Working Papers

FEDERAL RESERVE BANK OF BOSTON

No. 10-1

Insuring Consumption Using Income-Linked Assets

Andreas Fuster and Paul S. Willen

Abstract:
Shiller (2003) and others have argued for the creation of financial instruments that allow households to insure risks associated with their lifetime labor income. In this paper, we argue that while the purpose of such assets is to smooth consumption across states of nature, one must also consider the assets' effects on households' ability to smooth consumption over time. We show that consumers in a realistically calibrated life-cycle model would generally prefer income-linked loans (with a rate positively correlated with income shocks) to an income-hedging instrument (a limited liability asset whose returns correlate negatively with income shocks) even though the assets offer identical opportunities to smooth consumption across states. While for some parameterizations of our model the welfare gains from the presence of income-linked assets can be substantial (above 1 percent 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.

JEL Classifications: D91, E21, G11
Keywords: income risk, risk sharing, portfolio choice, financial innovation

---

Andreas Fuster is a graduate student in the economics department at Harvard University and a graduate fellow in the research department at the Federal Reserve Bank of Boston. His email address is afuster@fas.harvard.edu. Paul S. Willen is a Senior Economist and Policy Advisor at the Federal Reserve Bank of Boston. His email address is paul.willen@bos.frb.org

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at http://www.bos.frb.org/economic/wp/index.htm.

We are grateful to John Campbell, Nicola Fuchs-Schündeln, Rustom Irani, María Luengo-Prado, Claudio Michelacci, Ali Ozdagli, Julio Rotemberg, and seminar participants at Harvard, the First Annual Workshop of the Zurich Center for Computational Financial Economics, and the Household Finance and Macroeconomics conference at the Banco de España for helpful comments and discussions.

This version: February 24, 2010

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, ultimately owe their existence 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 payoffs 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 consumption and portfolio choice, in which households can invest in different assets and can borrow but at a substantial premium to the riskless rate of return. We find that the benefits of income-linked assets can be sizeable, but are highly sensitive to the precise design of the assets. For example, the seemingly innocuous decision of whether to link income negatively to the return on a savings instrument or positively to the interest rate on a loan has a large impact on the welfare gains our model predicts. The intuition here is that in a world with realistic frictions, one cannot separate the effect of an asset on consumption across states of nature with the effect on consumption across time. For instance, an asset that needs to be purchased today to provide insurance against future shocks reduces the variability of future consumption but also increases future consumption relative to current consumption. In a frictionless world, households can borrow to undo such an intertemporal distortion, but in a realistic world where borrowing is expensive, households may not want to invest in such an asset in spite of the reduction in future consumption variability.

The reason why such assets deserve consideration is that risk sharing between households is limited, as evidenced by the fact that much of observed household consumption volatility is due to idiosyncratic income shocks, not aggregate income fluctuations. This limited risk sharing can be seen either as a puzzle or as evidence of frictions. 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.

However, in principle there is scope for sharing part of one's labor income risk without inducing moral hazard, as part of this risk is *group-specific*. Groups could be defined, for instance, in terms of occupation, industry, region, or education level. As an example, an auto worker may experience an income shock because of his individual job performance, but also because the overall evolution of the demand for cars affects the average wage of auto workers, which is beyond his control. Shocks of the second kind are observable and verifiable, so that insurance-like contracts based on such risks can, in principle, easily be written. Indeed, Attanasio and Davis (1996) argue for the "puzzle" view of the limited risk sharing present in the data precisely because they find that a particular group-level shock — income variation at the education-sex-birth cohort level — appears not to be shared across households. One reason why the sharing of such risks that are immune to moral hazard may be difficult is the possible importance of another friction, namely limited commitment: if two individuals 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 the promise, even if afterwards he gets punished by being excluded from future similar transactions. This issue can be mitigated by moving from bilateral contracts to trading such risks through long-lived institutions that hold a diversified portfolio and have strong reputational concerns, so that they are very unlikely to default. This solution makes the market-based approach considered in this paper a potentially promising way to share household income risks.

The income-linked assets we consider take two basic forms. The first is a standard insurance-like contract in which an individual pays something now for an asset whose future payoff is negatively correlated with the individual's income innovations — we call this an "income-hedging instrument." Thus, if the individual's group receives a negative income shock, this asset will pay off more than if the group receives a positive shock. The second form we consider are "income-linked loans," where the required repayment is positively correlated with one's group's income shocks. Either way, the upshot of adding such assets to a household portfolio would be to reduce consumption fluctuations. Our contribution in this paper is a quantitative evaluation of what households' demand for such assets would be, and which design features of the assets this demand most strongly depends on. Furthermore, we assess the size of 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.[1] 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."[2] Of course, our model-based approach is precisely an attempt to predict what might happen in an alternate world. Understanding the potential welfare gains from such assets, and what these benefits 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 use a finite horizon, partial equilibrium model which roughly matches basic facts about households' 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 significant challenge for this research is the need to make assumptions about the return characteristics of assets that do not yet exist. For the mean returns of the income-linked assets, we make the baseline assumption that the risks upon which the payoffs are based are purely cross-sectional, such that the assets can be priced fairly. Thus, we assume that the mean return on the income-hedging instrument equals the mean return on the risk-free bond, and that the mean interest rate to be paid on the income-linked loan is equal to the interest rate on other unsecured

[1] We only consider a subset of Shiller's proposals. Our income-hedging instrument can be seen either as "livelihood insurance" or as a particular example of a "macro market." In addition to these two, and the income-linked loans that we also look at, Shiller's other suggestions 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."

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

household debt.[3] For the other return characteristics, we remain relatively agnostic and simply plug in different values for the volatility of the returns on the income-linked assets 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 calibrated model yields two main results. The first is that the benefits that income-linked assets could generate for households are very sensitive to the parameters of the return process. Most importantly, potential welfare gains are strongly convex in the assumed correlation between rates and income shocks. As a consequence, unless the correlation is 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. The attractiveness of our assets further depends on the assumed return volatility, with higher volatility providing "more bang for the buck" for households. The size of the cost differential between borrowing and lending is also very important: the larger it is, the less households gain from having access to the proposed income-linked assets.

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 percent (an increase in consumption of about 400 USD per year, in 2009 dollars) while the income-hedging instrument is essentially worthless. We also explore the boundaries of this result. For instance, we show that the attractiveness of the alternative investment option matters for the relative appeal of the two income-linked assets: the presence of equity (as in our baseline) makes the income-linked loan relatively more attractive (as households can invest some of the borrowed money in a high return asset) while the income-hedging instrument is in less

[3] 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 mean that misspecification of either the household decision problem or the general equilibrium could lead to inaccurate predictions about the benefits of the assets. 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.

demand than if equity were not available. If a household has access to borrowing at a cheap rate (lower than what it would have to pay on average on the income-linked loan), the ranking of the two assets may be reversed, such that the income-hedging instrument is more valuable. However, even under such assumptions, the gains from the income-hedging instrument remain moderate (below 1 percent). In sum, we find that under some assumptions, the gains to households from having access to income-linked loans could be significant, 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. However, even though the two assets are equally attractive in terms of *intra*temporal consumption smoothing, whether households will demand them also depends on how the assets square with households' desire for *inter*temporal smoothing. We show that over the life cycle, income-linked loans are more attractive in that regard than income-hedging instruments. The reason is that early in the life cycle, most households' main financial activity is high-interest rate borrowing (because they want to consume part of the higher income they expect in the future), for 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 on the income-hedging instrument must exceed the risk-adjusted return on equity, and if we set expected equity returns 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.

