Insuring Consumption Using Income-Linked Assets*

Andreas Fuster[†]
Harvard University and
Federal Reserve Bank of Boston

Paul S. Willen[‡]
Federal Reserve Bank of Boston
and NBER

October 6, 2009

PRELIMINARY AND INCOMPLETE

Abstract

Shiller (2003) and others have argued for the creation of financial instruments that allow individuals to insure risks associated with their lifetime labor income. In this paper, we analyze two such instruments in the context of a realistic life-cycle model of household portfolio choice. The instruments we consider are an income-hedging instrument (a limited liability asset whose returns correlate negatively with income shocks) and income-linked loans (with a rate positively correlated with income shocks). We find that income-linked loans would generally be more useful to households, as they have a lower (opportunity) cost. While for some parameterizations of our model the welfare gains from the presence of income-linked assets can be substantial (above 1% of certainty-equivalent consumption), the assets we consider can only mitigate a relatively small part of the welfare costs of labor income risk over the life cycle.

*We are grateful to John Campbell, Nicola Fuchs-Schündeln, Rustom Irani, María Luengo-Prado, Julio Rotemberg, and seminar participants at Harvard and the 1st Annual Workshop of the Zurich Center for Computational Financial Economics for helpful comments and discussions. The views expressed in this paper are solely those of the authors and not necessarily those of the Federal Reserve Bank of Boston or the Federal Reserve System.

†Address: Department of Economics, Harvard University, Littauer Center, 1805 Cambridge St., Cambridge, MA 02138. Phone: (857) 928-1122. E-mail: afuster@fas.harvard.edu

‡Address: Federal Reserve Bank of Boston, 600 Atlantic Ave., Boston, MA 02210. Phone: (617) 973-3149. Fax: (617) 973-2123. E-mail: paul.willen@bos.frb.org

1 Introduction

The problem of smoothing household consumption fluctuations lies at the heart of much public policy. A wide range of government programs and institutions, from central banks to unemployment insurance to Temporary Assistance for Needy Families, owe their existence ultimately to the goal of reducing household consumption volatility. In spite of these efforts, household consumption volatility remains significant.

In this paper, we analyze a market-based approach to reducing consumption risk: financial assets with payoff tied to households' labor income realizations or, as we call them, income-linked assets. We investigate the use of such assets in the context of a calibrated life-cycle model of household consumption in which households can invest in different assets and can borrow but at a substantial premium to the riskless rate. We find that the benefits of income-linked assets can be substantial but are highly sensitive to the precise design of the assets. For example, the seemingly innocuous decision of whether you link income negatively to the return on a savings instrument or positively to the interest rate on a loan has a large impact.

A significant proportion of observed household consumption volatility is due to idiosyncratic, not aggregate income fluctuations. This is a symptom of limited risk-sharing, which can be seen either as a puzzle or as evidence of frictions, depending on which economist you ask. Chief among the frictions is asymmetric information, in particular moral hazard: if one knew that one's consumption is independent of one's income, there is of course no longer a strong incentive to expend effort on trying to avoid negative income shocks such as being fired for shirking.

But there is in principle scope for sharing part of one's labor income risk without inducing moral hazard, because part of this risk is *group-specific*. For instance, an auto worker may receive an income shock because of his performance on the job, but also because of the overall evolution of the demand for cars affecting the average wage of car workers, which is beyond his control. Shocks of the second kind are observable and verifiable, so that insurance-like contracts based on them can in principle easily be written. Indeed, Attanasio and Davis (1996) argue for the "puzzle" view of the limited risk-sharing in the data precisely because they find that a particular group-level shock – income variation at the education-sex-birth cohort group – appears not be shared across households. One reason why sharing of such risks that are immune to moral hazard may be difficult (also discussed by Attanasio and Davis) is the possible importance of another friction, namely limited commitment: if two people sign a contract in which they promise to share their respective incomes with each other, the one receiving a high positive income shock may be tempted to renege on his promise, even if he afterwards gets punished by being excluded from future similar transactions.

This issue can be mitigated by moving from bilateral contracts to trading through long-lived institutions who hold a diversified portfolio and have strong reputational concerns so that they are very unlikely to default (or if they do, they may be backed by the government). This makes the market-based approach considered in this paper a potentially promising approach for sharing household income risks.

The income-linked assets we consider take two basic forms. The first is a standard insurance-like contract in which a household pays something now for an asset whose future payoff is negatively correlated with income – we call this an "income-hedging instrument." The second form we consider is somewhat more creative, and consists of loans where the required repayment is tied to income shocks – "income-linked loans." Either way, the upshot of adding such assets to a household portfolio would be to reduce consumption fluctuations. Our goal is to quantitatively evaluate what the demand for such assets would be, and which design features of the assets this most strongly depends on. Furthermore, we can then look at the welfare gains that the presence of such assets would generate for households.

Our undertaking, and the assets we consider, are inspired by two thought-provoking books by Robert Shiller (1993, 2003) in which he argues for the development of new household risk management instruments. He also furnishes the motivation for our study, as he writes that "Imagining the social and economic achievement that could come from a new financial order is difficult because we have not seen such an alternate world."[1] Of course, our model-based approach is precisely an attempt to predict what might happen in an alternate world. Understanding the potential gains from such assets, and what they depend on, is also important from a policy perspective, because Shiller argues that we require a concerted effort from the government and the private sector to facilitate the introduction of such assets.

To evaluate the demand for and the usefulness of the income-linked assets, we embed them in a realistic portfolio choice problem. We build on the work of Cocco, Gomes, and Maenhout (2005), Davis, Kubler, and Willen (2006) and others and use a finite horizon, partial equilibrium model which roughly matches basic facts about household risky asset holdings. Households receive stochastic labor income, which is subject to permanent and transitory shocks, and they can invest in bonds and stocks. Furthermore, they can also engage in unsecured borrowing at an interest rate that exceeds the return on the riskless bond.

A major challenge for this research is the need to make assumptions about the return characteristics of non-existing assets. For the mean returns of the income-linked assets, we make the baseline assumption that the risks upon which the payoffs

[1] Shiller (2003), p.10.

are based are purely cross-sectional, such that the assets can be priced fairly. Thus, we assume that the mean return on the IHI equals the mean return on the risk-free bond, and that the mean interest rate to be paid on the ILL is equal to the interest rate on other unsecured debt.[2] For the other return characteristics, we remain relatively agnostic and simply plug in different values for the volatility of the IHI and ILL returns and their correlation with the permanent shock to a household's labor income. We do, however, present some back-of-the-envelope calculations that lead us to adopt as our baseline assumption a correlation of 0.5 between individual permanent labor income shocks and the returns on income-linked assets.

Our model yields two main results. The first is that potential welfare gains are strongly convex in the assumed correlation between rates and income shocks. As a consequence, unless the correlation were very high, the income-linked assets can only eliminate a rather small part of the welfare cost imposed by income shocks over the life cycle. Furthermore, we find that the welfare gains that our model predicts are also sensitive to other details of the environment. The size of the cost differential between borrowing and lending, for instance, is very important: the larger it is, the less households gain from having access to the proposed income-linked assets. The risk-adjusted return, and hence the attractiveness of our assets likewise depends on the assumed return volatility, with higher volatility providing "more bang for the buck" for households.

The second main result is that income-linked loans are generally much more appealing and useful to households than the income-hedging instrument. For a baseline calibration in which the correlation between permanent income shocks and the interest rate on the income-linked assets is 0.5, and the volatility of the rate is 0.5, we find that income-linked loans would produce a welfare improvement of 1.4% (an increase in consumption of about USD 400 per year, in 2009 dollars) while the income-hedging instrument is essentially worthless. We also explore the boundaries of this result. We show, for instance, that the attractiveness of the alternative investment possibility matters for the relative appeal of the two income-linked assets: having the possibility to invest in equity (which is the case in our baseline) makes the income-linked loan relatively more attractive (as households can invest some of the borrowed money in a

[2] An alternative to our partial equilibrium approach would be to build a general equilibrium asset-pricing model to generate prices for the assets, but the well-documented problems with such models in generating prices even in-sample means that mis-specification of either the household decision problem or the general equilibrium could lead to inaccurate predictions about the benefits of the assets. As discussed below, another advantage of the partial equilibrium setting is that it allows us to explore the effects of different assumptions about the prices of other assets on the benefits of income-linked assets.

high return asset) while the income-hedging instrument is less demanded. If a household has access to borrowing at a cheap rate (below what it would have to pay on average on the income-linked loan), it may even be that the ranking of the two assets is reversed and that the income-hedging instrument is predicted to be more valuable. However, even under such assumptions, the gains from the income-hedging instrument remain moderate (below 1%). In sum, we find that under some assumptions, income-linked loans could be very useful to households, while it is more difficult to come up with a scenario in which income-hedging instruments would have an equally positive effect on welfare.

To understand these results, we turn to the theory of portfolio choice in the presence of constraints and focus on the risk-adjusted returns on assets. The higher the correlation of an asset with household income, the lower the risk-adjusted return on that asset. Thus, the negative correlation of the income-hedging instrument raises the risk-adjusted return and makes the asset more attractive to investors than a risk-free asset with the same mean return. Similarly, income-linked loans have a lower risk-adjusted cost of funds than borrowing at a fixed rate. We show that over the life cycle, the lure of borrowing at a lower risk-adjusted cost is more attractive to most households than the increased-return investing, so that they would use income-linked loans much more extensively than income-hedging instruments. Early in life, the reason is that most household's main financial activity is high-interest rate borrowing, to which income-linked loans provide a lower-cost alternative. To attract interest, the income-hedging instrument, in contrast, would need to offer a risk-adjusted return that exceeds the cost of unsecured debt. Later in life, the competition for funds comes not from high interest borrowing but from high equity returns. Then, the risk-adjusted return must exceed the risk-adjusted return on equity, and if we set expected returns on equity to match historical averages, that is a tall order as well. Meanwhile, the presence of high-return equity makes the income-linked loans relatively more attractive to households, as it means that they can insure at relatively low cost by taking out an income-linked loan and investing most of it in equity.

We now turn to the obvious question of why we observe little trade in the real world in income-linked loans if the benefits are as substantial as predicted by (at least some calibrations of) our model.[3] We first re-iterate that we assume here that the risks households are hedging are both observable and cross-sectional. The former implies

[3] There do exist some examples of income-linked loans in the real world. For instance, in Australia and some other countries, there exist education loans for which the necessary repayment is based on subsequent labor income. In the US, there have been recent examples of car firms (Hyundai, Ford) offering to make car payments for up to a year and/or take the car back with no loss in equity in case the buyer loses his job.

that there are no adverse selection or moral hazard problems and when combined with the latter, means that we can set the prices of the assets equal to their expectations, which we do. One can view this as an extreme assumption which stacks the deck in favor of income-linked assets and makes the failure of the income-hedging instrument even more surprising than it already is.

Shiller advances another reason for the current non-existence of the income-linked assets he envisions. He argues that until recently we did not have the technology necessary to collect and maintain the data underlying the various instruments. He points out that there will likely be a need for government help to establish what he calls "Global Risk Information Databases," and that new regulations may need to be enacted in order to make the "New Financial Order" possible, as for instance a change in the bankruptcy law, such that income-linked loans could not be canceled by declaring personal bankruptcy.

Finally, we remind the reader that firms traded few financial derivatives forty years ago. Interest rate swaps, exchange rate options and fed funds futures are traded today in highly liquid markets and form a central part of risk management for financial and non-financial firms, but they were only a theoretical concept to the typical corporate treasurer in 1965.

The rest of this paper is organized as follows: In the next two subsections we discuss Shiller's book and some of the other related literature, respectively. We then turn to a two-period model, in order to explain the theory of portfolio choice under constraints and to provide intuition for our results on the use of and gain from the assets we introduce. Section 3 then describes our life-cycle model and the quantitative results we obtain from it, which are further discussed in Section 4. Finally, Section 5 briefly concludes.

