TOWARDS UNDERSTANDING LIFE CYCLE SAVING OF BOUNDEDLY RATIONAL AGENTS: A MODEL WITH FEASIBILITY GOALS

By Johannes Binswanger

TILBURG ◆ UNIVERSITY

Towards Understanding Life Cycle Saving of Boundedly Rational Agents: A Model with Feasibility Goals*

Johannes Binswanger†

February 5, 2008

Abstract

This paper develops a new life cycle model that aims to describe the savings and asset allocation decisions of boundedly rational agents. The paper's main theoretical contribution is the provision of a simple, tractable and parsimonious framework within which agents make forward looking decisions in the absence of full contingent planning. Instead, agents pursue two simple so-called feasibility goals. The paper uses this framework to shed light on important empirical patterns of asset allocation that are puzzling from the point of view of existing models.

Key words: Behavioral economics, bounded rationality, equity shares, feasibility goals, life cycle saving, stock market participation.

JEL classification: D81, D91.

---

†Department of Economics, CentER and Netspar, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands; j.binswanger@uvt.nl.

1 Introduction

Within the last fifteen years it has become common to look at financial life cycle decisions through the lens of what may be dubbed the buffer-stock model of life cycle saving (Carroll, 1997).[1] This represents the modern form of the economic life cycle model. According to this model, individuals maximize the discounted sum of per-period expected utilities, while being exposed to non-insurable labor income shocks, and while facing borrowing constraints. This model has been partially successful in that its numerical predictions come relatively close to empirically observed consumption and savings profiles over the life cycle. The model has been much less successful when it comes to asset allocation. In this respect, the model is basically unable to explain the data in a reasonable way. In particular, the model predicts rates of stock ownership among the population that are far to high. Furthermore, it typically also predicts equity portfolio shares that are substantially higher than observed in the data. Importantly, this holds for a wide variety of circumstances, including the presence of fixed costs of stock market participation or the possibility of extreme labor income shocks, among others.[2]

Taking one step back, there is a more general reason why many economists are concerned about the modern life cycle model. As has been argued forcefully by Richard Thaler (1994, p. 187), "saving for retirement appears to be a domain where economists should be particularly worried about the issues raised by bounded rationality." This is likely to hold not only with respect to retirement saving but more generally with respect to life cycle saving and asset allocation. Subsequent studies have formally established that Richard Thaler has had a point. Allen and Carroll (2001) have explored whether boundedly rational agents could plausibly learn the life cycle model's predicted behavior, starting with trial and error. Their finding is staggering: Yes, but it would take about one million years! Intuitively, the reason is straightforward. There is little opportunity to learn how much to save and how to invest at the ages of 30, 40, 50 etc.

---

[1]Representative studies include, among others, Hubbard et al. (1995), Carroll (1997), Laibson et al. (1998), Cocco et al. (2005), Scholz et al. (2006)

[2]See Section 6 for a further discussion.

Moreover, positive or negative feedback from one's actions at particular ages lags substantially, which further slows down learning. Lettau and Uhlig (1999) demonstrate that it is even questionable whether consumers would ever learn the "right" rule. In other words, it is by no means guaranteed that learning would lead to convergence to the optimal solution of the model.

The main reason why decision making is so complex in the case of the modern life cycle model (henceforth called the standard model) is its requirement of full contingent planning. As mentioned above, the model assumes that agents maximize the discounted sum of expected per-period utilities from consumption. Suppose that an agent makes a savings or asset allocation decision today. In order to do so, she is required to anticipate her optimal decisions in any future contingency. Otherwise, the expected utility terms in the objective function would simply not be defined. In other words, it is not possible to decide on what is optimal today, independently of deciding on optimal actions in the future for all possible contingencies. In the presence of a realistic amount of uncertainty, decision making thus becomes extremely complex and full contingent planning may greatly exceed the capabilities of boundedly rational agents.[3]

This paper takes it as given that it is desirable to search for alternative models that offer a more plausible description of financial life cycle decision making by boundedly rational agents. The search for such models is also driven by the hope that they would be better able to explain observed asset allocation choices, while explaining consumption-saving choices at least equally well. The approach I pursue is to cut out the main source of complexity in the standard model, namely full contingent planning. The difficulty of developing such new models of bounded rationality stems from the fact that there is no existing framework to build on.[4] The provision of

---

[3]The importance of this has often been downplayed by referring to the argument that agents do not actually need to solve this complicated problem, but they may behave "as if" they knew how to solve it. Since Milton Friedman, it is common to refer to the example of a billiard player who may be very good at directing the movements of his balls, even if he is ignorant of the equations describing these movements. There is a key difference between life cycle saving and playing billiard, however. While the billiard player has the opportunity to train under identical circumstances as often as he likes, this possibility is absent in the case of life cycle saving. You are 30 years old just once. Moreover, if the as-if argument did hold, then we would expect it to be much better at explaining individuals' asset allocation choices.

[4]See, for instance, the following quote from Laibson et al. (1998, p. 101): "It is not clear how to [weaken assumptions about consumer sophistication] in a parsimonious and realistic fashion. (...) There are no well-

a new framework for modeling forward looking saving and investment behavior in the absence of full contingent planning is the main theoretical contribution of this paper.