Given that some calibrations of our model predict substantial benefits from income-linked loans, an obvious question is why such loans are not more frequently observed in the real world.[4] We first reiterate that we assume here that the risks households are hedging are both observable and cross-sectional. The former implies that there are no adverse selection or moral hazard problems and when combined with the latter, means that no risk premium needs to be added to the risk-free rates.[5] 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 (2003) advances another reason for the current nonexistence of the income-linked assets he envisions, arguing that until recently we did not have the technology necessary to collect and maintain the data underlying the various proposed instruments. He points out that there will likely be a need for government intervention to help 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—for instance, a change in the bankruptcy law, such that income-linked loans could not be canceled by declaring personal bankruptcy.[6]

The rest of this paper is organized as follows: In the next subsection, we discuss some of the related literature. 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.

[4] Some examples of particular forms of income-linked loans do exist in the real world. For instance, in Australia and some other countries, there exist education loans for which the required repayment is based on subsequent labor income. In the United States, 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.

[5] As our focus is on the benefits that income-linked assets could generate for households, we do not directly address what entities would be willing to offer these assets at such prices. One candidate would be investment firms that already manage the retirement funds of people in a wide variety of occupations. Alternatively, income-hedging instruments could be operationalized as exchange-traded products, or they could be offered by insurance companies (which would presumably add loadings and thus offer rates that are less than actuarially fair).

[6] Income-linked loans would then have a similar status as student loans. This change in the bankruptcy law may be necessary because otherwise, individuals are tempted to default on their income-linked loans after they receive a positive income shock.

1.1 Related Literature

Risk sharing is one of the fundamental topics of economics and finance and therefore much too vast to be surveyed here. We will therefore concentrate on referencing some classic and recent papers that we deem particularly relevant or related to our analysis. We begin by discussing papers that focus on measuring the extent of risk sharing or on the welfare cost of income or wage risk, and then consider papers on households' life-cycle portfolio choice that are more closely related to our model.

Classic empirical papers on the degree of risk sharing between households include Cochrane (1991), Attanasio and Davis (1996), and Hayashi, Altonji, and Kotlikoff (1996). These authors all attempt to assess the degree to which household consumption is insured against shocks to income, and find that such insurance is far from perfect (if it were not, there would of course be no need for new financial instruments that facilitate better insurance). In a related vein, 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 examine the evolution of income risk and inequality over the past decades, the degree of households' insurance against income risk, and the different channels through which such insurance can be achieved.

In quantitative dynamic macro models that are calibrated to match empirical data[7], the welfare cost of income uncertainty is typically very large. For instance, Storesletten, Telmer, and Yaron (2004) use a model that is calibrated to match the empirically observed evolution of household 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 percent of lifetime consumption in exchange for insurance against all life-cycle shocks. Pijoan-Mas (2006) studies a general equilibrium production economy with infinitely-lived agents, flexible labor supply, and stationary wage shocks, and finds that complete markets (meaning full insurance against wage risks) would lead to a welfare gain equivalent to increasing lifetime consumption by 16 percent. In a related paper, Heathcote, Storesletten, and

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

Violante (2008a) find that in a model with permanent and transitory wage shocks the welfare gain from complete markets would be almost 40 percent of expected lifetime consumption.[8] Their conclusion is thus the following: "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" (p. 520).

The papers referenced above and, in general, most quantitative general equilibrium macro models, only feature a rather simple asset market structure (often composed of 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, an assumption that is somewhat intellectually unsatisfying and may also yield misleading results in counterfactual exercises.[9] Our model builds on other computational analyses of optimal portfolio choice over the life cycle, some well-known examples of which include Bertaut and Haliassos (1997), 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. Perhaps closest in spirit to our work are recent papers by De Jong, Driessen, and Van Hemert (2008) and Cocco and Gomes (2009).[10] De Jong, Driessen, and Van Hemert consider the welfare benefits generated by the presence of housing futures and find these gains are small (mostly due to the significant fraction of house price risk that is idiosyncratic). Cocco and Gomes investigate the role that longevity bonds (for which no liquid market currently exists) could play in individual portfolios, what the welfare benefits from such bonds would be, and the optimal design of such bonds.

[8] They emphasize that in their model, 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.

[9] 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.

[10] Other examples include Campbell, Cocco, Gomes, and Maenhout (2001) and Gomes, Kotlikoff, and Viceira (2007, 2008).

2 Two-Period Model

To gain some intuition for how the introduction of the income-linked assets might affect households' asset portfolios 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[11], and then look at an example with a calibration similar to the one we will be using in the life-cycle model. As will be shown 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}. The 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, the optimal policy can be understood in terms of a simple algorithm. Start with any admissible asset allocation, which will imply a consumption stream $\{c_{t+1}^1, ..., c_{t+1}^S\}$ for the S states of the world. Based on this consumption stream, one can define the "risk-neutral" or "martingale" probability measure, which reweights the objective probabilities of the different 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 its "risk-adjusted" (gross) return using these risk-neutral probabilities: $\mathrm{E}_Q[\tilde{R}_i]$. Intuitively, this measure adjusts an asset's mean return by how

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

useful it is for consumption smoothing across states of the world. Thus, an asset that has a high payoff in states of the world in which consumption is low, and marginal utility therefore high, has a higher risk-adjusted return than an asset that has the same average return but pays off more in states of the world in which consumption is high.

Next, define the "shadow rate" as

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

This is the minimum interest rate at which an investor would be willing to decrease his consumption in period t by a small amount ϵ if in return he received $R\epsilon$ in period $t+1$ (or, equivalently, the highest rate at which he would be willing to borrow ϵ for consumption in t if he had to repay $R\epsilon$ in $t+1$). This rate is higher the more the investor expects consumption to grow between t and $t+1$, and (for utility functions such that marginal utility u' is convex) the less uncertain he is about consumption in $t+1$. Thus, the shadow rate is influenced by the investor's desire to smooth consumption across periods and across states of nature within a period.

Optimal portfolio choice can then be characterized by the following simple rule: investors should optimally add to (subtract from, leave unchanged) their 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, the portfolio's overall optimality requires $\mathrm{E}_Q[\tilde{R}_i] = R \; \forall i$. Furthermore, if the set of available assets includes one that is risk-free, so that investors 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 to fund more investment.

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 a world with incomplete markets, we would 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 very 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 the investors' shadow rate. An asset's risk-adjusted return depends on how helpful it is for consumption smoothing across states, while investors' shadow rates are driven by their desire to smooth consumption across states and time. The shadow rates thus depends on investors' 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 starts life with some cash-on-hand in period 1 and will receive a stochastic income in period 2 with mean 8 and standard deviation 1.5.[12] The investor has an isoelastic utility function with relative risk aversion of 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, and invest in equity with an expected return of $E(\tilde{r}_e) = 6$ percent and a standard deviation of 16 percent.[13]

The top left 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 if he is relatively poor in period 1, and saves (by investing in equity) if he is relatively rich. Using the terminology from 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 nor engages in risk-free saving (but he would do so if he were more risk averse or had very high cash-on-hand).