1.1 Shiller's *New Financial Order*

In his 2003 book, Shiller proposes various types of insurance or financial instruments that should be developed in order to help households manage and reduce risk. In this paper, we will be concerned with (examples of) a subset of them, namely "livelihood insurance" (which, in our implementation, can also be seen as a particular example of a "macro market"), and "income-linked loans."[4] These instruments are supposed to provide a hedge against fluctuations in a household's income, in particular permanent ones. Of course, Shiller is aware of the moral hazard issues that would arise if, for

[4] Other suggestions he makes include "home equity insurance" (which is now arguably available, through an exchange-traded product based on the Case-Shiller index), "inequality insurance", and "intergenerational social security".

instance, one were completely insured against a decline in one's personal income. Therefore, he proposes that the interest rates of these instruments depend partly or fully on an *index* of income in one's occupation, industry, or region (we will refer to this as an individual's "group"). That way, each individual's labor supply and effort decisions are (mostly) undistorted, but they are insured to some extent against shocks that affect the group they are in.

Shiller's idea of livelihood insurance is related to conventional insurance in the sense that people would be required to pay for a policy upfront which would then insure them against a decline (of a certain magnitude) in the income of their group, in which case the insurance would pay a stream of income as long as the index stays below the threshold. A related idea is to introduce markets for income by group (an example of a "macro market"), which could be conceived as (inverse) exchange-traded funds that track income indices and pay off more if the index goes down.[5] The "income-hedging instrument" we introduce in our model can be seen as encompassing both these ideas. It is a bond-like asset that households can invest in which tends to pay off more if they receive a negative income shock than if they receive a positive one (we say "tend" because the correlation is typically imperfect). A livelihood insurance contract can be seen as a special case of this, where the payoff of the asset is zero unless a certain threshold is hit by the underlying index, in which case the payoff is high. In that sense, livelihood insurance can be thought of as a volatile version of our income-hedging instrument.

Another related proposal is to introduce income-linked loans, whose repayment terms for an individual depend on the evolution of the income index of this individual's group (or perhaps partly also on the individual's income, as long as this does not lead to excessive moral hazard). If an individual's group is doing well, the individual will have to repay more than if his group is doing poorly. Shiller realizes that this may lead to another problem, namely that people may be tempted to default on their income-linked loans after they receive a positive income shock. As a consequence, he argues that bankruptcy law would need to be changed, such that income-linked loans, like student loans, cannot get canceled by declaring personal bankruptcy.

Our model focuses on the risk management benefits that such assets would have, taking the riskiness of household income as given. Shiller discusses some additional benefits that these instruments might have which are not part of our analysis. First, he thinks that the prices of these instruments (for instance, the borrowing rates on income-linked loans that the market offers to different professions) might aggregate

[5]Shiller notes that livelihood insurers and institutions making income-linked loans could also use these markets, in order to hedge their risk.

and reveal information which would facilitate more effective decision making, for instance regarding occupational choice. Second, he believes that the availability of such instruments might encourage occupational choices that may be beneficial for society but (in the absence of insurance) perceived as too risky by an individual (such as highly specialized areas of science which may be hit or miss).

A typical reaction to these proposals, especially by economists, is that "if these markets were needed or useful, they would already exist." Shiller disagrees which such statements, and instead argues that until recently we did not have the technology necessary to collect and maintain the data underlying the various instruments. He points out that there will likely be a need for government help to establish what he calls "Global Risk Information Databases," and that new regulations may need to be enacted in order to make the "New Financial Order" possible (as for instance a change in the bankruptcy law mentioned above).

1.2 Related Literature

The sharing of risk is one of the fundamental topics of economics and finance and therefore much too vast to be surveyed here. We will therefore mainly mention some classic and recent papers that we deem particularly relevant or related to what we are doing. We begin by papers that focus on the welfare cost of income or wage risk (or changes therein), and then move on to papers on households' life-cycle portfolio choice which are more closely related to our work.

Classic empirical papers on the degree of risk sharing between households include Cochrane (1991), Townsend (1994), and Attanasio and Davis (1996). These authors attempt to assess the degree to which household consumption is insured against income changes, and find that insurance is far from perfect (if it were not, there would of course be no need for new financial instruments that facilitate better insurance). Relatedly, papers by Blundell and Preston (1998), Krueger and Perri (2006), Blundell, Pistaferri, and Preston (2008), and Heathcote, Storesletten, and Violante (2008b) use both income and consumption data to look at the evolution of income risk and inequality over the past decades, and use the link between income and consumption inequality to learn something about the degree of household's insurance against income risk.

In quantitative dynamic macro models that are calibrated to match empirical data[6], the welfare cost of wage or income uncertainty is typically very large. For instance, Storesletten, Telmer, and Yaron (2004) use a model that is calibrated to

[6] For an excellent recent summary of this literature, see Heathcote, Storesletten, and Violante (2009)

match the empirically observed evolution of income and consumption inequality over the life cycle to determine the relative importance of initial conditions and life-cycle shocks for inequality, and find that in their model, an agent would be willing to give up 26% of lifetime consumption if that allowed him to insure against all life-cycle shocks. Heathcote, Storesletten, and Violante (2008a) find that the welfare gain from complete markets (i.e., insurance against wage risk) would be almost 40% of expected lifetime consumption.[7] Their conclusion is thus the following (p. 520): "From a policy perspective, an important implication is that the government should develop the legal and institutional frameworks that will allow new insurance markets to develop."

The papers above, and more generally most quantitative general equilibrium macro models, only feature a rather simple asset market structure (often only one asset). We opt to go the partial equilibrium route, which has the advantage of allowing for more realistic asset market structures, but at the cost of taking returns as exogenously given, which is somewhat intellectually unsatisfying and may also yield misleading results in counterfactual exercises.[8] Our life-cycle analysis builds on other computational analyses of optimal portfolio choice over the life cycle, some well-known examples of which include Cocco, Gomes, and Maenhout (2005), Gomes and Michaelides (2005), or Davis, Kubler, and Willen (2006), which is the model we will build on. Some papers in this literature explicitly investigate the welfare effects due to the presence or absence of certain assets or government policies.[9] Perhaps closest in its approach to our paper is a recent paper by Cocco and Gomes (2009), who investigate the role that an asset for which currently no liquid market exists (longevity bonds) could play in individual's portfolios, and what the welfare benefits from such bonds would be. They also discuss the optimal design of such bonds.

[7] In their infinite-horizon model with endogenous labor supply, this gain is more than twice as large as the gain from completely eliminating all risk (for instance, through distortionary taxation), because the latter would take away the opportunity to profit from temporarily high wages by increasing labor supply.

[8] For our exercise of evaluating the potential use and usefulness of new assets, one justification for using partial equilibrium instead of general equilibrium, which would allow for endogenous responses of the other asset returns, is that the world would most likely not move to the new general equilibrium very quickly. Rather, the new assets would need to be introduced into, and used in, the current equilibrium, which is captured by our calibrations.

[9] Examples include Campbell, Cocco, Gomes, and Maenhout (2001) and Gomes, Kotlikoff, and Viceira (2007, 2008).

2 Two-Period Model

To gain some intuition for the effects of the new assets we introduce on asset holdings and welfare, we first consider a simple two-period model. We start by briefly discussing the theory of optimal portfolio choice in the setting we are interested in[10], and then provide an illustration for a calibration that is similar to the one we will be using in the life-cycle model. As will be seen in later sections of the paper, the results from the two-period model largely carry over to the more complex setting.

2.1 Theory

Suppose an investor who lives for two periods has some cash-on-hand in period 1 and expects to receive a stochastic income in period 2 with mean $\mathrm{E}(Y_2)$ and standard deviation σ_{Y_2}. His objective is to maximize his overall expected utility, $u(c_1) + \beta \mathrm{E}[u(c_2)]$. The investor has access to I financial assets, with stochastic or deterministic returns. Finally, assume that the state-space is finite-dimensional. We first consider the optimal policy of an investor who faces no constraints (other than the budget constraint) on his asset holdings between the two periods. In this case, his optimal policy can be understood in terms of a simple algorithm. Start with any admissible asset allocation. This will imply a consumption stream $\{c^1_{t+1}, ..., c^S_{t+1}\}$ for the S states of the world. Then, one can define a new probability measure, called "risk-neutral" or "martingale" probabilities, which re-weighs the objective probabilities of the states by their relative marginal utilities:

$$Q = \begin{pmatrix} \frac{p_1 u'(c_1)}{\sum p_s u'(c_s)} \\ \vdots \\ \frac{p_S u'(c_S)}{\sum p_s u'(c_s)} \end{pmatrix}.$$

Then, for any asset $i \in \{1, ..., I\}$, we can define the "risk-adjusted" (gross) return of the asset using these risk-neutral probabilities: $\mathrm{E}_Q[\tilde{R}_i]$. The intuition is of course that an asset that pays off a lot in states of the world in which consumption is low, and marginal utility therefore high, is relatively more valuable to an investor than an asset that has the same average return but pays off a lot in states of the world in which consumption is high.

If we furthermore define the "shadow" rate as

$$R = \frac{u'(c_t)}{\beta \mathrm{E}(u'(c_{t+1}))},$$

[10]This discussion follows He and Pearson (1991) and Willen and Kubler (2006).

we can characterize optimal portfolio choice by the following simple rule: the investor should optimally add to (subtract from, leave unchanged) his position in an asset i if and only if the risk-adjusted return on that asset exceeds (falls short of, equals) the shadow rate. Thus, overall optimality of the portfolio requires $\mathrm{E}_Q[\tilde{R}_i] = R \; \forall i$. Furthermore, if the set of available assets includes a risk-free asset, so that the investor can borrow and lend potentially unlimited quantities at the risk-free rate, the shadow rate will be equal across all investors.

Classical unconstrained portfolio choice is easy to understand. However, limits on the quantities of assets the investor can hold, such as short-sales or borrowing constraints, make things more complicated. It may now be the case that the investor would like to sell an asset that has a low risk-adjusted return for him, but cannot do so because of short-sales constraints. Alternatively, he may want to buy more of an asset (because the asset has a high risk-adjusted return) but is unable to do so because he has already invested all his wealth in that asset and cannot borrow.

If such constraints are present, which is arguably the most realistic case, optimal asset holdings will depend on an investor's current wealth position and future income process, and shadow rates will differ across investors with different characteristics. Obviously, it also follows that for a given investor, risk-adjusted returns will generally differ across assets.

For instance, assume that the only available assets are risk-free borrowing and lending, at rates R_b and R_l respectively, with $R_b > R_l$ and the constraints that $b \leq 0$ and $l \geq 0$. Then, a relatively poor investor will borrow today, which means that his shadow rate R equals R_b and exceeds R_l – if he could, he would like to set $l < 0$, but he cannot do so. Likewise, a relatively rich investor lends today and has $R = R_l < R_b$, as he cannot set $b > 0$.

Now, suppose we add to this setting the possibility of investing in an income-hedging instrument (IHI) with $\mathrm{E}[\tilde{R}_{IHI}] = R_l$ and $corr(\tilde{R}_{IHI}, Y_2) < 0$. In an incomplete markets world, we would then typically have $\mathrm{E}_Q[\tilde{R}_{IHI}] > R_l$, because consumption tracks income. As a consequence, if relatively poor investors could borrow at R_l, it would always be worth it for them to do so in order to buy the IHI. Likewise, relatively rich investors who would otherwise save at R_l would now instead invest in the IHI. However, as discussed above, in the real world it is well possible that people are borrowing at a higher rate, such that $R = R_b$, or they may even be maxed out on their borrowing, such that $R > R_b$. In such cases, it is far from clear that $\mathrm{E}_Q[\tilde{R}_{IHI}] > R$, so that the investor may not want to hold the IHI. Similarly, relatively rich investors may have access to other investment opportunities, such as equity, which offer higher risk-adjusted returns than the IHI.

