Extended Abstract

The investment choice of savings is among the most important decisions over the life cycle. Not surprisingly, a considerable literature has examined how households should optimally allocate their savings between stocks and bonds over their lifecycle. Samuelson (1969) and Merton (1969) first addressed this problem in discrete and continuous time, respectively, assuming complete markets without labor income. Merton (1971) then introduced deterministic labor income into the continuous time framework while Bodie, Merton and Samuelson (1992) allowed for it to be elastically supplied. In more recent years, attention has shifted toward examining how uninsurable labor income risk and uncertain lifespans impact savings and portfolio allocations over the life cycle.

It is well known, however, that the standard expected-utility model does not appear to closely match the empirical evidence on portfolio choice. The actual empirical evidence suggests three "stylized facts:" (I) the share of a household's portfolio invested in equities is much less than 100% for most households; (II) the lifetime poor invest in fewer equities than richer households; and (III) the share of portfolio invested in risky assets tend to be "hump shape" (∩) in age.(Amerkis and Zeldes 2000; Heaton and Lucas 2000; Poterba and Samwick, 2002).In contrast, the standard model with homothetic preferences tends to generate a very high stock allocation across all income classes that declines with age as human capital depreciates; in many portfolio choice models, the allocation is counterfactually "U shaped" in age.

Several modifications to the standard model have been proposed in the past to deal with at least one of these empirical observations: liquidity constraints (Brown 1990; Amerkis and Zeldes 2000); saving for illiquid assets such as a house (Faig and Shum 2002); habit persistence (Polkovnichenko, 2007); and, incomplete trading markets between generations in a real business cycle economy where wages and stock returns are perfectly correlated (Storesletten, Telmer, and Yaron, 2007). In contrast, the current paper preserves the standard lifecycle expected utility model with homothetic utility but simply considers the importance of a more detailed fiscal setting: a progressive and wage-indexed social security system.

In the United States, a retiree's initial Social Security benefit level is calculated by first multiplying each previous wage earned by the participant with a wage-indexed factor. This factor tracks the average wage growth experienced across the entire economy since the time that particular wage was earned. For example, suppose that you retired today and you earned $30,000 twenty years ago. Also suppose that wages, averaged across the entire economy, grew on average of 5% (nominal) per year during the past 20 years. Your recorded wage of $30,000 would, therefore, be adjusted upward by a factor of $1.05^{20}=2.65$, or equal $79,598.93. A similarly constructed wage factor would be applied to each of your other previous wages as well. Your individual Average Indexed Yearly Earnings (AIYE) would then be determined by taking the average of your adjusted wages.

Your AIYE is then used to calculate your Primary Insurance Amount (PIA) using a nonlinear (progressive) function. Your PIA forms the basis of your Social Security benefit amount before some additional adjustments are made (e.g., a spousal benefit). Individuals with a larger AIYE receive a larger PIA,

thereby recognizing that they paid more into the system. However, the PIA / AIYE ratio -- the "replacement rate" -- is a declining function of the AIYE so that the lifetime poor receives a relatively larger replacement rate of their previous earnings.

Empirically, the Social Security wage index factors, which are used to adjust previous earnings, are highly correlated with stock returns. At first glance, this fact might be surprising since it is well known that average wages and stock returns are not highly correlated at "high" (annual) frequency even at the sectorial level (e.g., Davis and Paul 2000). However, there are two distinct differences that creates a higher correlation between the Social Security wage index factors and stock returns. First, "idiosyncratic" (including sector-specific) risks tend to cancel inside of the economy-wide average index wage factor that is used to adjust previous wages for calculating Social Security benefits. Second, and more importantly, this indexation is mostly calculated at a "low frequency," especially for previous wages earned early in the lifecycle. The economy-wide average wage and stock returns appear to be much more highly correlated at this lower frequency (Jermann, 1999).

Wage indexation has presumably not been considered in the previous literature due to the additional computational complexity it brings by adding an additional state variable. We show, however, that a wage-indexed social security system can potentially play an important role in reconciling all three stylized facts noted above. First, a wage-indexed social security system provides a retirement asset that is highly correlated with stock returns, thereby reducing the demand for stocks by most households (stylized fact (I)). Second, this effect is especially strong for lower-income households who derive most of their retirement income from social security due to its progressive benefit structure (stylize fact (II)). For poor households (many of whom don't have access to employer-based defined-contribution retirement plans), we also investigate how a realistic minimum investment account requirement in stocks reinforces the pattern observed in our model. Third, the correlation between social security and equity returns is larger at younger ages, reducing the demand for equity. As benefits begin to accrue at higher ages and the horizon toward retirement becomes shorter, social security benefits begin to become more substitutable with bonds, thereby increasing the demand for equities. Closer to retirement, enough human capital has depreciated, making stocks less attractive again (stylized fact (III)).

One might object to our positive findings by arguing that few people actually understand the complexities of how their social security benefits are calculated. We tend to agree. However, our results could also be interpreted from a normative angle as well. In particular, from an empirical perspective, the average household might be already holding something close to the optimal portfolio, at least prior to retirement. Immediately after retirement, our simulated optimal portfolio, however, sometimes diverges from the empirical data: the optimal portfolio actually again increases for almost a decade. These results are potentially consistent with retirees, with a large amount of wealth at risk, being more sensitive to traditional financial planning advice than younger people with less wealth at risk. Indeed, our results suggest that standard portfolio advice might be completely wrong on both sides of the age spectrum: younger people should devote a smaller percentage of their wealth to equities and retirees should increase their equity holdings.