The top right panel shows the optimal asset holdings if in addition to the assets from the benchmark model, the investor has access to an IHI with $E(\tilde{r}_{IHI}) = r_l = 2$ percent, standard deviation 25 percent, and a negative return correlation of 0.5 with second-period income. Thus, the IHI tends to pay off more when the investor experiences a negative income shock and pays less if his income exceeds expectations. The optimal policy features positive holdings of the IHI at low levels of cash-on-hand, financed by additional borrowing. As cash-on-hand increases, holdings of

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

[13] 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 other than losses from defaults for why borrowing costs exceed lending rates, such as transaction costs or the cost that lenders face in the screening of potential borrowers (with the goal of lowering default risk).

the IHI decrease, and for cash-on-hand levels between 5.5 and 7.3 equal zero. At higher levels of cash-on-hand, the IHI holdings become positive again, and eventually the investor simultaneously holds both 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 case.

In the lower left 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 25 percent, and a positive correlation of 0.5 with income in period 2. Thus, when taking out an ILL, the investor will need to repay a larger amount if his income is higher than expected in the next period and a lower amount if his 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, but does not lead to much additional total borrowing.[14] For a large intermediate range of cash-on-hand, however, there is now more borrowing (through the ILL) than there was at the fixed rate. Also, over this range, 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 cash-on-hand. While this may be surprising at first glance, it is a general feature of portfolio choice problems with short-selling constraints. The reason is that due to these constraints, and the types of assets available, the shadow rate may stay constant over some ranges of cash-on-hand but decreases over others, and sometimes discontinuously falls when a constraint is hit.

The lower right panel in figure 1 displays the welfare gains from having access to one of the two income-linked assets in this example. Welfare is measured in terms of certainty-equivalent (CE) consumption, which is defined as the constant consumption stream that would provide the same lifetime utility as the risky stream the investor actually expects. 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

[14] 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.

range depicted, the ILL provides higher welfare gains over the benchmark case than does the IHI. Here is some intuition for why this is the case. First consider 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 the ILL 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 his 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 such, the insurance provided by the ILL becomes somewhat more "expensive." For high cash-on-hand levels, ILL borrowing decreases to zero, so that investors in that range do not gain from having access to the ILL. 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 the 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).[15] 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 the 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 if only risk-free saving were available). Indeed, if no equity were available in our example, relatively rich investors (with cash-on-hand above 7.3 in this example), would hold more of the IHI, and borrow less through the ILL 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 shown in the figure.[16]

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 is certain to receive an income

[15]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.

[16]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.

of 8 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 CE 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 result? Part of the explanation is provided by the imperfect correlation of the rates of return with the income shock. As table 1 shows, if the rates are perfectly correlated with the income shock, the gains provided by the two income-linked 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, the ILL is more useful while for rich ones, the IHI leads to higher welfare gains.

Table 1's lines 5–8 further show the effect more volatile IHI or ILL returns will have on welfare gains. Clearly, given our earlier discussion of opportunity costs, it is not surprising 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 strongly affected by the increase in 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 show that the main points discussed in this section 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 we find will be rather small for the parametrizations we deem most realistic (particularly as compared with the hypothetical gain from completely eliminating all income risk), and income-linked loans generally appear to be more promising than the income-hedging instrument.

3 Life-Cycle Model

3.1 Setup

Our strategy in this part is as follows: 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 new assets into this model, one at a time, and analyze what the demand for these assets would be, how these would affect the demand for the other assets in the model, and what the predicted welfare gains from the new assets would be.

The specific portfolio choice model that we build on is the one by Davis, Kubler, and Willen (2006). This model explicitly accounts for 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 to earn higher incomes in the future, typically take advantage of this borrowing opportunity in order to smooth their consumption over time. However, this borrowing slows down the speed at which households accumulate wealth, and reduces their rate of participation in equity markets until about age 45. As a consequence, this model generates more realistic predictions than models that allow for 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: $\tilde{y}_t = \log(\lambda) + d_{t_R} + \eta_{t_R}$ for $t > t_R$.

The household maximizes expected utility over its remaining lifetime,

$$U(\alpha_t c_t) + E_t \sum_{s=t+1}^{T} \beta^{s-t} U(\alpha_s c_s), \tag{2}$$

in each period t, where $U(\cdot)$ is an isoelastic (power) utility function with curvature γ, β is the constant discount factor, and α is a "taste shifter" that we include mainly to account for the drop in consumption when entering retirement.[17] 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 following expression gives the budget constraint of a household at age t, in its most general form:

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

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

Households can always trade at least three financial assets. They can buy equity (e) with a stochastic return $\tilde{r}_e$ ($= \tilde{R}_{e,t} - 1$), save (l) at a net risk-free rate of return r_l, and borrow (b) at a fixed risk-free interest 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 a 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

[17] The taste shifter can be seen as a stand-in for a more elaborate model with labor supply. For instance, Cocco and Gomes (2009) use $\alpha_t = L_t^{\zeta}$, which generates a consumption drop at retirement, because leisure and consumption are substitutes in the utility function if $\gamma > 1$.

of the interest rate, to see how these parameters affect the demand for and the 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 only consider assets that correlate with the household's *permanent* income shock, 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.

We do not impose an exogenous borrowing constraint, but require that households be 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, in our model households never default on their debt. 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 leads 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.[18]

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 demand the model predicts for these assets (meaning 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 compute the (ex-ante) lifetime expected utility $\bar{U}$ in a given environment. Then, we find the

[18]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 primary residence, 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.

constant level of consumption, $\bar{c}$, that would yield the same level of utility:

$$\left(\sum_{t=0}^{T-1}\beta^t\right)\frac{\bar{c}^{1-\gamma}}{1-\gamma}=\bar{U}\Leftrightarrow\bar{c}=\left(\frac{1-\beta}{1-\beta^T}(1-\gamma)\bar{U}\right)^{\frac{1}{1-\gamma}}. \tag{3}$$

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

$$\phi_t^{\nu}=1-\frac{cov(\Delta c_{it},\nu_{it})}{var(\nu_{it})}, \tag{4}$$

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 one-for-one 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 experience.

3.3 Calibration and Discretization

Table 2 gives an overview of the parameter values we use to calibrate the model. For the labor income process, we use the parameters from Cocco, Gomes, and Maenhout (2005) for high school graduates, which in this literature have been accepted as somewhat of a standard. The deterministic component of income, 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 immediately after age 65, and die with certainty at age 80. Figure 2 displays the mean income over the life cycle as well as one simulated realization (to give a sense of the significant extent of income uncertainty households face).

The preference parameters in our main calibration are set as follows: relative risk aversion γ

(equal to the inverse of the intertemporal elasticity of substitution for the assumed isoelastic utility function) 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 model's benchmark case 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 1 before retirement and 0.9 afterwards; this generates a mean consumption drop at retirement of about 10 percent, which is consistent with most empirical estimates.

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. For simplicity, equity returns are assumed to be uncorrelated with labor income shocks, though relaxing this assumption and setting the correlation equal to an empirically reasonable value (for instance 0.15, as in Gomes and Michaelides 2005) has very little effect on our results. The interest rate on risk-free borrowing, r_b, is set equal to 8 percent, which is what Davis, Kubler, and Willen (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, Kubler, and Willen (2006). 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 shocks, 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]

[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.

[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 altogether. We believe that it is realistic to assume that there is a positive lower bound for income shocks, due to the presence of

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

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 gains from 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 model's computational solution easier.

The main feature of the permanent-transitory income process is that there is no 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 otherwise would be puzzling. In Guvenen's model, income shocks are less persistent, but individuals only learn about their personal d_t over time, through observation of their realized income. While his model intuitively makes a lot of sense, it is very computationally demanding; to our knowledge, nobody has solved a version that contains more than the risk-free asset. Furthermore, Hryshko (2009) 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

social safety nets.

comparability with the existing portfolio choice literature as well as computational tractability, we stick with the status quo.