To summarize, in order to determine whether investors will demand an income-linked asset (or any other asset), we need to know the risk-adjusted return on this asset and compare it to their shadow rate. This shadow rate will in turn depend on their current wealth position, their expected future income and its riskiness, and the return processes of the other assets they have access to.

2.2 Example

An investor expects to receive a stochastic income in period 2 with mean 8 dollars but subject to a stochastic shock with standard deviation 1.5.[11] The investor has an isoelastic utility function with relative risk aversion 2, and does not discount the future.

As a benchmark, suppose that he can borrow at $r_b = 8$ percent, save at $r_l = 2$ percent, or invest in equity with expected return $E(\tilde{r}_e) = 6$ percent and standard deviation of 16 percent.[12]

The first panel of Figure 1 displays the investor's optimal asset holdings as a function of his cash-on-hand in period 1. As his goal is to smooth consumption over the two periods, he borrows when he is relatively poor in period 1, and saves (by investing in equity) if he is relatively rich. In terms of the terminology of the previous subsection, the shadow rate equals 8 percent in the cash-on-hand region where the investor borrows, then falls to 6 percent (the mean return on equity) at the point where the investor starts investing in equity, and then further decreases in cash-on-hand. As $r_b > E(\tilde{r}_e)$, the investor does not borrow to invest in equity; also, he does not engage in risk-free saving (he would do so though if he were more risk averse or had very high cash-on-hand).

The second panel shows the optimal asset holdings if in addition to the assets from the benchmark model, the investor has access to an income-hedging instrument (IHI) with $E(\tilde{r}_{IHI}) = r_l = 2$ percent, standard deviation 0.25, and a negative return correlation of 0.5 with the income shock in the second period. Thus, the IHI tends to pay off more when the investor gets hit by a negative income shock and less if his income exceeds expectations. The optimal policy features positive holdings of

[11] Y_2 can take the values $\{5.4, 8, 10.6\}$ with respective probabilities $\{1/6, 2/3, 1/6\}$.

[12] There is no exogenously imposed borrowing limit, but we require the investor to pay back his debt in period 2. Given our assumed possible values for income in period 2, this means the investor can borrow at most $5.4/(1+r_b)$ (or more if he hedges his income risk). It may seem odd that $r_b > r_l$ even though there is no default in the model. However, there are many reasons why borrowing costs exceed lending rates, such as transaction costs or the cost that lenders face in the screening of potential borrowers (which should precisely result in lower risk of default).

the IHI at low levels of cash-on-hand, financed by additional borrowing. As cash-on-hand increases, holdings of the IHI decrease, and equal zero for cash-on-hand levels between 5.5 and 7.3. At higher levels of cash-on-hand, they become positive again, and eventually the investor simultaneously holds the IHI and equity. IHI holdings continue to increase in cash-on-hand even for higher levels of cash-on-hand than depicted in the graph, up to a point at which the variance of cash-on-hand in the next period cannot be decreased any further by higher IHI holdings. As a consequence, equity holdings are lower than in the benchmark.

In the third panel, we instead add the possibility of borrowing through an income-linked loan (ILL). We assume that this loan features a stochastic interest rate with mean $E(\tilde{r}_{ILL}) = r_b = 8$ percent, standard deviation 0.25, and a positive correlation of 0.5 with the income shock in the second period. Thus, when taking out an ILL, the investor expects to have to repay a larger amount if his income is higher than expected next period and a lower amount if income falls short of expectations. The figure shows that the investor makes quite extensive use of the ILL. For low levels of current cash-on-hand, borrowing through the ILL mostly replaces fixed-rate borrowing, without leading to much additional total borrowing.[13] Another feature of the optimal investment policy is to keep borrowing a decreasing amount as cash-on-hand increases. When borrowing through the ILL, the investor takes a larger position in equity as compared with the benchmark case.

Interestingly, demands for both the IHI in panel 2 and the ILL in panel 3 are non-monotonic in wealth. This may be surprising at first, but is a general feature of portfolio choice problems with short-selling constraints.

The final panel displays the welfare gains (in percent of certainty-equivalent consumption) from having access to one of the two assets in this example. As can be seen, both assets provide higher gains for relatively poor investors than for rich ones. Also, over most of the cash-on-hand range depicted, the income-linked loan provides higher welfare gains over the benchmark case than the income-hedging instrument. Here is some intuition for why this is the case. Consider first a case in which the investor has little cash-on-hand in the first period. When he has access to the income-linked loan, he uses it instead of risk-free borrowing. Given that it has the same expected cost (as we assume $E(\tilde{r}_{ILL}) = r_b$) but provides additional insurance benefits as compared with standard borrowing, the ILL clearly is a good deal for the investor. As cash-on-hand increases, he keeps borrowing through the ILL, but now also invests in equity, which has a slightly lower expected return than the expected interest rate on the ILL. As

[13] The ILL does not fully replace fixed-rate borrowing because that would be too risky, given the imperfect correlation of the interest rate with next period's income.

such, the insurance provided by the ILL becomes somewhat more "expensive." For high cash-on-hand levels, ILL borrowing goes to zero, so that investors in that range do not gain from having access to the ILL. The IHI holdings, on the other hand, are financed through expensive borrowing at low cash-on-hand levels, and reduce equity holdings at higher levels. Thus, intuitively, the (opportunity) cost of holding the IHI is higher than for ILL. However, note that IHI holdings do not go to zero as cash-on-hand increases – even for rich investors, hedging next period's income risk has some value (though the welfare gains in this example are minuscule).[14] Thus, for such investors, the IHI is preferred to the ILL.

The previous discussion hints at the role of equity in this model: it makes ILL relatively more attractive (by lowering its effective cost) and the IHI relatively less attractive (because the opportunity cost of investing in the IHI is higher than when only risk-free saving were available). Indeed, it can be shown that for relatively rich investors (with cash-on-hand above 7.3 in this example), IHI holdings would be larger and ILL holdings smaller when no equity is available than in the case depicted in Figure 1. As a consequence, for such investors the IHI would lead to larger welfare gains and the ILL to lower welfare gains than in the figure.[15]

One can also compare the welfare gains from the two assets with the welfare gain that would result from completely eliminating income risk (that is, the investor receives an income of 8 for sure in the second period). With the parameters we assumed, this gain would be much larger than the ones depicted; for an investor with no cash-on-hand, the gain in certainty-equivalent consumption would equal 9.2 percent, while an investor with cash-on-hand of 5 would gain 2.8 percent and one with cash-on-hand of 10 about 1.4 percent. Thus, the assets we introduce reap rather little of the potential gains. What accounts for this? Part of the explanation is provided by the imperfect correlation of the rates with the income shock. As Table 1 shows, if rates are perfectly correlated with the income shock, the gains provided by the two assets move significantly towards the gains that an elimination of income shocks would provide. This is particularly true for relatively poor investors; rich investors still gain less. Also, the table confirms that for relatively poor investors, ILL are more useful while for rich ones, the IHI leads to higher welfare gains.

[14] This is because the risk-adjusted return on equity declines in equity holdings and thus tends towards r_l as cash-on-hand increases. The risk-adjusted return on IHI decreases as well in holdings, and always remains weakly above r_l, so that optimal IHI holdings do not decline as cash-on-hand goes up.

[15] The gains in CE consumption from the IHI would average about 0.25 percent for an investor with cash-on-hand between 10 and 15, for instance, while the welfare gains from ILL would be zero for such an investor, as he would not borrow through the ILL at all.

Lines 5–8 of the table furthermore show the effect of more volatile returns of the IHI or ILL on welfare gains. Clearly, given our earlier discussion of (opportunity) costs, it is no surprise that these assets become more useful if their returns are more volatile, as one then has to hold less of them to obtain the same insurance. However, the table also demonstrates two additional points. First, the welfare gains from the ILL seem less affected by changes in the volatility than the ones from the IHI. Second, higher volatility has a relatively larger positive effect on welfare the more strongly returns and income are correlated. This is intuitive: an asset that is highly volatile but only offers an imperfect hedge against income risk also adds risk. Thus, even though an increase in volatility never lowers the welfare gains provided by an asset, it may be that it does not increase welfare gains or only slightly so.

In the remainder of the paper, we will see that the main points seen in this section will carry over to a more realistic life-cycle setting. The demand for income-linked assets, and the welfare gains achieved by their presence, will be very sensitive to the parameters of the return process. Also, the welfare gains will be found to be rather small for the parametrizations we deem most realistic (in particular as compared with the hypothetical gain from completely eliminating all income risk), and income-linked loans appear to be much more promising than the income-hedging instrument.

3 Life-Cycle Model

3.1 Setup

Our strategy in this part is the following: We start out with a life-cycle portfolio choice model with realistic borrowing and investment opportunities. We show that this model generates predictions regarding borrowing and equity holdings that are roughly consistent with the data. We then introduce a new asset into this model, and analyze what the demand for this asset would be, how it would affect the demand for the other assets in the model, and what the predicted welfare gains from the new asset would be.

The specific portfolio choice model that we build on is the one by Davis et al. (2006). This model explicitly takes into account the fact that the typical household has access to unsecured credit, albeit at a higher interest rate than the lending rate or the expected rate of return on equity. Young households, who expect higher income in the future, typically take advantage of this borrowing opportunity in order to smooth their consumption over time. Yet, this slows down the speed at which they accumulate wealth, and reduces the rate of participation in equity markets until about age 45.

As a consequence, this model generates more realistic predictions than models that allow no borrowing or, at the other extreme, borrowing at the risk-free lending rate. Furthermore, this is accomplished without a need to rely on implausible preference parameters.

The basic ingredients of our life-cycle consumption and portfolio choice model are the standard ones used in this literature. The household life cycle consists of two phases, work and retirement. Retirement age is assumed to be exogenous, at t_R. During working years, log labor income ($\tilde{y}_t$) evolves as the sum of a deterministic component (d_t), a random walk component ($\tilde{\eta}_t$), and an i.i.d. transitory shock ($\tilde{\varepsilon}_t$):

$$\tilde{y}_t = d_t + \tilde{\eta}_t + \tilde{\varepsilon}_t \text{ for } t \leq t_R \tag{1}$$

where $\tilde{\eta}_t = \eta_{t-1} + \tilde{\nu}_t$, with $\tilde{\nu}_t \sim N(-\sigma_\nu^2/2, \sigma_\nu^2)$, and $\tilde{\varepsilon}_t \sim N(-\sigma_\varepsilon^2/2, \sigma_\varepsilon^2)$. Thus, $\Delta\tilde{y}_t$ is an MA(1) process.

During retirement, it is assumed that the household receives a constant fraction λ of its permanent income in the last year of work:

$$\tilde{y}_t = \log(\lambda) + d_{t_R} + \eta_{t_R} \text{ for } t > t_R. \tag{2}$$

The household is assumed to have isoelastic preferences, and discounts the future at a constant rate. Thus, in each period t, it maximizes

$$U(\alpha_t c_t) + E_t \sum_{s=t+1}^{T} \beta^{s-t} U(\alpha_s c_s) \tag{3}$$

$U(\cdot)$ is an isoelastic (power) utility function with curvature γ. α is a "taste shifter" which we include mainly to account for the consumption drop at entry into retirement.[16] We assume that the household dies with certainty at age T, and do not include stochastic death or a bequest motive in our model.

The budget constraint of a household at age t is the following, in its most general form:

$$\underbrace{c_t}_{\text{Consumption}} + \underbrace{e_t}_{\text{Equity}} + \underbrace{l_t}_{\text{Saving}} + \underbrace{IHI_t}_{\substack{\text{Income-}\\\text{hedging}\\\text{instrument}}} - \underbrace{b_t}_{\substack{\text{Fixed-rate}\\\text{borrowing}}} - \underbrace{ILL_t}_{\substack{\text{Income-}\\\text{linked}\\\text{loan}}} =$$

$$\underbrace{Y_t}_{\substack{\text{Labor}\\\text{income}}} + \tilde{R}_{e,t} e_{t-1} + R_l l_{t-1} + \tilde{R}_{IHI,t} IHI_{t-1} - R_b b_{t-1} - \tilde{R}_{ILL,t} ILL_{t-1}.$$

[16] It can be seen as a stand-in for a more elaborate model with labor supply. For instance, Cocco and Gomes (2009) use $\alpha_t = L_t^\gamma$, which generates a consumption drop at retirement, because leisure and consumption are substitutes in that utility function.

