the effective increase in individual wealth. We display the effective increase in wealth for various levels of risk aversion, various lengths of the retirement period, and four different specifications of the available risky assets. The greatest gain in welfare from allowing risky investment naturally occurs for investors with low risk aversion and long retirement periods. For this group, the increase in effective wealth from allowing risky investment ranges from 55% to 75%. However, we find that there are notable, effective increases in wealth even for investors with higher risk aversion or shorter horizons. The effective increase in wealth drops below 5% only when the retirement horizon is less than 10 years, or when risk aversion exceeds $\sigma = 10$. An increase in retirement wealth between 5% and 75% would certainly be considered economically important by most households. We therefore conclude that there are sizable benefits to including risky assets in the retirement portfolios of most individuals.[24]

[24] Remember that these welfare gains only apply to allowing a specified asset pool to be invested in risky assets at the retirement point. This computation does not consider the potential for risky investment during the working years to deliver a larger value for the individual's assets at the retirement point.

5.3.2 Welfare costs of restrictions on retirement portfolios

Many proposals for Social Security reform call for investors to have the opportunity to invest in risky assets. These proposals differ, however, in their recommendations concerning the menu of risky assets available to individuals. The Feldstein/Samwick proposal, for example, calls for individuals to invest in a diversified portfolio of U.S. assets. The Kotlikoff/Sachs proposal, on the other hand, calls for retirement funds to be invested in a diversified world portfolio. This sub-section investigates the welfare implications of restricting investors to holding only U.S. assets in the risky part of their retirement portfolio. We also rule out short sales of all assets, including the risk-free bond.

Table 14 shows how these restrictions affect the welfare of the retired investor. We begin by noting that the "no short sales" constraint does not bind except for the least risk-averse investors ($\sigma = 2$). Panel A shows that the fraction of the portfolio invested in risky securities falls when investors are constrained to hold only U.S. securities: the $\sigma = 3$ investor reduces his risky holdings from 105% of his assets to 93% of his assets; the $\sigma = 10$ investor reduces risky holdings from 31% of assets to 28% of assets.

The reduction in welfare from the portfolio restrictions are shown in Table 14-B. The good news from Table 14-B is that allowing individuals to invest only in U.S. equities (as well as bonds) would deliver substantial welfare gains to retired individuals, relative to a situations in which retirees are required to hold just a risk-free bond. The bad news is that restricting the portfolio to only U.S. assets means that individuals receive only about 80% of the welfare increase that could be gained by permitting international diversification. These results reinforce the impression from Figure 3, which showed great improvements in the risk-return tradeoff from combining traditional Social Security with U.S. bonds and equities. It was unclear from that figure how important adding international investments might be in terms of generating increased welfare; this analysis suggests that the marginal contribution to individual welfare from international investment could be quite important to a retired individual.

5.4 Retirement portfolios with an uncertain lifetime

We now turn to investigating a situation in which individuals have uncertain lifetimes. We assume that individuals can purchase annuities from a competitive, zero cost insurance company. There are two effects of this modification, relative to our prior analysis of the certain-lifetime case. First, in the presence of less than certain survival, individuals effectively discount the future more heavily and wish to tilt their consumption profiles toward the present. Second, in the presence of less than certain survival, insurance companies price the longer term components of annuities more cheaply, thus encouraging individuals to substitute toward the present. When annuities are priced in an actuarially fair manner, as we will assume, then these two substitution responses offset each other. Individuals then can simply purchase a higher level of consumption in all periods in which they are alive. But the portfolio shares of risky assets–held by the insurance company in response to the preferences of households–are unaffected by uncertain lifetimes.

We assume that there are insurance companies which offer annuities which are based on the same three assets discussed above. Suppose that a competitive, zero cost insurance company faces a population of individuals with actuarial survival probabilities π_j. Then the insurance company could offer certain, real annuities (backed by bonds) to retirement-age individuals. In particular, the insurance company would offer annuities which took the following general form to an individual retiring at date t, considering the allocation of retirement wealth at t+1 so as to support consumption in the $T-R$ periods $t+1$, $t+2$,... $t+T-R$. The insurance company will provide the individual with a specified pattern of consumption beginning at $t+1$ (conditional on survival) and lasting for the remainder of his retirement life. In return, the individual or his estate would turn over

$$\sum_{j=1}^{T-R} \pi_j (\frac{1}{1+r_b})^{j-1} c_{t+j}$$

units of bonds to the insurance company at date $t+1$. That is: the price of the annuity increases with the survival probabilities and the level of real payments, but falls with the real interest rate.