In all other respects, not related to the issue of full contingent planning, I closely follow the standard approach. For instance, I do not dispense with the idea that behavior is generated by a kind of underlying preference structure and will hence not assume that behavior is driven by rules of thumb. I will also continue to assume that agents understand, at least intuitively, the arithmetic of intertemporal budget constraints.[5] The strategy I pursue is to depart from the standard model in steps, rather than switching the paradigm of the analysis in a radical way by assuming rule-of-thumb behavior.

The main idea of the new framework developed here is that individuals have some *feasibility* goals with respect to their future standards of living. Technically speaking, agents want to assure that certain choices lie within their future budget sets, and are hence feasible. In the standard model, agents are concerned about their *actual* future standards of living in different states of nature. Instead, in the feasibility goals (FG) framework, agents care only about the *feasibility* of certain future standards of livings. They never decide on their actual choices in advance. It is exactly this feature of the framework that eliminates the need for full contingent planning and substantially simplifies decision making.

When calibrating the model I find that it predicts consumption profiles that closely resemble the respective profiles predicted by the standard model. However, when it comes to asset allocation choices, the model is much better able to explain the data. Specifically, the model can explain why the young typically stay out of the stock market. Second, the model is consistent with relatively low equity shares for those who do participate in the market. Third, the model is consistent with the fact that low-income earners tend not to enter the stock market. Finally, conditional on stock market participation, equity shares increase with permanent income.

The paper is organized as follows. Section 2 briefly describes the new model developed in

developed bounded rationality models applicable to the problem of life-cycle saving."

[5]It would be straightforward to relax the latter assumption in future work.

this paper in an informal way. Section 3 discusses the relationship to the existing literature. Section 4 introduces the model formally. Second 5 discusses its calibration. Section 6 briefly discusses some stylized facts about how consumption, savings and asset allocation vary over the life cycle. It contrasts this facts with the predictions of existing models. Section 7 presents calibration results for a baseline specification as well as for a host of alternative specifications. Section 8 concludes. The computational algorithms that are used to solve the model are outlined in an appendix.

2   A Brief Informal Description of the Model

The model developed in this paper assumes two simple feasibility goals. These goals are specified in a way such that they give rise to both a desire for safety and the desire to enjoy high *expected* future standards of living. As a consequence, agents face simple risk-return trade-offs when making their savings and asset allocation decisions and they engage in precautionary saving.

The first goal is dubbed the insurance goal. It entails assuring that the minimally feasible levels of future standards of living are never below some fraction of current consumption. In other words, agents do not want to forgo the possibility of consuming at least some given fraction of current consumption in the future. This fraction may decline with the distance between the current and a particular future period. The insurance goal may be seen as reflecting habit formation or loss aversion. Alternatively, it may just be seen as a simple precautionary planning device for boundedly rational individuals that exhibit satisficing behavior. The insurance goal triggers a precautionary savings motive.

The insurance goal implicitly relates to a future worst-case scenario. Clearly, the latter should be understood in practical terms rather than literally. It reflects a scenario for which the probability that things turn our even worse than for this particular scenario is very low. In the model developed in this paper the worst-case scenario refers to low financial and non-financial

income levels in every future period of life.

The second feasibility goal refers to the feasibility of a certain (higher) standard of living in the future, provided that a "normal" or average scenario materializes. This goal is pursued by accumulating wealth and is therefore dubbed the accumulation goal. The accumulation goal captures the desire of enjoying an increasing standard of living, of "becoming rich," or the desire to save in order "to enjoy a sense of independence and the power to do things, though without a clear idea or definite intention of a specific action" (Keynes, 1936). A normal scenario is defined as the branch of the event tree where all random variables take on their expected values, conditional on current information.

Agents are assumed to apply an algorithm that determines how trade-offs are made between achieving the goal of a high level of current consumption, the insurance and the accumulation goal.

3   Relationship to Existing Literature

Existing behavioral life cycle models include the mental accounting model of Shefrin and Thaler (1988), models of hyperbolic discounting (Laibson, 1997), and the loss aversion model of Bowman et al. (1999). The issue of full contingent planning is not addressed in this literature. Gabaix and Laibson (2000, 2005) provide models of short-cuts that individuals may use when working through a decision tree, such as removing branches with low probability. An important difference between the FG model and the models of Gabaix and Laibson is that the latter deal with decisions where a choice is made only once at the beginning of a probabilistic event three. In contrast, there is generally a multitude of subsequent choices to be made in the FG model.

An early attempt to find an alternative to full contingent planning models has been made by Pemberton (1993). In his model agents maximize discounted expected utility of current consumption and so-called future sustainable consumption levels. The sustainable consumption levels refer to a flat consumption path that, in expectation, would just exhaust total life time

resources. It is thus equal to permanent income in its standard meaning. The main difference between the model of Pemberton and the one developed here is that his model is grounded in the discounted expected utility model. In contrast, the FG model more generically represents a model of bounded rationality where thinking about the future is confined to a few scenarios and decision making is determined by the desire to achieve certain goals that serve as reference points. While Pemberton does not provide any analysis of asset allocation choices, his model seems less promising for explaining such choices. It would be likely to face the same difficulties than standard expected utility models.