Households can always trade at least three financial assets. They can buy equity (e) with stochastic return $\tilde{r}_e$ ($= \tilde{R}_{e,t} - 1$), save (l) at a net risk-free rate r_l, and borrow (b) at a fixed risk-free rate r_b. We will refer to the version of the model in which only these three assets are available as the *benchmark case*.

We then add an additional financial asset to this model. The first possible addition is an *income-hedging instrument* which has stochastic return $\tilde{r}_{IHI}$ that is negatively correlated with the permanent income shock the household receives. We vary this correlation, as well as the volatility of the interest rate, to see how these parameters affect demand for and welfare gains from the asset. The other addition is *income-linked loans*, which offer another way for the household to borrow. They are different from risk-free borrowing in that their interest rate $\tilde{r}_{ILL}$ is stochastic and positively correlated with the permanent income shock the household receives.

We do not impose an exogenous borrowing constraint, but require that households are able to repay their debt with probability 1 by the time they die, so that $b_T = ILL_T = 0$ (this is usually referred to as the "natural debt limit"). Thus, households never default on their debt in our model. Another simplifying assumption of the model is that it ignores housing and secured (mortgage) borrowing. Given that a large proportion of households hold much of their wealth in housing, this will lead the model to overpredict equity holdings. However, apart from that, we do not believe that omitting housing and mortgages from the analysis has a large influence on our results.[17]

3.2 Welfare and Insurance Measures

We employ different measures to determine how "useful" the assets we introduce are for a household. First, we will analyze what the model predicts for the demand for these assets (how much households would hold on average at different stages of the life cycle). However, this does not tell us much about the welfare benefits from the new assets.

A better measure (which is standard in the literature) is the gain in certainty-equivalent (CE) consumption due to the introduction of a new asset. CE consumption is computed as follows: We first find the (ex-ante) lifetime expected utility $\bar{U}$ in a given environment. Then, we find the constant level of consumption, $\bar{c}$, that would

[17] Mortgages usually come at a lower interest rate than the unsecured borrowing we focus on in our model. However, except perhaps over the past few years, mortgages are usually taken out with the sole purpose of buying a house, not to smooth consumption or invest the borrowed money in equity or other assets. Nevertheless, in section 3.4.4 we consider a calibration in which households have access to low-cost borrowing, for instance through home-equity loans.

yield the same level of utility:[18]

$$\left(\sum_{t=0}^{T-1}\beta^t\right)\frac{\bar{c}^{1-\gamma}}{1-\gamma}=\bar{U}\Leftrightarrow\bar{c}=\left(\frac{1-\beta}{1-\beta^T}(1-\gamma)\bar{U}\right)^{\frac{1}{1-\gamma}} \tag{4}$$

Finally, we will also use the measure of partial insurance against permanent shocks proposed by Kaplan and Violante (2008). Define the insurance coefficient at age t as

$$\phi_t^\nu = 1-\frac{cov(\Delta c_{it},\nu_{it})}{var(\nu_{it})} \tag{5}$$

where c_{it} is log consumption, ν_{it} the innovation to the permanent component of log income, and variances and covariances are taken over the cross-section of simulated households at age t. The interpretation of this coefficient is intuitive: the lower it is, the more a permanent income shock translates into consumption changes. If $\phi_t^\nu = 0$, consumption adjusts 1-for-1 with permanent income. On the other hand, $\phi_t^\nu = 1$ would mean “perfect insurance” in the sense that households’ consumption growth is completely independent of the particular shock they receive.

3.3 Calibration and Discretization

For the labor income process, we use the parameters from Cocco et al. (2005) for high school graduates, which has become somewhat of a standard in this literature. d_t is given by a third-order polynomial in age, the standard deviations of the permanent and transitory shock are set to 0.103 and 0.272, respectively, and the replacement rate λ equals 0.682. Households enter the model at age 20, retire after age 65, and die with certainty at age 80.

The main preference parameters are set as follows: relative risk aversion γ is assumed to equal 2. The discount factor β is chosen such that the mean wealth-to-income ratio of households with a head aged 50 to 59 in the “base case” of the model where only equity, unsecured borrowing, and risk-free saving are available, matches its empirical counterpart of 2.6 (Laibson, Repetto, and Tobacman 2007).[19] This yields $\beta = 0.936$. The taste shifter α equals one before retirement and 0.9 afterwards; this generates a mean consumption drop at retirement of about ten percent, which is consistent with most empirical estimates.

[18] All results we report are based on simulation of the life cycle for 5,000 households, using the same random draws for all parameterizations.

[19] The empirical wealth measure used to obtain this number includes claims on defined contribution pension plans, but not Social Security wealth or claims on defined benefit plans, which are included in our retirement income measure.

For asset returns, we make the same assumptions as in the two-period example earlier. We set the annual return on risk-free saving, r_l equal to 2 percent per year, and the mean equity premium, $E(\tilde{r}_e) - r_l$ to 4 percent, which are customary values in this literature. The standard deviation of equity returns, σ_e, is set to 16 percent. The interest rate on risk-free borrowing, r_b, is set equal to 8 percent, which is what Davis et al. (2006) choose based on empirical data in which they find an interest rate differential between the risk-free lending rate and the mean rate on unsecured borrowing of approximately 6 percent, after adjusting for tax considerations and charge-offs.

We solve the model using numerical methods. The algorithm is similar to the one used by Davis et al.. Depending on the asset market assumptions, there are three or four sources of randomness in our model: the permanent income shock, the temporary income shock, the equity return shock, and the income-linked asset rate shock. We discretize the state space using Gaussian quadrature, with two nodes for the labor income shock, three for the equity return, and four for the income-linked asset return. This is not restrictive: using five nodes for each shock does not qualitatively alter the results (but significantly increases computation time).[20]

3.3.1 On the Labor Income Process

Clearly, if one wants to make a quantitatively appropriate assessment of the welfare burden of labor income risk, and the welfare gain of having access to financial instruments that can be used to hedge part of that risk, it is important to use a realistic labor income process with appropriate degrees of uncertainty. We follow the bulk of the existing literature and use a slightly simplified version of the labor income process introduced by MaCurdy (1982) and Abowd and Card (1989). This "permanent-transitory" process was popularized in the consumption literature by Zeldes (1989), Carroll (1997) and Gourinchas and Parker (2002), and has the advantage that the life-cycle optimization problem can be normalized by permanent income, which reduces the number of state variables and makes the computational solution of the model easier.

The main feature of the permanent-transitory income process is that there is no

[20]Results would change, however, if we increased the number of possible income shock realizations a lot, so that we would have the possibility of a shock in the very far left tail of the lognormal shock distribution. This would affect the natural debt limit, and in the extreme case of a possible zero-income shock, eliminate borrowing all together. We believe that it is realistic to assume that there is a positive lower bound for income shocks, due to social safety nets, and of course the fact that we do observe very significant borrowing in the real world is consistent with this.

individual heterogeneity in income growth rates beyond what is captured in the deterministic component d_t, which is typically estimated separately for different education levels. However, Guvenen (2007, 2009) has recently argued that this assumption may be overly restrictive, and that allowing for 'heterogeneous income profiles' (HIP) can account for features of the evolution of consumption inequality and the slopes of consumption profiles for different education groups over the life cycle that would otherwise be puzzling. In his model, income shocks are less persistent, but individuals only learn about their personal d_t over time, through observation of their realized income. While Guvenen's model intuitively makes a lot of sense, it is very computationally demanding; to our knowledge, nobody has solved a version with more than the risk-free asset. Furthermore, Hryshko (2008) argues that the PSID income data actually reject the HIP model when it is estimated in first differences, while the model with a permanent component that we use cannot be rejected. Thus, the question of which process is preferable is far from settled, and to maintain comparability with the existing portfolio choice literature as well as computational simplicity, we stick with the status quo.

Even once one has decided on the form of the income process, there remain calibration decisions of crucial importance for the extent of uncertainty and the consumption and asset profiles over the life cycle. In particular, the assumed variances for the permanent and transitory shocks matter a lot. As mentioned above, we use the estimates for high school graduates by Cocco et al. (2005). They estimate a variance of the permanent shock of $\sigma_\nu^2 = 0.0106$ and a variance of the transitory shock of $\sigma_\varepsilon^2 = 0.0738$. These estimates are quite different for instance from the ones by Gourinchas and Parker (2002), who find $\sigma_\nu^2 = 0.0277$ and $\sigma_\varepsilon^2 = 0.0431$, meaning that the relative volatility of permanent shocks is significantly larger. Feigenbaum and Li (2009) point out that the estimates strongly depend on sample length of the PSID data used.[21] They find, using the longest possible sample 1968–2001, $\sigma_\nu^2 = 0.009$ and $\sigma_\varepsilon^2 = 0.071$, which is close to the numbers we are using.[22] Furthermore, they compare what these numbers imply for income uncertainty over various future horizons to the results of a semi-parametric model. While the implied uncertainty of the permanent-transitory process with the variances they estimated is somewhat too high, its slope over different horizons seems much more appropriate than if the Gourinchas and Parker numbers were used.

[21] They point out that this fact by itself may be an indication that the persistent shock follows an AR process, and not a random walk.

[22] Gourinchas and Parker use the data of Carroll and Samwick (1997), which only comprises the PSID years 1981 to 1987. Cocco et al. use 1970 to 1992.

3.3.2 Correlation between Income-Linked Assets and Income Shocks

In our assessment of the use and usefulness of the income-linked assets, the assumed correlation between the return on the income-linked asset and an individual's labor income shock plays a crucial role. In this subsection, we briefly discuss the correlation that could realistically be achieved if the return on the income-linked asset were solely based on an occupation-level income index. The empirical question is: how much of an individual's income risk is specific to his occupation, and how much is completely idiosyncratic? In terms of our model, we can decompose individual i's permanent shock ν_{it} into a group-specific component $\xi_t \sim N(-\sigma_\xi^2/2, \sigma_\xi^2)$, and an independent idiosyncratic component $\omega_{it} \sim N(-\sigma_\omega^2/2, \sigma_\omega^2)$:[23]

$$\nu_{it} = \xi_t + \omega_{it}. \tag{6}$$

Then, if we assume that the return on the income-linked asset (ILA) is perfectly correlated with the group-specific permanent shock ξ_t, we have that the correlation of the return with individual i's permanent income shock is given by

$$corr(\tilde{r}_{ILA,t}, \nu_{it}) = \frac{\sigma_\xi}{\sigma_\nu}. \tag{7}$$

Thus, in addition to the total standard deviation of a person's permanent income innovations, we need to know the standard deviation of permanent shocks to an occupation's income series. In Davis, Fuster, and Willen (2009), we use repeated cross sections of the Current Population Survey to construct occupation-level components of individual income shocks (after removing predictable components of individual income due to age, education, etc.) for 17 occupations for which the occupational classification has remained largely unchanged over 35 years or more and for which we have a relatively large number of individuals in each survey year.[24] Although these occupations are not necessarily a representative set of occupations for the US population, we can at least get an estimate of the order of magnitude of occupation-specific income shocks. If we assume for simplicity that all occupation-level income shocks are permanent (an assumption that is not too far from the truth in our data – see Davis et al. for details), we can get an estimate of σ_ξ simply by looking at the standard deviation of annual changes to the occupation-level income index. For

[23]This decomposition is similar to the one used in Cocco et al. (2005), who concentrate on an aggregate component instead of a group-specific component.