Even once the form of the income process has been determined, there remain calibration decisions that are crucially important 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, Gomes, and Maenhout (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 the 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 estimate is somewhat too high, its slope over different horizons seems much more appropriate than if the Gourinchas and Parker numbers were used.

Importantly, Cocco, Gomes, and Maenhout (and most other articles in this literature) use a broad definition of labor income to estimate the variances of permanent and transitory income shocks. In particular, the following additional sources of income are included in the measure of labor income they use: unemployment compensation, workers' compensation, social security, supplemental social security, other welfare, child support, and total transfers (mainly help from relatives). Also, the incomes of both the household head and the spouse (if present) are included. Thus, this income measure accounts for many implicit and explicit insurance mechanisms other than asset accumulation that are already available to households.

[21] They note that this fact by itself may be an indication that the persistent shock follows an autoregressive 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, Gomes, and Maenhout 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.[23] 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 $\omega_{it} \sim N(-\sigma_\omega^2/2, \sigma_\omega^2)$, such that $\nu_{it} = \xi_t + \omega_{it}$.[24] If we assume that the return on the income-linked asset (ILA) is perfectly correlated with the group-specific permanent shock ξ_t, then the correlation of the return with individual i's permanent income shock is given by

$$corr(\tilde{r}_{ILA,t}, \nu_{it}) = \frac{\sigma_\xi}{\sigma_\nu}. \tag{5}$$

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 (2010), 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) for 17 occupational classifications that have remained largely unchanged for at least 35 years and for which we have a relatively large number of individuals in each survey year.[25] Although these occupations are not necessarily a representative set of oc-

[23] In a related exercise, Shiller and Schneider (1998) use PSID data from 1968–1987 to construct group-level income indices after first identifying occupation-industry groups such that only few people transition from one group to another over time. Their preferred grouping procedure yields seven distinct groups, such as "Professional/Technical" or "Agriculture/Labor." They find that changes in the index of the group an individual belongs to explain 40–50 percent of individual nominal income changes at a one-year horizon (and more at a five-year horizon) after controlling for hedonic variables, while the consumer price index only explains 20 percent.

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

[25] 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.

cupations for the U.S. population, we can at least get an estimate of the order of magnitude of occupation-specific income shocks. If for simplicity we assume that all occupation-level income shocks are permanent (an assumption that is not too far from the truth in our data — see Davis, Fuster, and Willen 2010 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 the 10 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 0.103, this average implies 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.[26] 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" indexes could be constructed which would be more highly correlated with individual income shocks (for instance, "Plumbers located 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, Kubler, and Willen (2006): households borrow substantial amounts while they are young (on average, 50 percent of their annual income between ages 20 and

[26] 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.'s estimate of the permanent shock standard deviation for high school graduates in the industry in brackets): Manufacturing (0.068), Construction (0.120), Trade (0.106) and Transportation (0.067). 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.

30)[27] and only start making substantial investments in the stock market after age 35. The predicted equity market participation rate starts out around 20 percent for young households and increases through mid age, reaching 95 percent at age 45. Average equity holdings at retirement amount to about three times annual income; this is about twice as high as in the data. However, this is arguably not a major failing of the model, as the model does not feature home equity, which in reality is a risky asset held by most households. Other features of the model are that it predicts practically no borrowing for households older than 40 years, 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.[28] (In section 3.4.4, we will consider a version of our model in which households are forced to invest at least 50 percent 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.[29]

For our benchmark case, CE consumption equals 19,638 USD.[30] The partial insurance coefficient averages 0.09 over the life cycle. This is significantly lower than what Kaplan and Violante (forthcoming) find in their model, which features more redistributive social security, and is even further below the baseline insurance coefficients that Blundell, Pistaferri, and Preston (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. On the other hand, the main insurance coefficients reported by Blundell, Pistaferri, and Preston (2008) and used as a benchmark by Kaplan and Violante (forthcoming) may give too optimistic a view of "true" insurance, as these coefficients are computed only from nondurable consumption. However,

[27]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 percent for below-30-year olds in 1995 and 1998 (Davis, Kubler, and Willen 2006). However, Zinman (2009) finds that the SCF misses around one-half of revolving debt.

[28]The continued credit card borrowing might also be due to consumers having self-control problems, as in Laibson, Repetto, and Tobacman (2007).

[29]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.

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

households may respond to income shocks largely by changing their expenditures on durables, which then affects the utility flows they get from these goods.[31]

The welfare cost to households imposed by income shocks is fairly high in our 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 percent more.[32] Thus, the ex-ante cost of income shocks is high, and of a similar order of magnitude as what is found in the quantitative macro literature discussed earlier. Thus, one would hope that introducing financial assets for households to hedge the risk of these income shocks could yield high welfare gains.[33]

3.4.2 Income-Hedging Instrument

Our baseline assumption for the mean return on the IHI 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, using the empirical evidence discussed in section 3.3.2 as a guide. The baseline for the standard deviation of returns is 0.5, meaning that the return on the IHI 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 to be arbitrarily volatile, we chose this high volatility as our benchmark.[34]

Table 3 summarizes the results, while figure 4 shows the mean life-cycle holdings of the IHI, equity, and unsecured borrowing (denoted by "CC", which stands for "credit cards") for the baseline case and also for two cases with higher (absolute) correlation.[35]

[31] Blundell, Pistaferri, and Preston present some evidence that is consistent with this idea. As part of their sensitivity analyses, they consider a measure of total expenditure and find that, at least for low-wealth households, there appears to be much less (indeed, no) insurance against permanent shocks in this case than when only nondurable consumption is considered.

[32] 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.

[33] 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 percent less than in the benchmark with equity.

[34] The volatile version of our IHI also approximates a classical insurance contract, which pays off at all only in a small number of states of the world.

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

Our first finding is that there will not be high demand among households for the baseline IHI. As the top left panel of figure 4 shows, mean holdings of the IHI never go much above 5,000 USD, and as the top right and lower left 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 of this asset until age 33). Mean IHI holdings peak around age 50, then slowly decline as households move 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 against permanent income shocks. Particularly for young households, which are not insured against shocks to permanent income, the degree of insurance is virtually unchanged by the availability of the baseline IHI.[36] Even for older households, the increase in the degree of insurance is rather small.

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 the IHI is higher, but still starts out relatively low for younger households. 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 higher future income. As the bottom right panel shows, in this case insurance against the permanent income shock is much improved, even though it is still only around 0.5 on average over the working life.

IHI return volatility matters greatly for mean holdings of the IHI as well as the effect on equity holdings and borrowing.[37] For the baseline correlation of –0.5, for example, there are practically no IHI holdings when the standard deviation of IHI returns is only 0.3 (nobody participates in

[36]It may be surprising that the insurance coefficient is slightly negative for young households in the benchmark. Kaplan and Violante (forthcoming) explain that this is due to the interaction of permanent and transitory shocks in this model (see their footnote 30).

[37]This can be seen by comparing lines 5–7 of table 3 to lines 1–3.

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 USD. The partial insurance coefficients are also very significantly reduced.

Figure 5 displays the welfare gains over the benchmark case (in percent of CE 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 between the IHI return and the permanent shock, as well as the volatility of the IHI return. Welfare gains are convex in the strength of the correlation.[38] The welfare gains are tiny (below 0.1 percent) if the correlation 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. As in the two-period model in section 2, the reason behind this surprising finding lies in households' effective cost of funds — the borrowing costs for young households and the opportunity costs due to the possibility of investing in equity for older households.