Binswanger (2006) develops a simple goal-based two-period life cycle model which is dubbed threshold goal model. This model coincides with the FG model in the special case where there are only two periods, such that decisions are only made once in the initial period. The two-period setup is basically static in nature and the issue of full contingent planning does not arise. In Binswanger (2006) it is demonstrated that the static goal model explains the *cross-section* of savings and asset allocation choices very well in comparison to other existing models. In contrast, the issue of this paper is to provide a model that explains how savings and asset allocation vary *over the life cycle* within a truly dynamic setting.

4   The Model

For simplicity, the basic decision making unit of the model is a household with a constant number of members.[6] I will refer to this household as the "agent." Let t denote the current period or, equivalently, current age. Over the agent's life t runs from T_0, where she starts to work, to age T, after which death occurs with certainty. The analysis abstracts from lifetime uncertainty, since this would raise additional issues that are unrelated to the main topic of the paper.

At each age t the agent makes three choices. She decides how much to consume, how much

---

[6]See Section 7.3 for a specification of the model where the size of the household changes over time.

to invest in a risk-free bond, and how much to invest in risky stocks. This is a prototypical decision problem that has been analyzed in the literature[7] and captures the nature of financial life cycle decisions in a stylized way. An important omission is housing. The inclusion of housing would not provide any conceptual difficulties, but it would render the model less parsimonious and less comparable to the existing literature. It is thus natural to omit housing in a first step.

Denote current consumption by c_t, current bond investments by b_t, and current stock investments by s_t. The crucial feature of the model is that at age t the decision maker does not know what her future decisions will be at different knots of the event tree. However, she is concerned about what consumption levels would be *feasible* in the future. Denote by $c_{t|t+i}$ the (random) level of consumption that is feasible at a future age $t + i$, $i = 1, 2, \ldots, T - t$, from the perspective of age t, given the current level of wealth and given current information about future income streams. Similarly, denote by $b_{t|t+i}$, $s_{t|t+i}$ the levels of bond and stock investment that are feasible at time $t + i$ from the perspective of age t. Concerning the subscripts, I apply the convention that $t|t + i \equiv t - |i|$ for $i \leq 0$. Thus, $t|t = t$, for instance. In accordance with the relevant literature it is assumed that

$$b_{t|t+i} \geq 0, \ s_{t|t+i} \geq 0 \tag{1}$$

for all t and $0 \leq i \leq T - t$. This means that borrowing or short-selling of stocks is excluded. Clearly, future feasible choses depend on scenarios or states of the world that will materialize. For simplicity, the notation does not make this dependence explicit.

Denote the constant risk-free real per-period return of bonds by r. Furthermore, denote risky real per-period stock returns by r_t^s. It is assumed that stock returns are distributed identically and independently (*iid*) over time. Thus, only the realizations of r_t^s depend on time, while the time subscript for the random variable can be dropped. It is assumed that $Er^s > r$ and $r^{s\,\min} < r$, where $r^{s\,\min} \equiv \min r^s$. Thus, expected stocks returns are higher than bond returns,

---

[7]See e.g. Cocco et al. (2005) and Gomes and Michaelides (2005).

but minimum stock returns are lower.

At each age t the agent earns an exogenous and possibly risky stream of non-financial income Y_{t+i}, $0 \leq i \leq T - t$. This income may be composed of transitory and permanent shock components and, in expectation, follow a typical hump-shaped curve over the life cycle. Thus, the process for Y_t, Y_{t+i} is generally not *iid*. It will be specified more fully in Section 5. Denote future non financial income conditional on information at age t by $Y_{t|t+i}$. The intertemporal budget constraint is then given by

$$c_{t|t+i} + b_{t|t+i} + s_{t|t+i} = Y_{t|t+i} + b_{t|t+i-1}r + s_{t|t+i-1}r_t^s \tag{2}$$

for all t and $0 \leq i \leq T - t$. It is assumed that the agent has unbiased knowledge of the distribution of future income streams. Furthermore, she is assumed to understand (at least intuitively) the arithmetics of budget constraints. It would be straightforward to relax these assumptions, but this is beyond scope of this paper.

I turn now to the definition of the insurance and the accumulation goal which, in turn, are used for the statement of the FG algorithm. This algorithm plays the role of *preferences* in traditional analyses of intertemporal decision making. It proves convenient to introduce a more compact notation. Let $\mathbf{z}_{t|t+i} := \left(c_{t|t+i}, b_{t|t+i}, s_{t|t+i} \right)$, $0 \leq i \leq T - t$ denote a choice that is feasible at age $t + i$, and denote the set of all feasible choices at age $t + i$ by

$$\mathbf{B}_{t|t+i} \left(\mathbf{z}_{t|t+i-1}, Y_{t|t+i} \right) := \left\{ \mathbf{z}_{t|t+i} \in \Re_+^3 : \ c_{t|t+i} + b_{t|t+i} + s_{t|t+i} \leq Y_{t|t+i} + b_{t|t+i-1}r + s_{t|t+i-1}r_t^s \right\}.$$

The insurance goal is then defined as follows.