Panel A of Figure 4 shows the US survival probabilities, conditional on attaining 65 years of age, based on the general population.[25] The expected lifetime for such an individual is about 17.5 years. We will use these survival probabilities in all of our computations below, although it would be interesting to explore how certain computations would differ for subgroups of the population that have alternative survival probabilities. To get a feel for the relevant annuity factors under actuarial fair insurance and at various real interest rates, we produced table 15.[26] This table shows that a real interest rate of 2% and the US survival probabilities imply that the price of an annuity paying a level real stream is \$14.33 per dollar of annuity payment. The table also reports prices of real annuities that grow at various rates.

To help us later think about more complicated annuities that the insurance company might offer, it is useful to think about the insurance company's balance sheet under the assumption that such riskless annuities are "backed" by insurance company investments in the safe bond. The insurance company would begin with b dollars per investor at date $t+1$ and it would pay each of the surviving investors c at this

[25] Source:

[26] The elements in this table are

$$\sum_{j=1}^{T-R} \pi_j (\frac{1+\gamma}{1+r_b})^{j-1}$$

for various real grow rates and real interest rate, using the U.S. survival probabilities.

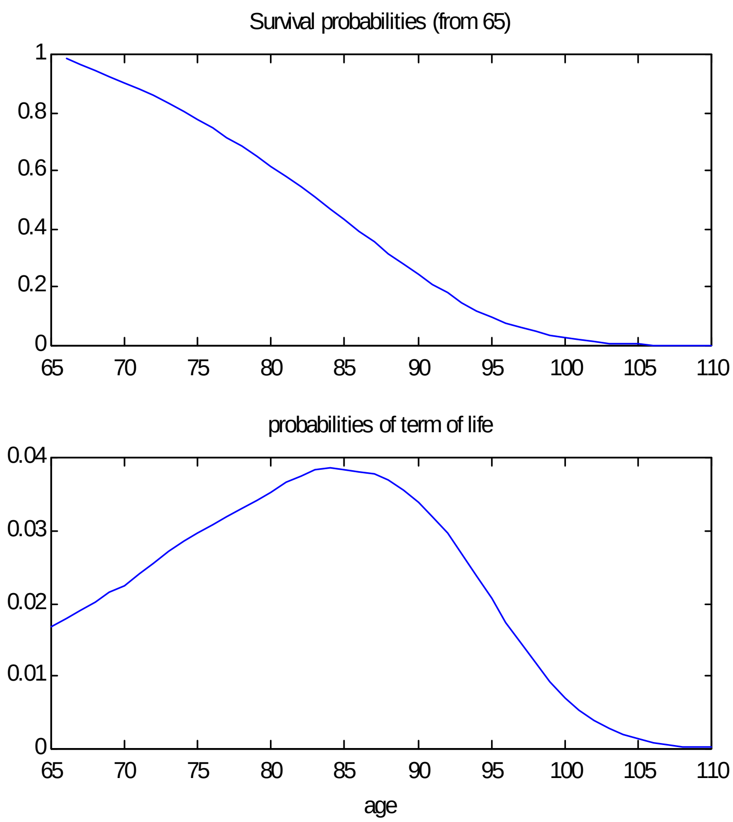

Figure 4: U.S. Survival Probabilities

date, so that it would have $b_{t+2} = (1+r_b)(b_{t+1} - \pi_1 c_{t+1})$ in the next period (period $t+2$). In later periods, the evolution of the insurance company's bonds would satisfy $b_{t+j+1} = (1+r_b)(b_{t+j} - \pi_j c_{t+j})$.[27]

More generally, such an insurance company could hold any underlying assets and make any feasible payment pattern to its group of retirees. For example, if the insurance company held a portfolio of bonds, domestic assets and risky assets which bore random return $\widetilde{r}$, then the disbursements to the annuity holders would be constrained by the sequence of wealth accumulation constraints $a_{t+j+1} = (1+\widetilde{r}_{t+j+1})(a_{t+j} - \pi_j c_{t+j})$. We can accordingly think about the optimal annuity package that our household would select given its utility function.