[24]These occupations are: Accountants and Auditors, Electrical Engineers, Registered Nurses, Elementary School Teachers, Cashiers, Secretaries, Police and Detectives, Waiters and Waitresses, Cooks, Janitors and Cleaners, Auto Mechanics, Carpenters, Electricians, Plumbers, Machinists, Welders and Cutters, and Truck Drivers.

the ten occupations in our data for which at least two-thirds of the individuals are high school graduates but not college graduates, these standard deviations range from 0.021 (secretaries) to 0.059 (plumbers), with an average of 0.038. Given our σ_ν of .103, this average would imply a correlation of individual permanent income shocks with the return on an asset that is based on an occupation-income index of slightly below 0.4.[25] This estimate comes from the best data currently available; yet it is possible that, if better and broader data sources became available in the future, "finer" indices could be constructed which would be more highly correlated with individual income shocks (e.g., "Plumbers in New England"). We take a somewhat optimistic baseline assumption, namely a correlation of 0.5. One of the main results from our analysis, however, will be how sensitive the usefulness of income-linked assets is to this correlation.

3.4 Results

3.4.1 Benchmark Case

In the benchmark case, households can borrow at a rate r_b, and invest either in the risk-free asset with a fixed return r_l or in equity with a stochastic return $\tilde{r}_e$. The life-cycle profiles, displayed in Figure 3, mirror the ones in Davis et al. (2006): households borrow substantial amounts while they are young (on average, 50% of their annual income between ages 20 and 30)[26] and only start making substantial investments in the stock market after age 35. The predicted equity market participation rate starts out around 20% for young households and increases through mid age, reaching 95% at age 45. Average equity holdings when reaching retirement amount to about three

[25]Ideally, one would also want to estimate σ_ν separately by occupation, rather than simply taking the estimate for all high school graduates. Unfortunately, the PSID does not contain a large enough number of observations to do that. However, Campbell, Cocco, Gomes, and Maenhout (2001) split households into 36 different industry-education cells, and estimate separate labor income profiles for each cell. The four industries that are most relevant for our occupations (which coincidentally have the largest cell sizes in the PSID), are the following (with Campbell et al. (2001)'s estimate of the permanent shock standard deviation for high school graduates in the industry in brackets): Manufacturing (.068), Construction (.120), Trade (.106) and Transportation (.067). The differences in the magnitude of the permanent shocks is in accordance with what we find on the occupation level: occupations that are mainly active in construction or trade tend to have a higher standard deviation of earnings changes than the ones in manufacturing and transportation.

[26]This is somewhat higher than the average unsecured borrowing (credit card balances plus installment loans and other unsecured borrowing) as a percentage of income reported in the Survey of Consumer Finances (SCF), which was 28% for below-30-year olds in 1995 and 1998 (Davis et al. 2006). However, Zinman (2009) finds that the SCF misses around one-half of revolving debt.

times annual income; this is about twice as high as in the data. However, this is arguably not a major failure of the model, as it does not include home equity, which is a risky asset most households are invested in in reality. Other features of the model are that it predicts practically no borrowing for households after age 40, and no significant bond holdings at any age. Both of these predictions are somewhat at odds with reality; this may be due to liquidity motives that are missing from the model.[27] (In the robustness section, we will consider a version of our model in which households are forced to invest at least 50% of their financial wealth in bonds.) Another possible shortcoming of the model is that it produces the consumption hump that is typically observed in empirical data only for median consumption, while mean consumption increases until retirement.[28]

For our benchmark case, CE consumption equals 19,638 USD.[29] The partial insurance coefficient averages 0.09 over the life cycle. This is significantly lower than what Kaplan and Violante (2008) find in their model, which features more redistributive social security, and even further below the insurance coefficients that Blundell et al. (2008) estimate in empirical data. This means that we may be overestimating the welfare cost of income uncertainty and therefore also the potential gain that new assets would be able to provide.

The welfare cost to the typical household of income shocks is fairly high in this model: if there were no life-cycle income shocks, but all asset return characteristics remained the same, CE consumption would equal 22,861 USD, or 16.4% more.[30] Thus, the ex-ante cost of income shocks is high, and of similar order of magnitude as what is found in the quantitative macro literature discussed earlier. Thus, one would hope that introducing financial assets to hedge the risk of these income shocks could yield high welfare gains. Meanwhile, the welfare gain of having access to stocks is fairly modest in this model: without equity, CE consumption equals 19,424 USD, and thus only 1.1% less than in the benchmark with equity. It is useful to keep this in mind when evaluating the potential gains from other assets.

We will now introduce two different assets that the households can use to hedge their labor income risk: first, an income-hedging instrument whose return corre-

[27] The continued credit card borrowing might also be due to consumers having self-control problems, as in Laibson et al. (2007).

[28] Mean consumption is so much higher than median consumption in our model because some of our households get very rich thanks to positive income and asset return shocks. It may be that the introduction of flexible labor supply would reduce this disparity.

[29] All dollar amounts are expressed in 1992 dollars.

[30] In this counterfactual, we do not alter the income shocks in the first period of the working life, which can be seen as a "fixed effect," for instance due to inherent differences in ability, and could never be insured against in our model.

lates negatively with the permanent income shock a household receives, and second, income-linked loans (where the interest rate paid on a loan is positively correlated with the household's income shock). We throughout only consider assets that correlate with the household's *permanent* income shock. This is because in models such as ours, the transitory shock is usually smoothed out easily by the household and has very little effect on welfare or asset allocations.

3.4.2 Income-Hedging Instrument

Our baseline assumption for the mean return on the income-hedging instrument available to a household is that it is equal to the rate on risk-free bonds, or "actuarially fair": $E(\tilde{r}_{IHI}) = r_l$. We vary the correlation of the return with the household's permanent income shock from –0.25 to –1, taking –0.5 as our baseline, based on the empirical evidence discussed in Section 3.3.2. The baseline for the standard deviation of returns is 0.5, meaning that the return is much more volatile than the return on equity. This may seem excessive, but given that the IHI is more useful to the household if its return is more volatile (at least up to some point), and as in principle this asset could be created arbitrarily volatile, we chose this high volatility as our benchmark.

Figure 4 shows the mean life-cycle holdings of IHI, equity, and unsecured borrowing (which we denote by "CC", standing for credit card) for the baseline case and also for two cases with higher (absolute) correlation.[31]

Our first finding is that the baseline IHI is not highly demanded by households. As the top left panel of the figure shows, mean holdings of the IHI never go much above 5,000 USD, and as the second and third panels show, this investment is financed almost exclusively by a reduction in equity holdings, not by additional borrowing. As a consequence, young households hold only very little of the IHI (less than half the households hold any until age 33). Mean IHI holdings peak around age 50, then slowly decline towards retirement, while equity holdings keep increasing until retirement.

The bottom right panel of the figure shows how the presence of the IHI affects the degree of partial insurance. What can be seen is that particularly for young households, which are not insured against shocks to permanent income, the degree of insurance is virtually unchanged.[32] Even for older households, the increase in the degree of insurance is rather small.

[31] We do not display the case with lower correlation ($\rho = -0.25$) because in this case there are no IHI holdings.

[32] It may be surprising that insurance is slightly negative for young households in the benchmark. Kaplan and Violante (2008) explain that this is due to the interaction of permanent and transitory shocks in this model (see their fn. 27).

Next, the figure shows how strongly the demand for IHI, as well as the impact on borrowing and equity holdings, depend on the correlation between the IHI return and the permanent income shock. With $\rho = -0.75$, demand for IHI is higher, but still starts out relatively low. It is financed by a combination of reduced equity holdings and additional borrowing. On the other hand, with $\rho = -1$, households start holding high amounts of IHI much earlier in the life cycle, and borrow massively higher amounts – much more than what is needed to finance their IHI holdings. This is because there is now less need for precautionary wealth, and households can consume more in anticipation of their higher future income. As the bottom right panel shows, in this case insurance against the permanent shock is much improved, even though it is still only around 0.5 on average over the working life.

Volatility matters greatly for mean holdings of the IHI as well as the effect on equity holdings and borrowing.[33] For instance, for the baseline correlation of –0.5, there are practically no IHI holdings when the standard deviation of IHI returns is only 0.3 (nobody participates in the IHI market until age 50; the maximum participation rate is 5 percent, right before retirement), and mean IHI holdings also decrease for the other assumed correlations. Also, while with perfect correlation and volatility 0.5, households' average borrowing over the life cycle is above 25,000 USD, with volatility 0.3 the corresponding number is below 4,000. The partial insurance coefficients are also very significantly reduced.

Figure 5 displays the welfare gains over the benchmark (in percent of certainty-equivalent consumption) that having access to the IHI would generate for consumers. The first thing one notes is how strongly the gains depend on the correlation and the volatility. Welfare gains are convex in the strength of the correlation.[34] The welfare gains are tiny (below 0.1 percent) if the correlation between the IHI return and the permanent shock is 0.5 or less, while if the correlation is perfect and volatility high, the gain reaches almost 2.4 percent (with low volatility, on the other hand, the corresponding gain is only 0.3 percent).

Overall, the results in this section indicate that unless the IHI had volatile returns that are highly correlated with a household's permanent income shock, the welfare gains it generates are very small. This may be surprising, given that our assumptions are such that we likely obtain an upper bound of the welfare gains. As we will further discuss in section 4, the reason behind the small welfare gains lies in household's effective cost of funds, mainly their borrowing costs, and the opportunity costs due to the possibility to invest in equity.

[33] This can be seen by comparing lines 5–7 of Table 3 to lines 2–4.

[34] They are also convex in the *square* of the correlation.

3.4.3 Income-Linked Loans

For the income-linked loan, our baseline assumption is that the mean interest rate a borrower needs to pay on it is the same as for other unsecured ("credit card") borrowing: $E(\tilde{r}_{ILL}) = r_b$. As in the previous section, we again make different assumptions about the volatility of the rate and its correlation with the household's permanent shock. Our baseline assumption is to set both parameters equal to 0.5.

Figure 6 shows mean borrowing through the ILL, equity holdings, and other borrowing under this baseline assumption as well as for higher correlations. The first panel shows that mean ILL borrowing increases in age for young households, then peaks between ages 30 and 35, and decreases towards retirement. The higher the correlation between the rate on the ILL and the permanent income shock, the more extensively households borrow through the ILL.

It is interesting to consider the effect of the presence of income-linked loans on other borrowing. The bottom left panel shows that when $\rho = 0.5$, ILL borrowing reduces other borrowing early in life, as one would expect, given that it clearly has a lower risk-adjusted cost. What may be more surprising at first is that as ρ increases, households massively increase their fixed-rate borrowing, even though they also engage in a lot of borrowing through the ILL. This is again because the uncertainty about future resources is now much smaller, so that young households can borrow to consume their future (higher) income. What they do not consume is invested in equity, so that mean equity holdings are higher for young households when ILL are available, while being lower for households approaching retirement, which had less of a need to accumulate precautionary wealth.

The bottom right panel displays the insurance coefficients, and we see that especially for young households, the ILL is much better at improving insurance than the IHI in the previous subsection. With $\rho = 1$, households approach perfect insurance, especially in the early parts of the life cycle.

Lines 12–15 in Table 3 show what happens to mean asset holdings (and welfare gains) if the ILL rate is less volatile (0.3 instead of 0.5). It is interesting to note that the effect of lowering volatility on mean ILL borrowing does not go in the same direction for all assumed correlations. For some ρ, mean ILL borrowing increases when the ILL return volatility decreases (which makes sense, because one needs to borrow more through the ILL to get the same degree of insurance) while other times mean ILL borrowing decreases (which also makes sense because the ILL is less attractive when its volatility is lower). The effects on equity holdings are small, while other borrowing is reduced.

The welfare gains from different types of ILL are shown in Figure 7. As was the case for the IHI, they are again convex in the strength of the correlation of rates and permanent income shocks. However, they are now much higher, and also less dependent on highly volatile rates. For the baseline case of a correlation of 0.5, the welfare gain from ILL is 1.36 percent if volatility is 0.5 and 0.95 percent if volatility is 0.3. This would be quite a substantial welfare gain.