Definition 1 *A current choice* (c_t, b_t, s_t) *satisfies the insurance goal if, for a given sequence of nonnegative numbers* $\{\alpha_i\}_{i=0}^{T-t}$ *(where* $\alpha_0 \equiv 1$*), there exist* $b_{t|t+i}^*$*,* $s_{t|t+i}^*$*,* $b_{t|t+i+1}^*$*,* $s_{t|t+i+1}^*$ *(where* $b_{t|T}^* \equiv 0$*,* $s_{t|T}^* \equiv 0$*), such that*

$$\left(\alpha_i c_t, b_{t|t+i}^*, s_{t|t+i}^* \right) \in \mathbf{B}_{t|t+i}$$

8

and

$$\left(\alpha_{i+1}c_t, b^*_{t|t+i+1}, s^*_{t|t+i+1}\right) \in \mathbf{B}_{t|t+i+1}$$

for all i with $0 \leq i \leq T - t - 1$ and all values of r^s and Y_{t+i} that are taken on with positive probability, given information available in t. (The latter two random variables are implicit in the definition of $\mathbf{B}_{t|t+i}$).

Achieving the insurance goal means that a specific fraction α_i of current consumption remains feasible in the future under all circumstances. This includes the worst-case scenario which corresponds to a sequence of lowest possible values for r^s and $Y_{t|t+i}$. The fraction α_i may depend on the distance between the current period t and a particular future period $t + i$. By definition, $\alpha_0 \equiv 1$. In principle, the sequence $\{\alpha_i\}$ is unrestricted and could even be non-monotonic (except that the elements are to be nonnegative). However, when it comes to a baseline specification for simulating the model, it is desirable to severely restrict the flexibility of the model and thus its degrees of freedom. For the baseline specification it will thus be assumed that $\alpha_i = (\alpha^*)^i$, $0 < \alpha^* < 1$, which just offers one free parameter. This specification includes an element of "discounting."[8]

Similar to models of habit formation or loss aversion, the insurance goal refers to a reference point that is proportional to c_t. This represents a particularly simple specification. It assures that choices fulfilling $c_{t|t+i} \geq \alpha_i c_t$ in all future contingencies are always feasible if financial markets provide an asset with a strictly positive minimum return in all possible future states of the world.[9]

An example may illustrate how knowledge of the budget constraints is sufficient to determine whether a choice fulfills the insurance goal or not. Consider the case of three periods, where an agent earns w in the first period and earns no non-financial income in the other two periods. Suppose, for simplicity, that the agent can only invest in bonds, not in stocks.

---

[8]See Section 5 fore a more detailed discussion.

[9]While, in principle, the minimum return of any financial asset is zero since even the U.S. government could go bankrupt, it is common and practical to assume that there exists a risk-free interest rate.

The present value of future "habit consumption" is then $\frac{\alpha_1 c_1}{r} + \frac{\alpha_2 c_1}{r^2}$. Thus, any c_1 for which $w - c_1 < \frac{\alpha_1 c_1}{r} + \frac{\alpha_2 c_1}{r^2}$ will not fulfill the insurance goal. Thus, the insurance goal imposes a lower bound on current consumption. In the general case where an agent can also invest in stocks, the insurance goal also imposes a lower bound on the proportion of savings invested in stocks, as stock returns have a higher downside risk than bonds. It follows that it is possible to decide whether a choice fulfills the insurance goal or not in the absence of any knowledge about future actual consumption levels, i.e. in the absence of full contingent planning.

The next step is to define the accumulation goal. The accumulation goal represents a target standard of living that an agent wants to be feasible from some future age T^* on in case that a normal scenario evolves.[10] This normal scenario is defined as the branch of the decision tree where every random variable takes on its expected value, given information at age t.

The standard of living that will be feasible from some future age $T^* > t$ depends not only on choices made at age t, but also on savings and asset allocation choices between t and T^*. Thus, in the absence of knowledge of future actual decisions, the only way to determine what standard of living will be feasible from age T^* on is to go through a hypothetical accumulation scheme that is feasible from time t on and, in turn, leads to feasibility of a particular standard of living from age T^* on.

This hypothetical accumulation scheme works as follows. Start with a particular level of current consumption c_t. This fixes the level of age-t savings. Identify now all portfolios, corresponding to different equity shares, for which the insurance goal is met. Among these portfolios, identify the one with the highest expected return and thus the highest equity share. Choose the latter portfolio for the determination of the hypothetical accumulation scheme. The accumulation goal captures thus a "speculative" motive of savings.

Making the transition to the following period, assume that stock returns as well as non-financial income will be drawn according to the normal scenario. Assuming a hypothetical

---

[10]An agnostic specification that will be used when calibrating the model is $T^* = \frac{T-t}{2}$. This should be understood as the expected value of a random variable that is uniformly distributed over the remaining ages. See Section 7.3 for a different specification.

level of consumption kept fixed at c_t for the following periods, this accumulation scheme can be iterated until age T^*. The accumulation goal is then defined as the hypothetically accumulated balance at age T^*, divided by the remaining number of lifetime periods, to convert it in per-period units. This approximates a standard of living that is feasible from age T^* on under the accumulation scheme just discussed, provided that the normal scenario will materialize. Since the level of the accumulation goal depends on the initial choice of c_t, the value of the goal should be understood as a function of c_t. The FG algorithm discussed below will determine how the agent makes a trade-off between the conflicting goals of achieving a high c_t and a high level of the accumulation goal.