3.4.3 Income-Linked Loans

For the ILL, 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 interest rate and its correlation with the household's permanent income shocks. 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

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

ILL borrowing increases by 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 that the presence of the ILL has on other borrowing. The bottom left panel shows that when $\rho = 0.5$, ILL borrowing reduces other borrowing early in the life cycle, as one would expect, given that the ILL 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 because the uncertainty about future resources is now much smaller, so that young households want to borrow to consume from 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 borrowing is available, while equity holdings are lower for households that are approaching retirement, as they had less 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 discussed 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 interest 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 (because households needs to borrow more through the ILL to get the same degree of insurance) while at other times mean ILL borrowing decreases (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, these are again convex in the strength of the correlation of rates and permanent income shocks. However, the gains 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 Alternative Calibrations

In this subsection, we investigate further what is driving the differences in welfare gains between the two income-linked 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 shortcoming 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.[39] One possible explanation for the lack of demand for bonds is that our model ignores potential liquidity benefits to holding the risk-free asset, or that participation costs exist 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 an 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 the IHI and overstating the gains from the 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 equity, which will make the portfolios generated by our model look more like what is observed empirically. For our calibration, this assumption is equivalent to replacing equity by a 50/50 stock-bond 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 such as to match the mean wealth-to-income ratio before retirement, which yields $\beta = 0.947$. As compared with the

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

benchmark with no income-linked assets, the gain in CE consumption from having access to the baseline IHI with correlation –0.5 and volatility 0.5 is now 0.33 percent and thus, as expected, higher than if no bond holdings are required. In a related exercise that is more favorable to the IHI, we assume that households could invest in 50/50 stock-bond funds and in 50/50 stock-IHI funds (which assumes that households would view the IHI as a direct substitute for bonds in terms of liquidity advantages). The predicted welfare gain from having access to such a vehicle would be 0.71 percent. Meanwhile, the baseline ILL now produces a welfare gain of only 0.85 percent, instead of the 1.36 percent without required bond holdings. Thus, even though the ILL is still more attractive than the IHI, the differential welfare gain is now significantly smaller than when no bond holdings are required.

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 percent while for the baseline ILL with $E(\tilde{r}_{ILL}) = 0.05$ the gain is a very substantial 3.04 percent. 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 percent. 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 the 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 in reality most household borrowing other than through mortgages occurs at a rate significantly above the return on risk-free saving.

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, and also corresponds to the most commonly made assumption in macro models. However, the finance literature often assumes a 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 used 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 at any given wealth level. Hence, if we want to match the empirically observed wealth-to-income ratio (or the observed debt holdings; see Davis, Kubler, and Willen 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 percent, which is significantly larger than the 0.04 percent 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 percent, 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 percent 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. Furthermore, the welfare gains from the income-hedging instrument are 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 usually 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 2 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 baseline 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 pay 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 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 an expected return equal to the riskless rate. If the investor could borrow at the riskless rate, then he would perfectly hedge away his income risk, at no cost. 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.

A frequently heard comment is that if people were as concerned with labor income risk as in our model, and as rational and sophisticated as we assume them to be, they would hedge their income risk by shorting an equity index of the industry they work in, or even the their employer's stock. There are different reasons for why such a strategy is not as appealing as one might think. First, it is costly to take short positions, as one needs to post money in a margin account, and this has an opportunity cost to the investor (in particular if he needs to borrow). Second, and perhaps more surprisingly, the data on the correlation between industry stock returns and labor income shocks to workers in this same industry reveal that the correlation is often near zero and unstable over time (Davis, Fuster, and Willen 2010). Thus, even ignoring the costs of taking short positions, such a strategy might not provide a good hedge against labor income shocks.[40]

Another important issue is how the proposed income-linked instruments relate to currently existing social insurance mechanisms such as unemployment and disability insurance, or the possibility of declaring bankruptcy to clear one's debts after a negative shock. Given these mechanisms, are income-linked assets even needed? One answer is that the existing institutions provide insurance only temporarily (in the case of unemployment insurance) or against extreme negative shocks (disability insurance). Likewise, bankruptcy is very costly for defaulting households (in terms of legal costs as well as limited access to credit in the future) such that the option to default on one's debt after an adverse income shock usually will only be exercised after an extreme negative income shock, and due to its high costs declaring bankruptcy may provide only a moderate welfare gain. Income-linked assets, on the other hand, would facilitate insurance against less extreme shocks, potentially at a low cost.

In evaluating income-linked assets, using a model with realistic borrowing and investment

[40]The correlation between income shocks and stock returns may be more robustly positive at the firm level (it certainly is for most executives), but shorting the stock of one's own company may be undesirable for other reasons, such as the possibility of being accused of insider trading. Also, a significant fraction of the workforce does not work for publicly traded companies.

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 percent, and the mean return on the income-linked asset were also 2 percent and had a correlation of 0.5 with permanent income shocks, the predicted welfare gains would exceed 4 percent.

5 Conclusion

Income-linked assets such as the ones we consider in this paper have the potential to be useful for households' income risk management, but as we have emphasized, the devil is in the details. Folding the insurance against negative income shocks into a loan product makes it more useful to households than letting them purchase the insurance directly.[41] Furthermore, the correlation of the income-linked assets with households' permanent income shocks is a crucial determinant of the predicted size of the welfare gains.

This latter point highlights the importance of measurement issues, a point made also by Shiller (2003). To make good use of financial instruments such as the ones we have considered, one must be able to precisely 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 financial 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 individual household income. But the dramatic improvements in computing power and the wide availability of large disaggregated datasets mean that these challenges can be overcome.

[41] Incidentally, we believe that the intuition from our model likely carries over to insurance against house price risk: the welfare gains from folding such insurance directly into mortgage contracts (for instance, by reducing the mortgage principal if a — preferably local — house price index decreases) may be much larger than what households gain from having access to a housing derivatives market in which they can purchase such insurance separately.

Our model focuses on the risk management benefits that income-linked assets would provide, taking the riskiness of household income as given. Shiller (2003) discusses some additional benefits that these instruments might have which are not part of our analysis. First, the prices of these instruments (for instance, the borrowing rates on income-linked loans that the market offers to different professions) might aggregate and reveal information which would facilitate more effective decision making—for instance, when choosing which occupation to enter. Second, 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). Analyzing the potential benefits from these channels could be a fruitful avenue for future research.

References

Abowd, John M., and David Card. 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57(2): 411–445.

Attanasio, Orazio, and Steven J. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104(6): 1227–1262.

Bertaut, Carol C., and Michael Haliassos. 1997. "Precautionary Portfolio Behavior from a Life-Cycle Perspective." *Journal of Economic Dynamics and Control* 21(8-9): 1511–1542.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98(5): 1887–1921.

Blundell, Richard, and Ian Preston. 1998. "Consumption Inequality and Income Uncertainty." *Quarterly Journal of Economics* 113(2): 603–640.

Campbell, John Y., Joao F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout. 2001. "Investing Retirement Wealth: A Life-Cycle Model." In *Risk Aspects of Investment-Based Social Security Reform*, eds. John Y. Campbell and Martin Feldstein, 439–473. Chicago: University Of Chicago Press.

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

Carroll, Christopher D., and Andrew A. Samwick. 1997. "The Nature of Precautionary Wealth." *Journal of Monetary Economics* 40(1): 41–71.

Cocco, Joao F., and Francisco J. Gomes. 2009. "Longevity Risk, Retirement Savings, and Financial Innovation." Working Paper. Available at http://faculty.london.edu/fgomes/cg.pdf.