3.4.4 Robustness Checks

In this subsection, we investigate further what is driving the differences in welfare gains between the two assets we consider, and how sensitive our results are to the assumed risk aversion of agents. A summary of the results is provided in Table 4.

Equity returns and borrowing rates. One weakness of our benchmark model is that it does not match empirically observed bond holdings – indeed, it predicts practically no bond holdings at all, because equity is so much more attractive. This is a common feature of portfolio choice models such as the one we use.[35] One possible explanation for this is that our model ignores potential liquidity benefits to holding the risk-free asset, or that there exist participation costs in equity markets that we have not modeled. In our model, the presence of equity makes the IHI relatively less attractive (because it has to "compete" with equity to enter households' portfolio) while the ILL is made relatively more attractive (because what a household borrows through ILL in order to insure against income fluctuations can be invested in a high return asset). As a consequence, our model may be understating the gains from IHI and overstating the gains from ILL.

To address this issue, we solve a version of the model in which households are required to invest at least as much money in bonds as they invest in stocks, which will make the portfolios generated by our model look closer to what is observed empirically. For our calibration, this is equivalent to replacing equity by a 50-50 stocks-bonds fund with expected return of $0.5 \cdot (E(\tilde{r}_e) + r_l)$ and standard deviation $0.5 \cdot \sigma_e$. We again choose β in order to match the mean wealth-to-income ratio before retirement, which yields $\beta = 0.947$. As compared with the benchmark with no income-linked assets, the gain from having access to the baseline IHI with correlation –0.5 and volatility 0.5 is now 0.33% of certainty-equivalent consumption and thus, as expected, higher than if no bond holdings are required. In contrast, the baseline ILL now produces a welfare gain of only 0.85%, instead of the 1.36% without required bond holdings. Thus, even

[35]Some papers in the literature predict positive bond holdings, particularly later in the life cycle. They often do so by assuming a much higher risk aversion than we do. However, this leads them to predict too much wealth accumulation as compared with US data.

though our main results still hold, the differential welfare gain from ILL over IHI is now significantly smaller.

Next, we consider what would happen if the interest rate wedge between borrowing and lending were smaller, and solve a model with $r_b = 0.05$. This may be applicable if households have access to funding that is cheaper than credit card borrowing, for instance through home equity loans. As borrowing becomes cheaper, the gains from both income-linked assets increase: the baseline IHI now produces a welfare gain of 0.8% while for the baseline ILL with $E(\tilde{r}_{ILL}) = 0.05$ the gain is a very substantial 3.04%. If the mean rate on the ILL is instead assumed to equal 0.08 as before, while $r_b = 0.05$, the ILL produces a welfare gain of 0.52%. Thus, the size of the borrowing wedge is an important determinant of the absolute size of the predicted welfare gain, and if the mean rate on the ILL is much above the rate on other borrowing possibilities the household has access to, the usefulness of ILL is reduced. In this example, in fact, the IHI now generates a higher welfare gain than the ILL. However, we think that this case is less realistic than our benchmark case where $r_b = E(\tilde{r}_{ILL}) = 0.08$, because most household borrowing other than through mortgages occurs at a rate significantly above the return on risk-free saving.

Finally, if we combine the two previous exercises and assume a borrowing rate of 0.05 and that 50/50 stock-bonds funds must be held, the predicted gains are 0.77% for the IHI and 2.24% for the ILL with $E(\tilde{r}_{ILL}) = 0.05$.

Preferences. We additionally quantify the sensitivity of our results to the assumed coefficient of relative risk aversion. Our baseline assumption is that this coefficient equals 2, which seems reasonable from micro studies of consumption behavior, as well as from what it implies for choices over hypothetical gambles. However, the finance literature often assumes much higher risk aversion, in order to justify the observed equity premium. It is important to point out here that in life-cycle portfolio choice models such as the one in this paper, equity holdings usually *increase* in risk aversion over the range of risk aversion parameters that are at least somewhat plausible (say, from 2 to 8). This is because more risk averse individuals accumulate more precautionary wealth to self-insure against their labor income fluctuations, and this more than offsets their lower willingness to invest in risky assets. Hence, if we want to match the empirically observed wealth-to-income ratio (or the observed debt holdings; see Davis et al. 2006), we need to lower the discount factor when increasing the coefficient of relative risk aversion.

Here, we check what happens if we increase the coefficient of relative risk aversion γ from 2 to 3. For the benchmark case without income-linked assets, we need to lower the discount factor β to 0.920 in order to match our target wealth-to-income ratio before retirement. In this case, the welfare gain from our baseline IHI (with correlation –0.5 and volatility 0.5) is now 0.42%, which is significantly larger than the 0.04% in the case of $\gamma = 2$, but still rather small. For the ILL, the effects are again more dramatic. The baseline ILL now produces a welfare gain of 2.42%, which is definitely very substantial for the standards of such models. Nevertheless, the gains from the income-linked assets are still far below the welfare cost of labor income risk, which is 27.4% of certainty-equivalent consumption if $\gamma = 3$.

4 Discussion

In our view, the most important results from the previous section can be summarized as follows: First, the potential use and usefulness of both income-linked assets strongly depends on the characteristics of the assumed return process. For both assets, welfare gains are convex in the assumed correlation between rates and permanent income shocks.[36] The welfare gains from the income-hedging instrument are furthermore very sensitive to the assumed volatility of the return, while the same is the case in less extreme form for income-linked loans.

Second, income-linked loans generally generate higher welfare gains than the income-hedging instrument. The extent of the difference is sensitive to the assumptions about other assets the households can invest in. If an asset with high expected return (such as equity) is available, the welfare gains from income-linked loans become relatively larger while those from income-hedging instruments become smaller.

Third, under some assumptions, namely a low borrowing wedge (a rate on borrowing that is not much higher than the risk-free lending rate) or high risk aversion, the income-linked assets can generate very substantial welfare gains, in excess of two percent of certainty-equivalent consumption.

Our fourth and final main finding is that none of the assets we consider generate a welfare gain that comes close to the 16.4 percent of certainty-equivalent consumption that would be attained under our parameter assumptions if life-cycle income risk were completely eliminated.

To understand these results, we reemphasize the intuition from the earlier two-period model: in a world where borrowing rates are higher than returns on saving, risk management is expensive, especially if a relatively poor household has to put up money upfront to insure against a future contingency (this is the case for the income-hedging instrument, but not income-linked loans). It may then not be worth

[36] The welfare gains are even convex in the *square* of the assumed correlation.

it to do so, especially if insurance is imperfect. And even if one does not need to borrow, putting money in a risk management asset may be inferior to just investing it in other risky assets, such as equity.

Households would not eliminate all income risk even if the correlation between the returns on the income-linked asset and income shocks were perfect. This might seem surprising in light of our assumption that the assets are fairly priced. However, in our constrained life-cycle framework, actuarial fairness must be considered relative to the household's cost of funds. Suppose that there is an income-hedging instrument perfectly correlated with an investor's labor income, and with expected return equal to the riskless rate. If the investor could borrow at the riskless rate, then she would at no cost perfectly hedge away her income risk. But if the investor needs to use a credit card and pays a much higher interest rate, then the insurance requires a costly investment and is no longer actuarially fair to this investor. To be sure, the investor would prefer the income hedging instrument to the riskless asset, but an investor that is borrowing on a credit card should not invest in the riskless asset anyway.

In evaluating income-linked assets, using a model with realistic borrowing and investment opportunities for the households is crucial. If one instead relied on a simpler model in which there is only one other asset, which households can go long or short in, and still assumed that the income-linked assets were priced fairly, one would get very different results. In particular, there would be no difference between income-linked loans and income-hedging instruments, return volatility would not matter, and such a model would predict large welfare gains. For instance, if households could borrow and save at 2%, and the return on the income-linked asset were also 2%, the predicted welfare gains would exceed 4%.

5 Conclusion

In the 1920s, U.S. auto manufacturers targeted cash-poor, lower-income customers as an area of potential growth. While they understood that financial innovation was key to enabling these households to transition to car ownership, the two largest players came up with different designs. General Motors offered a standard loan contract in which the buyer made a downpayment, received a car and committed to make a series of monthly payments for the balance of the price. Ford, in contrast, created a "cash-payment plan" by which the buyer made a series of payments and then received the car. Commentators argue that the relative success of the GM plan was one of the main factors that led Ford by the end of the decade to relinquish its dominant place in the U.S. auto market to GM.[37]

Thus our insight that the specific design of a financial contract, the main finding of this paper, is central to its success has a historical precedent. In our examples, offering income-linked loans typically dominates income-hedging instrument. Just as consumers in the 1920s wanted the car first and to pay it off later, consumers in our model want help with borrowing, not with saving.

But more broadly, the exact form that a financial innovation will take is difficult to predict. In the 1970s, models like ours might have suggested that households should use derivatives to manage interest rate risk. Arguably that prediction would have been right, but the innovation was invisible to households as the derivatives were embedded in fixed-rate mortgages that lenders could only originate because they or their investors could hedge the interest rate risk.

Our paper also highlights the importance of measurement issues, a point made also by Shiller (2003). To use these instruments, one must measure both the risks households face and the covariance of those risks with other risky financial assets. This remains a challenge, and arguably it is this problem that has prevented intermediaries from offering income-linked assets. Measuring the interest-rate risk exposure of a portfolio of financial assets is much easier than measuring the exposure of household income. But the dramatic improvements in computing power and the wide availability of large disaggregated datasets mean that researchers can overcome these challenges.

[37] See Calder (1999), pp. 189-199 for a discussion.

References

Abowd, J. and D. Card (1989). On the covariance structure of earnings and hours changes. *Econometrica 57*, 411–445.

Attanasio, O. and S. Davis (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy 104*, 1227–1262.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review 98*(5), 1887–1921.

Blundell, R. and I. Preston (1998). Consumption Inequality and Income Uncertainty. *Quarterly Journal of Economics 113*(2), 603–640.

Calder, L. (1999). *Financing the American Dream: A Cultural History of Consumer Credit*. Princeton University Press.

Campbell, J., J. Cocco, F. Gomes, and P. Maenhout (2001). Investing Retirement Wealth: A Life-Cycle Model. In J. Y. Campbell and M. Feldstein (Eds.), *Risk Aspects of Investment-Based Social Security Reform*, pp. 439–473. University Of Chicago Press.

Carroll, C. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics 112*(1), 1–55.

Carroll, C. and A. Samwick (1997). The nature of precautionary wealth. *Journal of Monetary Economics 40*(1), 41–71.

Cocco, J. and F. Gomes (2009). Longevity risk and retirement savings. *Working Paper*.

Cocco, J., F. Gomes, and P. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies 18*(2), 491–533.

Cochrane, J. (1991). A simple test of consumption insurance. *Journal of Political Economy*, 957–976.

Davis, S., A. Fuster, and P. Willen (2009). Occupation-Level Income Shocks and Asset Returns: Covariance Structure and Portfolio Choice Implications . *Work in progress*.

Davis, S., F. Kubler, and P. Willen (2006). Borrowing costs and the demand for equity over the life cycle. *The Review of Economics and Statistics 88*(2), 348–362.

Feigenbaum, J. and G. Li (2009). A Semiparametric Characterization of Income Uncertainty over the Lifecycle. *Working Paper*.

Gomes, F., L. Kotlikoff, and L. Viceira (2007). The Excess Burden of Government Indecision. *NBER Working Paper*.

Gomes, F., L. Kotlikoff, and L. Viceira (2008). Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds. *American Economic Review 98*(2), 297–303.

Gomes, F. and A. Michaelides (2005). Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance 60*, 869–904.

Gourinchas, P. and J. Parker (2002). Consumption over the life cycle. *Econometrica 70*, 47–89.

Guvenen, F. (2007). Learning your earning: Are labor income shocks really very persistent? *American Economic Review 97*(3), 687–712.