More formally, given a choice $c_{t|t+i}$, denote $\mathbf{I}_{t|t+i}\left(c_{t|t+i}\right)$ the set of all feasible portfolios $\left(b_{t|t+i}, s_{t|t+i}\right)$ such that the insurance goal is satisfied for $\left(c_{t|t+i}, b_{t|t+i}, s_{t|t+i}\right) \in \mathbf{B}_{t|t+i}$. The set $\mathbf{I}_{t|t+i}$ may be empty for high values of $c_{t|t+i}$. In this case, $\alpha_j c_{t|t+i}$, $1 \leq j \leq T - t - i$, may not be reached in the future even if all savings are invested in bonds. Given $c_{t|t+i}$ and, implicitly, a level of accumulated wealth of $A_{t|t+i} := Y_{t|t+i} + b_{t|t+i-1}r + s_{t|t+i-1}r^s_{t+i}$ which determines $\mathbf{B}_{t|t+i}$, define $\left(b'_{t|t+i}, s'_{t|t+i}\right)$ as the element of $\mathbf{I}_{t|t+i}\left(c_{t|t+i}\right)$ with the maximal value for $s_{t|t+i}$. This represents the portfolio with the highest expected return for which the insurance goal is still met. If $\mathbf{I}_{t|t+i}$ is nonempty such a portfolio exists. It is assumed that, given $c_{t|t+i}$, the agent considers to choose $\left(b'_{t|t+i}, s'_{t|t+i}\right)$ in order to "implement" (hypothetically) the accumulation goal at age $t + i$. (It will become clear from the FG algorithm discussed below that $\mathbf{I}_{t|t+i}$ will always be nonempty along the "optimal" trajectory.) Furthermore, the agent projects a choice of $c_{t|t+i} = c_t$. Wealth is accumulated (hypothetically) according to

$$A_{t|t+i+1} = E_t Y_{t+i+1} + b'_{t|t+i}r + s'_{t|t+i}Er^s, \tag{3}$$

which corresponds to the normal scenario. It is now possible to give a precise definition of the accumulation goal.

Definition 2 *Take $c_t \in \mathbf{B}_t$ as given. If all sets $\mathbf{I}_{t|t+i}\left(c_t\right)$, $0 \leq i \leq T^* - t - 1$, are nonempty,*

then the value of the accumulation goal, denoted by $a_{t|T^*}$, is equal to $A_{t|T^*}/(T - T^*)$, where $A_{t|T^*}$ is defined by equation (3). Otherwise $a_{t|T^*} = 0$.

Note that, according to this definition, $a_{t|T^*}$ is well defined for all values of $c_t \in \mathbf{B}_t$. It is now possible to give a formal definition of the FG algorithm.

Definition 3 *The feasibility goals (FG) algorithm is implemented by working through the following steps.*

Step 1. Identify the element of $\mathbf{B}_t$ with the highest level of c_t such that $\mathbf{I}_t(c_t)$ is nonempty. Denote this element by $\hat{\mathbf{i}}_t$ and the corresponding value of c_t by $\hat{c}_t$.

Step 2. If $\hat{c}_t \leq \bar{c}$ (where $\bar{c} \geq 0$), then implement $\hat{\mathbf{i}}_t$. Otherwise go to step 3.

Step 3. Choose the element of $\mathbf{B}_t$ which maximizes

$$\ln(c_t - \bar{c}) + \beta^{T^* - t} \ln\left(a_{t|T^*}\right), \tag{4}$$

where $0 < \beta < 1$.

The first step just means maximizing c_t, subject to meeting the insurance goal. Necessarily, step 1 leads to a choice for which stock investments are zero. The reason is that minimum stock returns are lower than bond returns. If stock holdings were positive and the insurance goal is fulfilled, stocks could be substituted by bonds and, at the same time, consumption could be increased, such that the insurance goal would still be fulfilled. As a result, the original choice with positive stock holdings was not optimal.

The parameter $\bar{c}$, which appears in step 2 and is assumed to be greater or equal to zero, represents a "normal standard of living." Think of this parameter as the real dollar value of the expenditures on an "average" consumption bundle. The parameter $\bar{c}$ is treated as exogenous. It acts as a reference point or goal value for current consumption if budgets are relatively low. To see this, note that for $\bar{c} > 0$ we will always have $\hat{c}_t < \bar{c}$ if the budget is sufficiently small. In this case step 1 implies that the decision maker has two goals. The first is to achieve the

12

insurance goal and the second, conditional on the first, to bring c_t as close as possible to $\bar{c}$. It is only when this latter goal is achieved that the accumulation goal is triggered in step 3. From a mathematical point of view, $\bar{c} > 0$ has the effect of separating out a set of small budgets for which the accumulation goal is not activated. The latter thus has the status of a "luxury" goal.

Clearly, the most parsimonious specification of the model is the one where $\bar{c} = 0$.[11] In this case the model would not be able to explain that equity shares or wealth accumulation increase with permanent income, which is observed in the data.[12] Beyond this, I believe that it is of natural interest to consider the effect that such a reference standard of living has on savings and asset accumulation choices.

The objective function (4) in step 3 specifies the way in which the decision maker is assumed to solve the trade-off between her goal of enjoying a high level of current consumption and of the accumulation goal, in case that the latter is activated. (Note that the desire to achieve the insurance goal is implicit in the definition of the accumulation goal.) The parameter β measures the preference weight the decision maker assigns to the accumulation goal relative to present consumption. The weight depends on the distance between current age and the future date T^*, which is the age the accumulation goal refers to. The parameter β should not be understood as a genuine discount factor, as it does not specify overall time preferences. Rather, it refers more narrowly to the accumulation goal and captures thus the strength of the desire to enjoy a high standard of living in the future.[13] Note that (4) does not include any element of risk aversion. The only source of risk aversion in the FG model comes from the desire to meet the insurance goal.