5.4.1 Annuity demand with only mortality risk

To understand the demand for annuities, we begin by considering the case in which there is only survival risk but no rate of return risk. In return for the individual's accumulated retirement assets (a units of wealth), the insurance company is willing to supply any pattern of consumption over time which satisfies the constraint

$$a_{t+1} = \sum_{j=1}^{T-R} \pi_j (\frac{1}{1+r_b})^{j-1} c_{t+j}$$

This individual will choose the time pattern of consumption to maximize

$$\sum_{j=1}^{T-R} \pi_j \beta^{j-1} u(c_{t+j}, \underline{c}_{t+j})$$

so that the demands for consumption at each date satisfy the first order conditions

$$\beta^{j-1} \pi_j [(c_{t+j} - \underline{c}_{t+j})^{-\sigma} = \Lambda \pi_j (\frac{1}{1+r_b})^j \text{ for } j = 1, 2, ...T-R$$

where Λ is a multiplier on the present value budget constraint. These conditions, together with the budget constraint, determine the desired annuity profile. Since the annuities are actuarially fair and since the interest rate is constant over time, there is a constant growth rate of "net" consumption, determined as $\gamma = [\beta(1+r)]^{1/\sigma}$. Using this result, the annuity path then can be shown to be:

$$c_{t+j} = \underline{c}_{t+j} + \gamma^{j-1} \frac{[a_{t+1} - \sum_{j=1}^{T-R} \pi_j (\frac{1}{1+r_b})^{j-1} \underline{c}_{t+j}]}{\sum_{j=1}^{T-R} \pi_j (\frac{1+\gamma}{1+r_b})^{j-1}}$$

This expression shows that there is a natural modification of the standard life-cycle model when there is an actuarially fair annuities market. First, in the absence

[27] Imposing $c_{t+j} = (1+\gamma)^{j-1} c_{t+1}$ and solving this difference equation to the terminal condition $b_{T+1} = 0$ then leads to the present value annuity formula reported in the text.

of subsistence consumption, we have the standard Fisherian result that the growth rate of consumption is positively influenced by the gap between the rate of time preference and the real interest rate, with the intensity of this effect governed by the intertemporal rate of substitution ($1/\sigma$). Second, the level of consumption per unit of wealth is raised by mortality risk. [28] Third, subsistence consumption has two effects: (i) it acts to reduce the wealth that can be allocated to support other consumption; and (ii) it raises the annuity that is required at a given date.

5.4.2 Demand for variable payment annuities

The demand for annuities with variable payments, based on the returns on risky investments, combines portfolio management with the sharing of life risks. There are a variety of ways of determining the demand for variable payment annuities, but we find it simplest to make use of the analytical results on optimal portfolios under certain lifetimes as derived in section 5.1 above.

To begin, we note that the survival probabilities conditional on reaching $R = 65$ as shown in panel A of Figure 4 imply a conditional distribution of the length of lifetime, as shown in Figure 4-B. Let χ_j be the probability of living for j years past R, which can be formed as $\chi_1 = 1 - \pi_1$, $\chi_2 = \pi_1 - \pi_2$, ... $\chi_{T-1} = \pi_T - \pi_{T-1}$ and $\chi_T = \pi_T$.

We have already learned about the value functions which describes an individual with a certain j periods of life, since these were intermediate products in our analysis above. Age was an important state variable in the value function only because it stood in for the length of life time. The date $t+1$ expected utility for an individual who knows he is going to live j periods takes the form:

$$v_j(a_{j,t+1}) = \frac{1}{1-\sigma}[q_j a_{j,t+1}]^{1-\sigma}$$

with $a_{j,t+1}$ being the wealth of this agent and q_j satisfying the recursion ${q_j}^{1-\sigma} = [{\mu_j}^{1-\sigma} + (1-\mu_j)^{1-\sigma} E\{\beta(1+\widetilde{r})^{1-\sigma} {q_{j-1}}^{1-\sigma}]$ so that it depends explicitly on the parameters of preferences (β, σ) and implicitly on portfolio allocations (through $\widetilde{r}$).