Guvenen, F. (2009). An Empirical Investigation of Labor Income Processes. *Review of Economic Dynamics 12*(1), 58–79.

He, H. and N. Pearson (1991). Consumption and portfolio policies with incomplete markets and short-sales constraints: The finite dimensional case. *Mathematical Finance 1*, 1–10.

Heathcote, J., K. Storesletten, and G. Violante (2008a). Insurance and Opportunities: A Welfare Analysis of Labor Market Risk. *Journal of Monetary Economics 55*, 501–525.

Heathcote, J., K. Storesletten, and G. Violante (2008b). The Macroeconomic Implications of Rising Wage Inequality in the United States. *NBER Working Paper*.

Heathcote, J., K. Storesletten, and G. Violante (2009). Quantitative Macroeconomics with Heterogeneous Households. *Annual Review of Economics 1*, forthcoming.

Hryshko, D. (2008). RIP to HIP: The Data Reject Heterogeneous Labor Income Profiles. *Working Paper*.

Kaplan, G. and G. Violante (2008). How Much Insurance in Bewley Models? *Working Paper*.

Krueger, D. and F. Perri (2006). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory 1. *Review of Economic Studies 73*(1), 163–193.

Laibson, D., A. Repetto, and J. Tobacman (2007). Estimating discount functions with consumption choices over the lifecycle. *NBER Working Paper*.

MaCurdy, T. (1982). The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of Econometrics 18*(1), 83–114.

Shiller, R. J. (1993). *Macro Markets*. Clarendon Press Oxford.

Shiller, R. J. (2003). *The New Financial Order: Risk in the 21st Century*. Princeton University Press.

Storesletten, K., C. Telmer, and A. Yaron (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics 51*(3), 609–633.

Townsend, R. (1994). Risk and insurance in village India. *Econometrica 62*, 539–591.

Willen, P. and F. Kubler (2006). Collateralized Borrowing and Life-Cycle Portfolio Choice. *NBER Working Paper 12309*.

Zeldes, S. (1989). Optimal consumption with stochastic income: Deviations from certainty equivalence. *The Quarterly Journal of Economics 104*, 275–298.

Zinman, J. (2009). Where is the Missing Credit Card Debt? Clues and Implications. *Review of Income and Wealth forthcoming*.

Figure 1: Asset Holdings and Welfare Gains in Two-Period Model

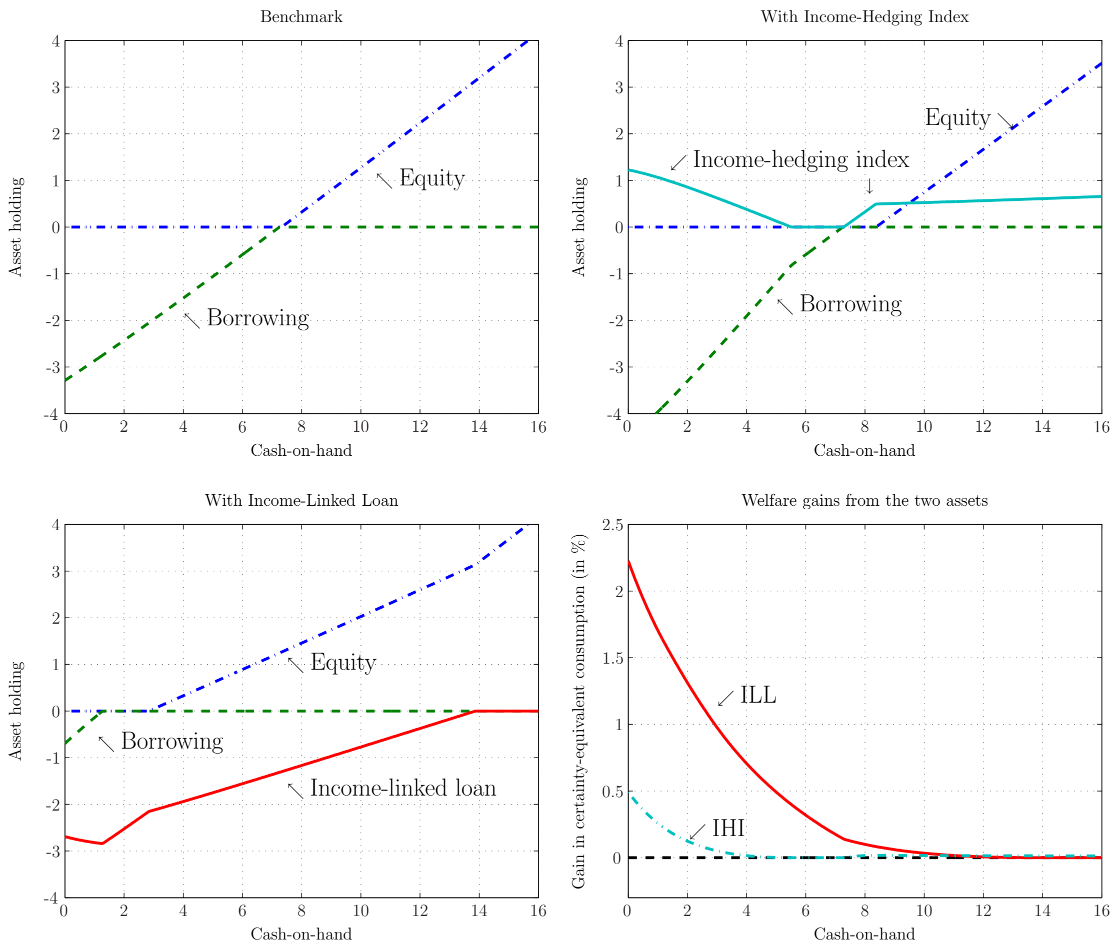

Figure 2: Income Process: Mean Profile and One Realization

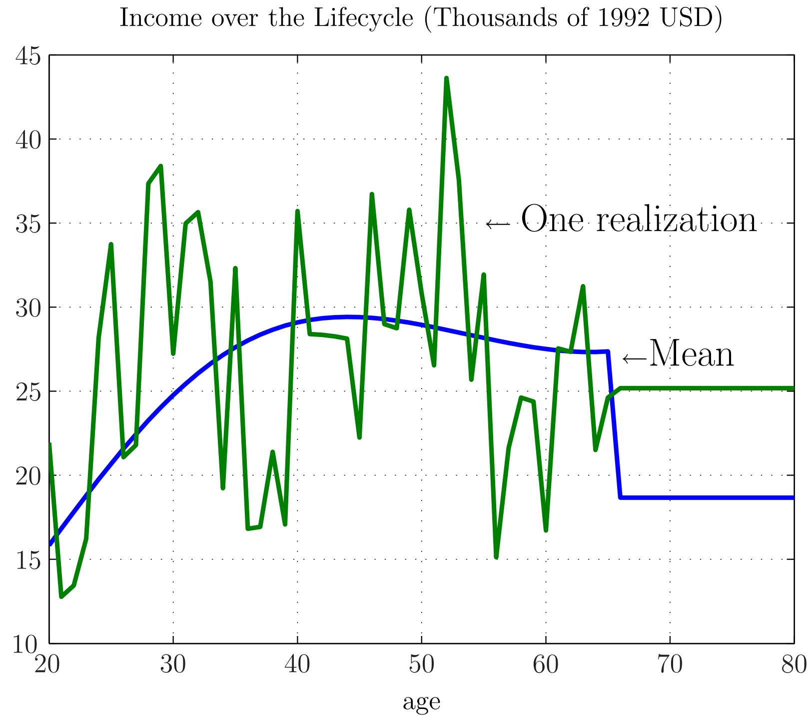

Figure 3: Life-cycle profiles in Benchmark Model

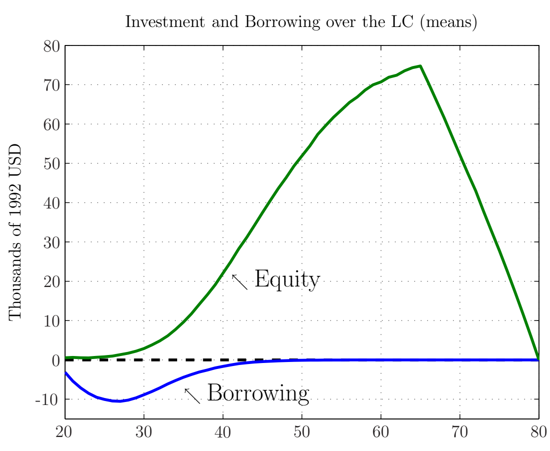

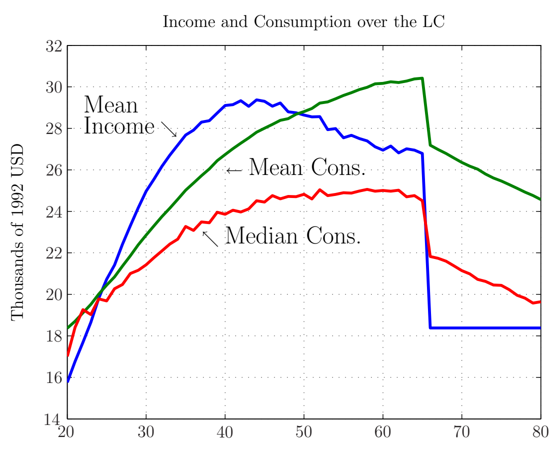

Figure 4: Income-Hedging Instrument: Asset Holdings and Insurance Coefficients for Different Correlations

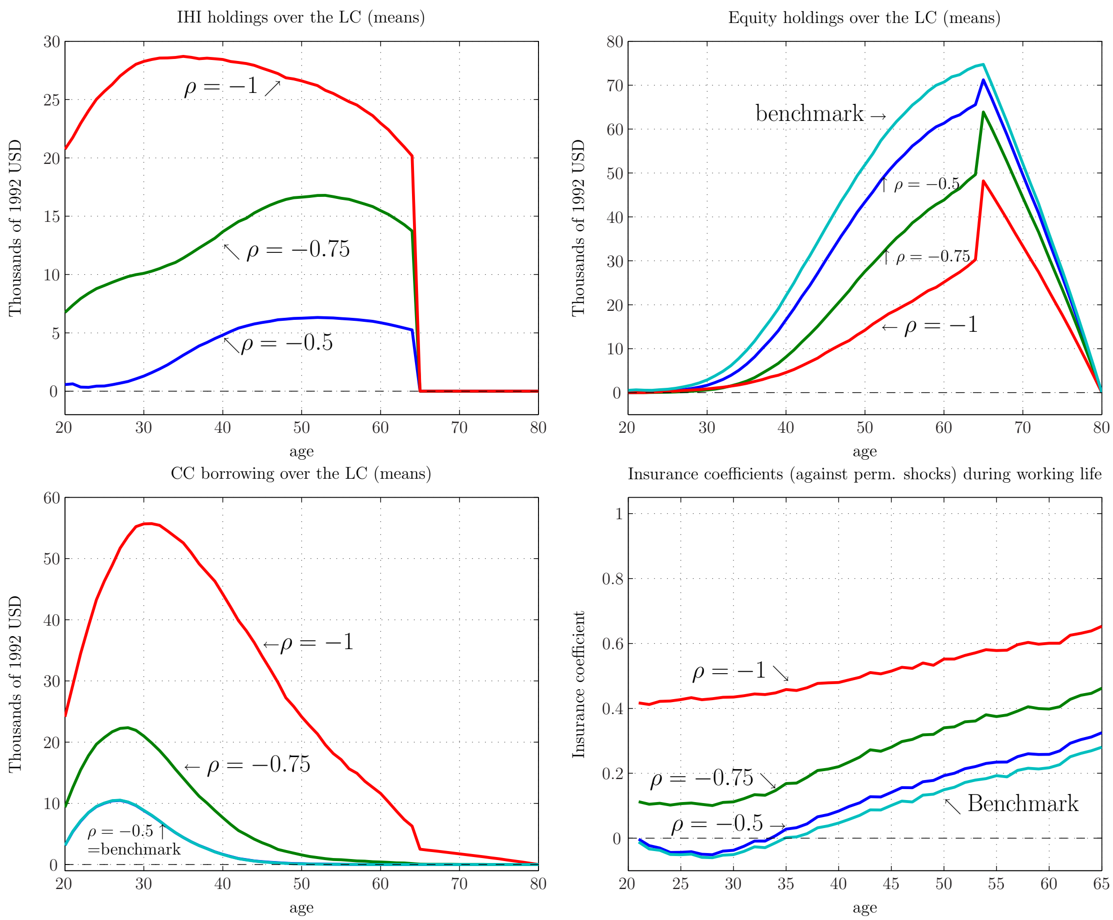

Figure 5: Income-Hedging Instrument: Welfare Gains (over Benchmark)