As has already been mentioned, the FG algorithm plays the role of preferences in standard analyses. The decision maker is assumed to follow the FG algorithm at each age t anew,

---

[11]See Section 7.3.

[12]See Sections 6 and 7.

[13]Note that the sequence $\{\alpha_i\}_{i=1}^{T-t}$ includes some element of discounting if it is falling over time, such that the decision maker is less concerned about the downside risk of her standard of living in the far-distant future than in the near future. See Section 5.

irrespective of her earlier choices and plans. Thus, decision making can be seen as time-inconsistent. However, this inconsistency arises from incomplete planning rather than from a convex pattern of discounting. At each age t the decision maker focuses only on a worst-case scenario, by means of the insurance goal, and on a "normal" scenario, by means of the accumulation goal. Reducing the complexity of an event tree to a few representative scenarios is a plausible feature of decision strategies for boundedly rational agents (Gabaix and Laibson, 2000, 2005).

It is of interest to briefly discuss whether the outcome of the FG algorithm is well defined. Note first the budget sets $\mathbf{B}_{t|t+i}$ are compact and convex. Furthermore, it is quite straightforward to show that the sets $\mathbf{I}_{t|t+i}$ are compact and convex if they are nonempty. It can be shown that $r > r^{s\min}$ implies that there is a unique element of $\mathbf{B}_t$ for which c_t is maximal under the constraint that $\mathbf{I}_t(c_t)$ is to be nonempty.

Step 3 is somewhat more tricky. Note first that (4) is equal to minus infinity if $a_{t|T^*} = 0$. Thus, the agent will never choose a level of c_t for which one or several of the sets $\mathbf{I}_{t|t+i}$ are empty (see Definition 2). In the absence of the borrowing constraints in (1), $a_{t|T^*}$ would simply be a linear function of c_t in the relevant domain where the insurance goal is feasible. However, its shape is more complex in the presence of borrowing constraints. If c_t is sufficiently low, then the insurance goal would be feasible even with an equity share larger than one. However, an equity share larger than one is not feasible in the presence of borrowing constraints. As a result, $s'_{t|T^*}$ becomes a piecewise linear function of c_t. Similarly, $b'_{t|T^*}$ is a piecewise linear function of c_t, since bonds are zero for low levels of c_t and increase linearly at higher levels of c_t. It can be shown that $b'_{t|T^*}r + s'_{t|T^*}E\tilde{y}$ is a monotonically decreasing, piecewise linear and concave function of c_t in the relevant domain where the insurance goal is feasible (such that the sets $\mathbf{I}_{t|t+i}$ are nonempty). An induction argument can then be used to infer that the accumulation goal $a_{t|T^*}$ is a monotonically decreasing, piecewise linear and concave function of c_t. Since the composition of a concave function and an increasing and concave function is concave, $\ln\left(a_{t|T^*}\right)$ is a concave function of c_t. This finally implies that the objective function (4) is concave. It is

even strictly so, which follows from the piecewise linearity of $a_{t|T^*}$ and the strict concavity of the logarithmic function. As a result, there is a unique level of c_t that maximizes (4).

5 Calibrating the Model

Although the FG model could be solved analytically for some stylized cases, it is more interesting to study its quantitative predictions, in order to compare them to empirical estimates and to simulation results obtained in the literature for conventional preference models. It should be mentioned here that I will only report simulation results for choices made during working life. Since the FG model as presented here does not take into account mortality risk nor bequest motives, it would not be well suited for explaining savings and asset allocation choices during retirement. As a result, their analysis is beyond scope of this paper.

Calibrating the model requires a specification of the parameters of the FG algorithm, of the process for non-financial income, and of bond and stock returns. They shall be discussed in turn. The computational solution of the model is sketched in an appendix.

Concerning the time setup of the model, an individual is assumed to start working at age T_0, which is set to 21, to retire at the beginning of age R, set to 66, and to die with certainty at the beginning of age T, set to 86. The model is simulated at an annual frequency.

5.1 The parameters of the FG algorithm

Applying the FG algorithm requires specification of the sequence $\{\alpha_i\}$, of β, $\bar{c}$, and T^*. Starting with the sequence $\{\alpha_i\}$, I choose $\alpha_i = (\alpha^*)^i$ as a baseline specification, where $0 < \alpha^* < 1$. There are two reasons for this specification. First, it is very parsimonious and requires specification of only one parameter. Second, it captures in a very simple way the idea that decision makers may be less concerned, or may think less carefully, about negative events in the far-distant future than in the near future. This is most likely to happen on an intuitive or semi-conscious level. It is rather unlikely that many individuals literally engage in careful financial

life cycle planning, at least at early ages. In this sense, the specification of the sequence $\{\alpha_i\}$ captures thus a particular element of discounting.