Cocco, Joao F., Francisco J. Gomes, and Pascal J. Maenhout. 2005. "Consumption and Portfolio Choice over the Life Cycle." *Review of Financial Studies* 18(2): 491–533.

Cochrane, John H. 1991. "A Simple Test of Consumption Insurance." *Journal of Political Economy* 99(5): 957–976.

Davis, Steven J., Andreas Fuster, and Paul S. Willen. 2010. "Occupation-Level Income Shocks and Asset Returns: Covariance Structure and Portfolio Choice Implications." Work in progress.

Davis, Steven J., Felix Kubler, and Paul Willen. 2006. "Borrowing Costs and the Demand for Equity over the Life Cycle." *The Review of Economics and Statistics* 88(2): 348–362.

De Jong, Frank, Joost Driessen, and Otto Van Hemert. 2008. "Hedging House Price Risk: Portfolio Choice With Housing Futures." Working Paper. Available at http://ssrn.com/paper=740364.

Feigenbaum, James, and Geng Li. 2009. "A Semiparametric Characterization of Income Uncertainty over the Lifecycle." Working Paper.

Gomes, Francisco, and Alexander Michaelides. 2005. "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence." *Journal of Finance* 60(2): 869–904.

Gomes, Francisco J., Laurence J. Kotlikoff, and Luis M. Viceira. 2007. "The Excess Burden of Government Indecision." Working Paper No. 12859. Cambridge, MA: National Bureau of Economic Research.

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

Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. "Consumption over the Life Cycle." *Econometrica* 70(1): 47–89.

Guvenen, Fatih. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review* 97(3): 687–712.

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

Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff. 1996. "Risk-Sharing Between and Within Families." *Econometrica* 64(2): 261–294.

He, Hua, and Neil D. Pearson. 1991. "Consumption and Portfolio Policies with Incomplete Markets and Short-Sales Constraints: the Finite-Dimensional Case." *Mathematical Finance* 1(3): 1–10.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008a. "Insurance and Opportunities: A Welfare Analysis of Labor Market Risk." *Journal of Monetary Economics* 55(3): 501–525.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2008b. "The Macroeconomic Implications of Rising Wage Inequality in the United States." Working Paper No. 14052. Cambridge, MA: National Bureau of Economic Research.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2009. "Quantitative Macroeconomics with Heterogeneous Households." *Annual Review of Economics* 1: 319–354.

Hryshko, Dmytro. 2009. "RIP to HIP: The Data Reject Heterogeneous Labor Income Profiles." Working Paper. Available at http://www.arts.ualberta.ca/ econweb/hryshko/.

Kaplan, Greg, and Giovanni L. Violante. Forthcoming. "How Much Consumption Insurance Beyond Self-Insurance?" *American Economic Journal: Macroeconomics* .

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

Laibson, David, Andrea Repetto, and Jeremy Tobacman. 2007. "Estimating Discount Functions with Consumption Choices over the Lifecycle." Working Paper No. 13314. Cambridge, MA: National Bureau of Economic Research.

MaCurdy, Thomas E. 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis." *Journal of Econometrics* 18(1): 83–114.

Pijoan-Mas, Josep. 2006. "Precautionary Savings or Working Longer Hours?" *Review of Economic Dynamics* 9(2): 326–352.

Shiller, Robert J. 1993. *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*. Oxford: Oxford University Press.

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

Shiller, Robert J., and Ryan Schneider. 1998. "Labor Income Indices Designed for Use in Contracts Promoting Income Risk Management." *Review of Income and Wealth* 44(2): 163.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. "Consumption and Risk Sharing over the Life Cycle." *Journal of Monetary Economics* 51(3): 609–633.

Willen, Paul, and Felix Kubler. 2006. "Collateralized Borrowing and Life-Cycle Portfolio Choice." Working Paper No. 12309. Cambridge, MA: National Bureau of Economic Research.

Zeldes, Stephen P. 1989. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics* 104(2): 275–298.

Zinman, Jonathan. 2009. "Where is the Missing Credit Card Debt? Clues and Implications." *Review of Income and Wealth* 55(2): 249–265.

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

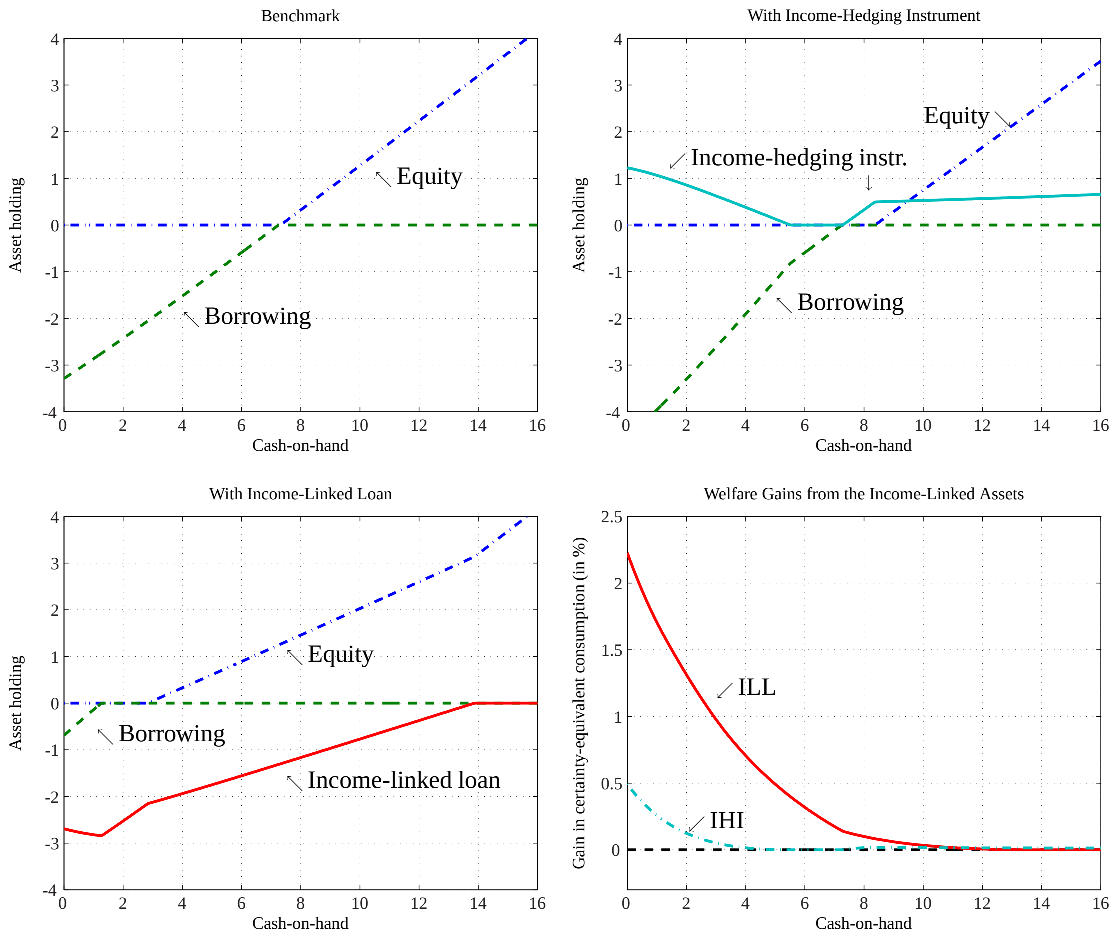

Notes: "Cash-on-hand" refers to the sum of wealth and income available to the investor in period 1. Income in period 2 can take the values $\{5.4, 8, 10.6\}$ with respective probabilities $\{1/6, 2/3, 1/6\}$. Source: Authors' calculations.