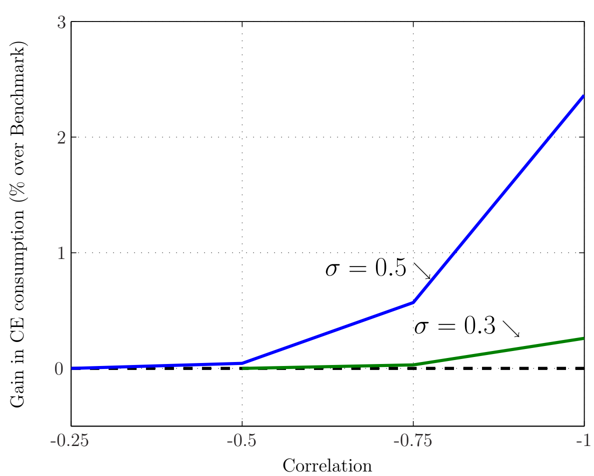

Figure 6: Income-Linked Loans: Asset Holdings and Insurance Coefficients for Different Correlations

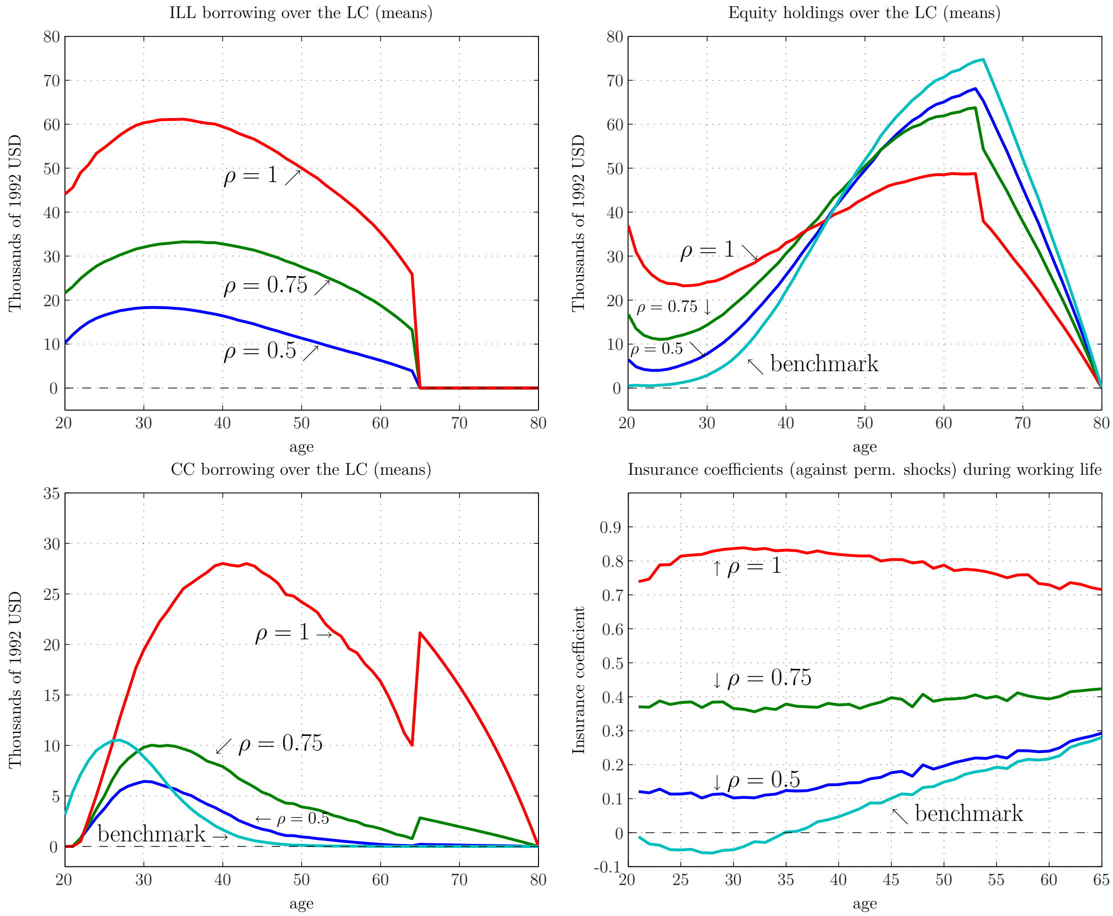

Figure 7: Income-Linked Loans: Welfare Gains (over Benchmark)

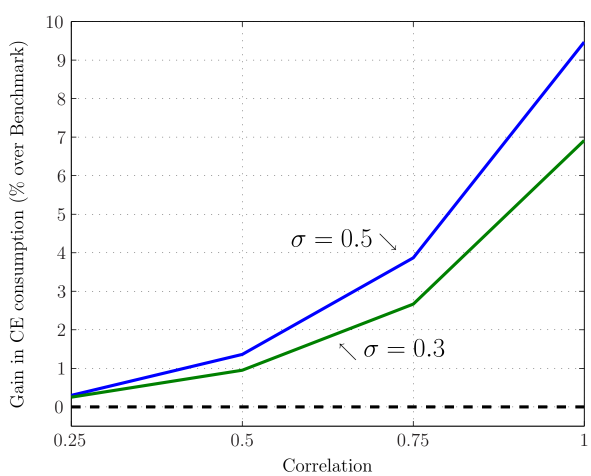

Table 1: Welfare Gains (over Benchmark) in Two-Period Model

Gains in CE consumption, in %, for an investor with cash-on-hand of...				
	0	*5*	*10*	*15*
1) IHI with $\rho = -0.5$, $\sigma = 0.25$	0.49	0.00	0.02	0.01
2) ILL with $\rho = 0.5$, $\sigma = 0.25$	2.23	0.49	0.03	0.00
3) IHI with $\rho = -1$, $\sigma = 0.25$	5.03	0.75	0.29	0.25
4) ILL with $\rho = 1$, $\sigma = 0.25$	8.52	1.87	0.47	0.12
5) IHI with $\rho = -0.5$, $\sigma = 0.5$	1.20	0.20	0.13	0.08
6) ILL with $\rho = 0.5$, $\sigma = 0.5$	2.23	0.69	0.15	0.04
7) IHI with $\rho = -1$, $\sigma = 0.5$	6.94	1.61	0.74	0.51
8) ILL with $\rho = 1$, $\sigma = 0.5$	9.21	2.52	0.91	0.40
No income risk	9.21	2.81	1.40	0.84

Table 2: Parameter values

Parameter	Baseline	Alternative values
RRA γ	2	3
Discount factor β	0.936	0.92, 0.939, 0.947, 0.95
Age of labor force entry	20	
Age of retirement	65	
Age of death	80	
St. dev. of permanent shock σ_ν	0.10296	
St. dev. of transitory shock σ_ε	0.27166	
Replacement rate λ	0.682	
Risk-free lending rate r_l	0.02	
Risk-free borowing rate r_b	0.08	0.05
Mean equity return $E(\tilde{r}_e)$	0.06	0.04
St. dev. of equity returns σ_e	0.16	0.08
Mean return on income-hedging instrument $E(\tilde{r}_{IHI})$	0.02	
St. dev. of IHI return σ_{ILL}	0.5	0.3
Correlation$(\tilde{r}_{IHI}, \tilde{\nu})$	-0.5	-0.25, -0.75, -1
Mean rate on income-linked loan $E(\tilde{r}_{ILL})$	0.08	0.05
St. dev. of ILL rate σ_{ILL}	0.5	0.3
Correlation$(\tilde{r}_{ILL}, \tilde{\nu})$	0.5	0.25, 0.75, 1

Table 3: Summary Table of Results

	Parameters						*Welfare measures*			*Asset holdings (in thousands)*			
Description	$E(\tilde{r}_{IHI})$	ρ_{IHI}	σ_{IHI}	$E(\tilde{r}_{ILL})$	ρ_{ILL}	σ_{ILL}	$\overline{c}$	*Gain (%)*	$\overline{\phi}^{\nu}$	$\overline{IHI}$	$\overline{ILL}$	$\overline{Eq}$	$\overline{CC}$
Benchmark							19638		0.09			33.82	2.33
1) IHI	0.02	-0.50	0.50				19646	0.04	0.12	2.98		29.51	2.34
2) IHI with higher corr.	0.02	-0.75	0.50				19750	0.57	0.26	9.71		21.61	6.32
3) IHI with perfect corr.	0.02	-1.00	0.50				20102	2.36	0.51	19.22		13.81	25.28
4) IHI with lower corr.	0.02	-0.25	0.50				19638	0.00	0.09	0.00		34.00	2.34
5) IHI with lower volatility	0.02	-0.50	0.30				19638	0.00	0.09	0.01		33.99	2.34
6) ", higher corr.	0.02	-0.75	0.30				19644	0.03	0.12	4.24		27.56	2.34
7) ", perfect corr.	0.02	-1.00	0.30				19689	0.26	0.25	11.45		18.44	3.99
8) ILL				0.08	0.50	0.50	19905	1.36	0.17		9.68	33.20	1.86
9) ILL with higher corr.				0.08	0.75	0.50	20398	3.87	0.39		20.20	33.97	4.19
10) ILL with perfect corr.				0.08	1.00	0.50	21497	9.47	0.79		37.50	31.85	17.20
11) ILL with lower corr.				0.08	0.25	0.50	19697	0.30	0.09		2.00	32.92	1.35
12) ILL with lower volatility				0.08	0.50	0.30	19825	0.95	0.11		7.75	30.91	0.47
13) ", higher corr.				0.08	0.75	0.30	20161	2.67	0.26		23.21	32.86	1.11
14) ", perfect corr.				0.08	1.00	0.30	20996	6.91	0.59		51.52	33.14	6.77
15) ", lower corr.				0.08	0.25	0.30	19687	0.25	0.09		2.43	32.77	0.62

Table 4: Summary Table of Robustness Checks

Preferences		*Assets*				*Welfare measures*		
γ	β	Stocks/Bonds	ILA	$E(\tilde{r}_{ILL})$	r_b	$\overline{c}$	*Gain (%)*	$\overline{\phi}^u$
2	0.947	50/50	-	-	0.08	19515	-	0.08
2	0.947	50/50	IHI	-	0.08	19579	0.33	0.17
2	0.947	50/50	ILL	0.08	0.08	19681	0.85	0.12
2	0.939	free	-	-	0.05	20092	-	0.05
2	0.939	free	IHI	-	0.05	20254	0.80	0.13
2	0.939	free	ILL	0.05	0.05	20703	3.04	0.23
2	0.939	free	ILL	0.08	0.05	20196	0.52	0.12
2	0.950	50/50	-	-	0.05	19688	-	0.07
2	0.950	50/50	IHI	-	0.05	19839	0.77	0.18
2	0.950	50/50	ILL	0.05	0.05	20129	2.24	0.22
3	0.920	free	-	-	0.08	18244	-	0.08
3	0.920	free	IHI	-	0.08	18320	0.42	0.16
3	0.920	free	ILL	0.08	0.08	18686	2.42	0.18