The baseline value for α^* is set to .96, which is chosen as follows. As mentioned in the introduction, Binswanger (2006) provides a cross-sectional analysis of savings and asset allocation in a static model, corresponding to the FG model in the degenerate case where there are only two periods and it is only in the first period where agents make any decisions. By construction, such a setup does not allow for any variation of consumption, savings or asset allocation over age. In Binswanger (2006) this model is calibrated in a stylized way, assuming that the length of the planning horizon amounts to 30 years. Fitting the model to *cross-sectional* data on savings and asset allocation choices, the value obtained for the parameter corresponding to α_i at a time horizon of 30 years is .27.[14] In order to impose discipline and to limit flexibility, baseline parameter values in this paper are set such that they correspond to the parameters that best fit cross-sectional data in the case of the static model. I choose thus a baseline value of $\alpha^* = 0.96$ such that $(\alpha^*)^{30} \approx 0.27$.

Turning to $\bar{c}$, the best-fit value found in Binswanger (2006) amounts to 27,000 year-2001 U.S. dollars. Section 7.3 also provides simulations for the more parsimonious case where $\bar{c} = 0$. Concerning β, the best-fit annualized value for the cross-section case is .94. While this value leads to modest equity shares in the static setup, it leads to relatively high equity shares in the dynamic model of this paper (see Section 7.3). By construction the static model with a planning horizon of 30 years restricts people to hold the same portfolio for 30 years. In contrast, the dynamic model with period lengths of one year allows agents to rebalance their portfolio every year. It is well known from the literature that increases in flexibility increase the willingness to take risk in standard expected utility models (see e.g. Gollier, 2001). The same holds for the FG model. Thus, a value for β of .94 should be scaled down in order to obtain a suitable counterpart value for the dynamic model. I thus set the baseline value for β to .90. Remember that β should

---

[14]This number should not be interpreted as implying that the consumption level that individuals would like to have at all costs during retirement would be only .27 of their working life consumption. See Binswanger (2006).

not be understood as a genuine discount factor but more narrowly as the annualized preference weight of the accumulation goal. Thus, a value of .90 does not seem particularly low.

The baseline value for the parameter T^*, representing the age that the accumulation goal refers to, is set to $(t + T)/2$, rounded downwards. This can be understood as the mean of a uniform distribution over all remaining life ages from time t on. It represents thus an agnostic specification. (See Section 7.3 for an alternative specification.)

5.2 Non-Financial Income

Non-financial income is thought to include mainly labor income, but also unemployment benefits or income from Social Security etc. Following Cocco et al. (2005), I assume that

$$Y_t = F_t P_t V_t, \tag{5}$$

where

$$P_t = P_{t-1} U_t, \tag{6}$$

during working life. F_t represents a non-stochastic component of non-financial income that determines the general shape of this income over the life cycle in the absence of shocks (see Figure 1). V_t and U_t are mutually independent iid random variables with a mean of one, where $V_t \in \left[\underline{V}, \overline{V}\right]$ and $U_t \in \left[\underline{U}, \overline{U}\right]$. U_t represents a shock component which has a permanent effect on non-financial income. V_t represents a transitory shock component.

Cocco et al. (2005) report estimations for (5), (6), as well as for the standard deviations of the logs of V_t and U_t.[15] The calibrations here are based on their estimations for high school graduates. Starting with the assumption that U_t and V_t are distributed lognormally, I truncate the distributions of U_t, V_t by setting $\underline{U}$ and $\underline{V}$ to the first percentile of their original lognormal distribution, while keeping expected values at a value of one.[16] The resulting values for $\underline{U}$ and

---

[15]See their Tables 1, 2, 3.

[16]Note that the FG algorithm requires only specification of the minimum and expected values of all random

17

$\underline{V}$ are .97 and .81, corresponding to a decline in non-financial income of 3 and 19 percent, respectively.

During retirement, it is assumed that non-financial income is given by

$$Y_t = \rho Y_{R-1} M_t, \tag{7}$$

The parameter ρ represents a "replacement rate "and is set to .6. $M_t \in \left[\underline{M}, \overline{M}\right]$ represents a medical expenditure shock. The idea is that the medical expenditure shock reduces the amount of income that is available for normal spending. The baseline value of $\underline{M}$ is set to .6. For simplicity, it is assumed that this shock is realized at age R and is fully permanent. This means that $M_{R+i} = M_R$ for $1 \le i \le T - R$. Furthermore, it is assumed that $E_t M_R = 1$ for $T_0 \le t \le R - 1$. Clearly, this specification is very stylized. The advantage of this is that it keeps the simulations transparent. Since the paper focuses on savings and asset allocation decisions during working life, this stylized specification should not be seen as problematic.

All simulations in Section 7 are run for the case where, ex post, no shocks occur. This, again, keeps the simulation results particularly transparent and easy to interpret.

5.3 Bond and Stock Returns

Bond returns are assumed to be risk free and to amount to 2 percent per year. Furthermore, I set annual expected stock returns to 6 percent and their standard deviation to its historical value of .157.[17] Stock returns are assumed to be *iid* over time. For a simulation of the FG algorithm it is necessary to specify minimum stock returns. I do so starting with the assumption that stock returns are lognormally distributed with a mean and standard deviation as mentioned above. The distribution is then truncated by setting minimum stock returns to the first percentile of

variables involved.

[17] These assumptions are fairly common. See e.g. Campbell and Viceira (2002).

the original lognormal distribution, while keeping expected returns at 6 percent.[18] The gross minimum return amounts then to .74 per year, which corresponds to a net return of minus 26 percent. Again, all simulations in Section 7 are run for the case where, ex post, stock returns always take on their expected values.