Figure 2: Income Process: Mean Profile and One Realization

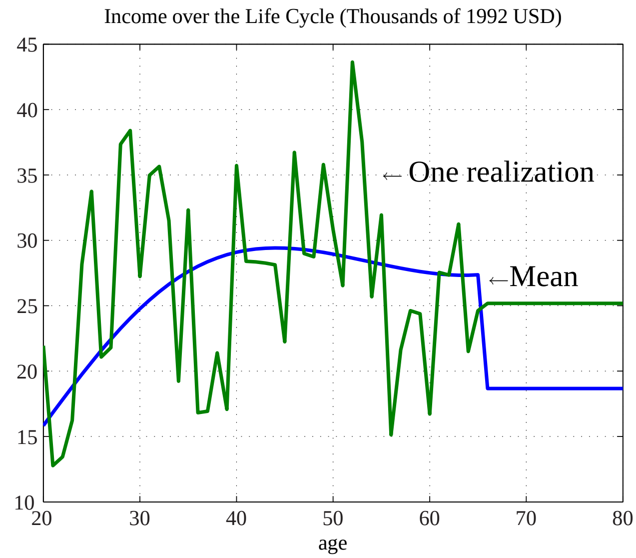

Source: Authors' calculations, based on income profiles for high school graduates estimated by Cocco, Gomes, and Maenhout (2005).

Figure 3: Life-Cycle Profiles in Benchmark Model

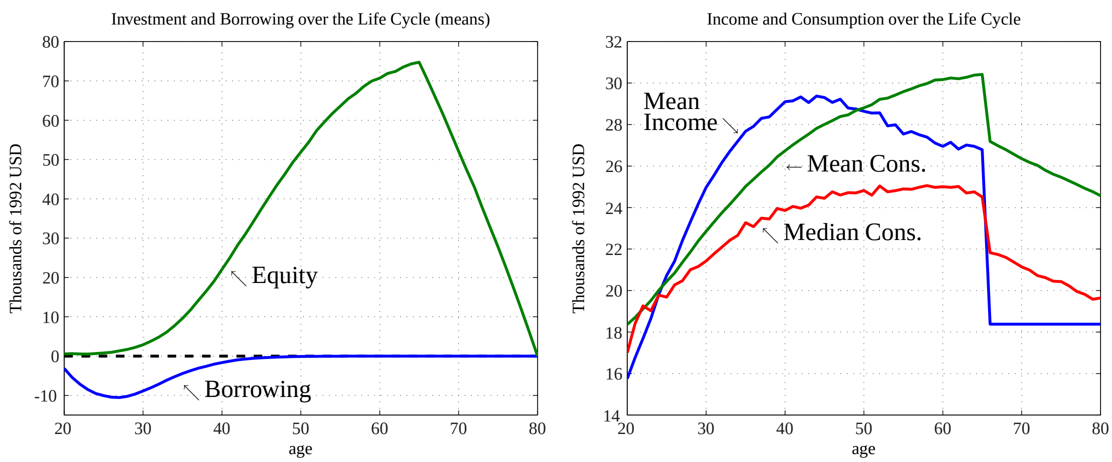

Source: Authors' calculations.

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

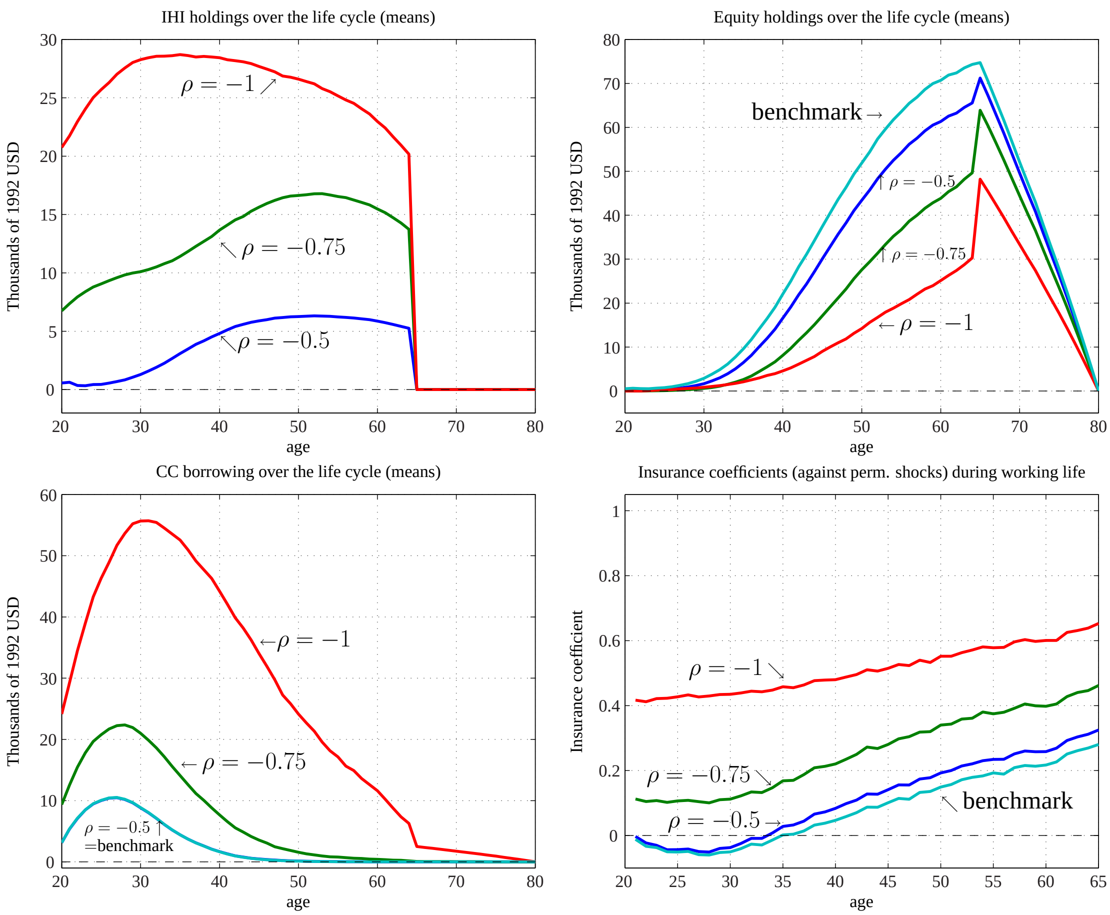

Notes: ρ denotes the correlation of the rate of return on the IHI with the permanent shock to labor income. All other parameters are as in the "Baseline" column of table 2. "Benchmark" refers to the case without IHI. Source: Authors' calculations.

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

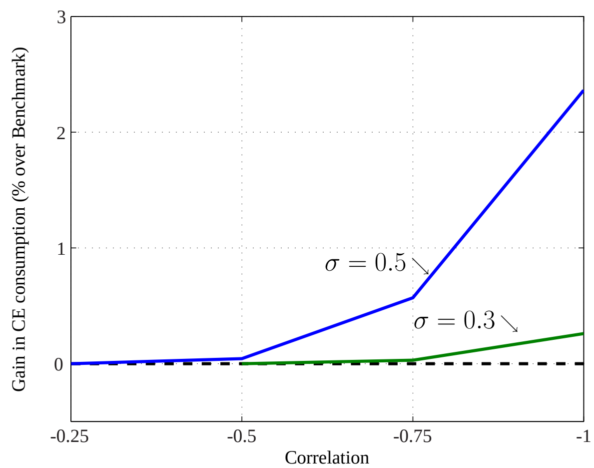

Notes: "Correlation" stands for the correlation of the rate of return on the IHI with the permanent shock to labor income, and σ denotes the volatility of the rate of return on the IHI. Source: Authors' calculations.