These agents have no disagreements about how wealth should be invested: they all want the same shares in risky and riskless assets. Accordingly, we can write expected utility at the retirement funds allocation point as:

$$\sum_{j=1}^{T-R} \chi_j v_j(a_{j,t+1})$$

Now, of course, the v_j do not really apply to different agents but instead to different lifetimes that the individual may experience. The individual must allocate his wealth

[28] This can be seen by noting that $\sum_{j=1}^{T-R} \pi_j (\frac{1+\gamma}{1+r_b})^{j-1}$ is less than the comparable expression that prevails without mortality risk, $\sum_{j=1}^{T-R} (\frac{1+\gamma}{1+r_b})^{j-1}$, so that the level of consumption is higher.

to these different lifetimes, which involves a budget constraint of the form

$$\sum_{j=1}^{T-R} \chi_j a_{j,t+1} \leq a_{t+1}$$

An efficient allocation across lifetimes implies that

$$\chi_j \Lambda = \chi_j D v_j(a_{j,t+1}) = \chi_j q_j^{1-\sigma}(a_{j,t+1})^{-\sigma}$$

with Λ being the multiplier on the wealth allocation constraint. These conditions then imply that $a_{j,t+1} = [q_j^{\frac{1-\sigma}{\sigma}} / \sum_{j=1}^{T-R} \chi_j q_j^{\frac{1-\sigma}{\sigma}}]a_{t+1}$ so that wealth is allocated proportionately to various lengths of lifetime. A more familiar risk-sharing condition is that consumption is smoothed across various contingencies. With a little algebra, we can confirm that $c_{j,t+1} = \mu_j a_{j,t+1}$ is independent of j and equal to

$$c_{t+1} = [\frac{1}{\sum_{j=1}^{T-R} \frac{\chi_j}{\mu_j}}]a_{t+1}$$

Thus, there is a modified annuity factor under uncertainty, which is a natural generalization of several other constructs. First, if there is a certain lifetime, the annuity factor is just the marginal propensity to consume μ. Second, if the marginal propensity to consume in each lifetime plan is equal to the length of the lifetime, $\mu_j = \frac{1}{j}$, then the annuity factor is just the expected lifetime, $\sum_{j=1}^{T-R} j\chi_j$.

With annuitization of retirement wealth, the effects of portfolio distortions can be explored in the same manner that we used in our analysis of certain lifetimes above. These results are reported in Table 16. Since the life table used in constructing Figure 4 implies that the expected lifetime is about 17.5 years, the results are in between those for certain retirement periods of 15 and 25 years.

6 Conclusions

To be added.

References

[1] Baxter, Marianne and Urban J. Jermann (1997) "The international diversification puzzle is worse than you think," American Economic Review.

[2] Campbell, John Y. (1996) "Understanding risk and return," Journal of Political Economy 104;2 (month 1996), 298-345.

[3] ________________ (1998) Handbook chapter

[4] Davis, Steven and Paul Willen (1998) "Using financial assets to hedge labor income risks: Estimating the benefits," manuscript, University of Chicago, December 1998.

[5] Fama, Eugene and William Schwert (1977)

[6] Feldstein, Martin and Andrew Samwick (1998)

[7] Friedman, Benjamin and Mark Warshawsky (1990)

[8] Friend, Irwin and Marshall E. Blume (1975) "The demand for risky assets," American Economic Review 65:5 (December 1975), 900-922.

[9] Gustman, Alan and Thomas Steinmeier (1997) "Effects Of Pensions On Savings: Analysis With Data From The Health and Retirement Study"

[10] Gustman, Alan and Thomas Steinmeier (1998)

[11] Hakansson

[12] Kennickell, et al. (1997)

[13] Levhari and Srinvasan

[14] Miron, Jeffrey and David Weil (1997)

[15] Mitchell, Olivia; James Poterba and Mark Warshawsky (1997)

[16] Samuelson

[17] Sargent, Thomas (1998) Comment on "Simulating the Privatization of Social Security in General Equilibrium" by Laurence J. Kotlikoff; in: Martin Feldstein, ed., Privatizing Social Security, University of Chicago Press, 1998.

[18] Shiller, Robert J. (199x) Macro Markets