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

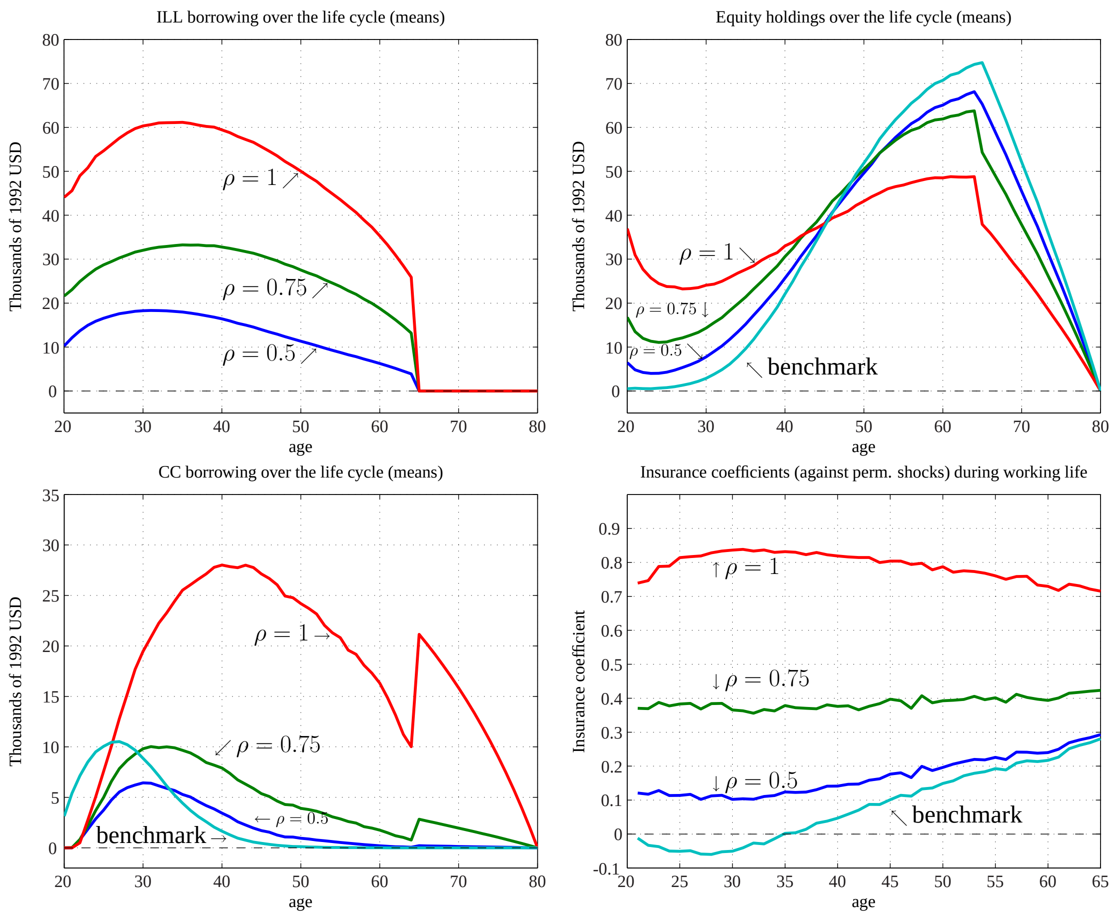

Notes: ρ denotes the correlation of the interest rate on the ILL with the permanent shock to labor income. All other parameters are as in the "Baseline" column of table 2. "Benchmark" refers to the case without ILL. Source: Authors' calculations.

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

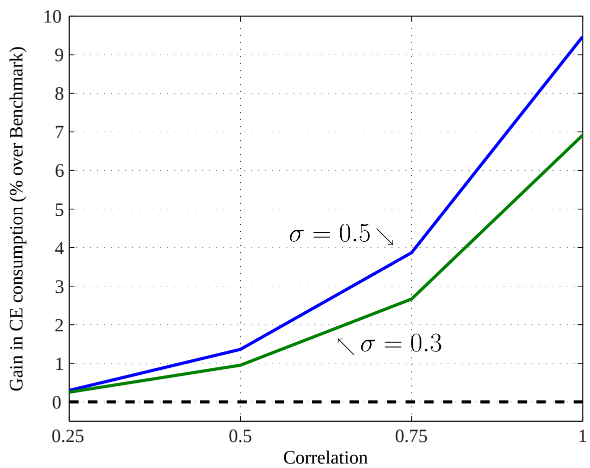

Notes: "Correlation" stands for the correlation of the interest rate on the ILL with the permanent shock to labor income, and σ denotes the volatility of the interest rate on the ILL. Source: Authors' calculations.

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

Gains in certainty-equivalent consumption, in percent, 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

Notes: ρ denotes the correlation of the return on the income-linked asset with income in period 2. σ is the standard deviation of this return.

Table 2: Parameter Values

Parameter	Baseline	Alternative values
Relative risk aversion γ	2	3
Discount factor β	0.936	0.92, 0.939, 0.947
Age of labor force entry	20	
Age of retirement	65	
Age of death	80	
Std. dev. of permanent shock σ_ν	0.10296	
Std. 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
Std. dev. of equity returns σ_e	0.16	0.08
Mean return on income-hedging instrument $E(\tilde{r}_{IHI})$	0.02	
Std. 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
Std. 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 positions (in thousands)			
Description	$E(\tilde{r}_{ILA})$	$\rho_{ILA,\tilde{\nu}}$	σ_{ILA}	$\overline{c}$	*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

Notes: “ILA” = income-linked asset, refers to the income-hedging instrument (IHI) for cases 1) to 7) and to the income-linked loan (ILL) for cases 8) to 15). $\rho_{ILA,\tilde{\nu}}$ = Correlation$(\tilde{r}_{ILA}, \tilde{\nu})$. $\overline{c}$ is certainty-equivalent consumption as defined in equation (4). Gains are assessed with respect to the benchmark case. $\overline{\phi}^{\nu}$ is the (unweighted) average of the insurance coefficient against permanent shocks (defined in equation (5)) over the working life. Asset positions are unweigthed means over the entire life cycle, with “Eq” referring to equity holdings and “CC” referring to unsecured borrowing (“credit cards”). Parameters other than the ones for income-linked assets are as in the “Baseline” column of table 2.

Table 4: Summary Table of Results from Alternative Calibrations

Preferences		*Assets*				*Welfare measures*		
γ	β	Stocks/Bonds	ILA	$E(\tilde{r}_{ILL})$	r_b	$\overline{c}$	*Gain (%)*	$\overline{\phi}^{\nu}$
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	IHI/Eq.	-	0.08	19653	0.71	0.20
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
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

Notes: In the “Stocks/Bonds” column, “free” refers to the case where no minimum investment in bonds is required, while “50/50” refers to the case where the household most hold at least as much money in bonds as in stocks. In the “ILA” (= income-linked asset) column, “IHI/Eq.” refers to the case where households have access to 50/50 stock-IHI funds.