Table 1 summarizes the parameter values for the baseline specification.

6 The "Facts" about Saving and Asset Allocation and the Predictions of Existing Models

One major aim of this paper is to explore whether the FG model provides a good explanation of individuals' observed consumption, savings and asset allocation choices during their working life. Before turning to the simulation results it seems therefore useful to discuss some of the empirical "facts" regarding the variation of consumption/savings and asset allocation choices over the life cycle. These are contrasted with the predictions of existing life cycle models.

6.1 The Empirical "Facts"

Researchers estimating age profiles for consumption/savings and asset allocation choices face a fundamental identification problem. Both choices are likely to be influenced not only by age, but also by the fact that a person is member of a certain cohort, as well as by the state of the economy at a particular time. From an econometric point of view, it is not possible to identify these three effects separately without invoking specific assumptions. The reason is that an individual's age is a deterministic linear function of its birth year (which determines his cohort) and the time at which data is collected. As a result, it is difficult to infer "true" age profiles for consumption/savings or asset allocation choices. This point is discussed in a illuminating way in Ameriks and Zeldes (2004).

---

[18]Thus, all minimum values of the exogenous random variables of the model, i.e. shocks to non-financial income and stock returns, are set to the first percentile of their originally assumed distribution.

In a recent study, Fernández and Krueger (2007) explore how consumption expenditures vary over the life cycle. Their study is based on the U.S. Consumer Expenditure Survey. The authors solve the identification problem mentioned above by assuming that time effects are orthogonal to a time trend. Their main finding is that consumption expenditures are hump-shaped over the life cycle. When not controlling for the fact that household size typically varies over the life cycle, household consumption expenditures take on their peak value around the late forties. The peak value is about 60 percent higher than consumption expenditures during the early twenties. From the fifties on, consumption falls steadily, reaching the level of the early twenties around the age of 65. When filtering out the effect of changing family size, the hump reduces to about half its size. For total expenditures the peak value is now achieved during the early fifties and is about 30 percent higher than consumption expenditures during the early twenties. Again, expenditures decrease steadily from about the age of 55 on.

The variation of asset allocation over age is addressed by Ameriks and Zeldes (2004). Again, they are confronted with the problem of separately identifying age, cohort and time effects. The authors tackle this problem by estimating age profiles twice, once excluding time effects and once excluding cohort effects.

Ameriks and Zeldes (2004) provide estimates for the variation over age of both, the decision of whether or not to participate in the stock market and, second, of what share of savings to invest in stocks, conditional on participating in the stock market. The evidence on stock market participation can be summarized as follows. It is fair to conclude that stock market participation is low during the first half of the twenties, say between 10 and 30 percent. Thus, most young people do not participate in the stock market. The participation rate increases up to about 50 percent at the age of 40. If there were no cohort effects participation would then slightly decrease to about 40 percent at the age of 65. If there were no time effects the participation rate would increase up to 80 percent.

Turning to equity shares conditional on stock market participation, the picture is as follows. In the absence of cohort effects equity shares are roughly constant at 40 percent over the entire

life. In the absence of time effects, equity shares start below 20 percent in the twenties and steadily increase to about 90 percent at the age of 65. While in reality both time and cohort effects are likely to play a role, Ameriks and Zeldes (2004) provide a number of reasons why they give priority to their specification without cohort effects over the specification without time effects. According to their most preferred estimation, equity shares conditional on stock market participation are thus basically flat.[19]

Table 2 presents some evidence on the variation of equity shares over *income*. Based on the 2001 wave of the SCF, the table reports median equity shares for particular age/income cells. Note that equity shares are zero for those who do not participate in the stock market. The table suggests that, for all age categories, members of the two bottom income quintiles typically do not participate in the stock market. In contrast, members of the upper quintiles do participate. While the evidence of Table 2 is descriptive and thus only suggestive, potential biases would have to be very strong in order to overturn the marked gradient effect concerning stock market participation. There is also rather strong evidence that equity shares increase with income among stock market participants.

6.2 Predictions of Existing Models

Existing life cycle models come in several variants. Most importantly in the current context, some more recent studies allow for investments in two different financial assets such as stocks and bonds. Other studies assume that there is only one investment possibility, typically a risk-free bond. Since the FG model aims to explain both the general level of savings and the share of

---

[19]The estimates of Ameriks and Zeldes (2005) discussed so far are based on the U.S. Survey of Consumer Finances (SCF). One potential concern may be that the estimated profiles may be confounded by the fact that people beyond a certain age are more likely to have access to employer-sponsored pension plans. This does not seem to be a main issue, however. Ameriks and Zeldes also provide estimations based on data from the TIAA-CREF institution, thus only including individuals who do have access to pension plans. While the stock market participation rate is substantially higher for the TIAA-CREF sample, the general shape of both the participation profile and the profile for equity shares conditional on participation are remarkably similar to the curves obtained for SCF data. This does not only hold for equity shares referring to total balances, but importantly so also for *inflows*.

savings invested in stocks (which may be zero), I only discuss here the predictions of existing life cycle models which allow for both stock and bond investments.

Such models successfully explain that consumption increases with age up to some peak level. However, the peak occurs to late. Empirically, the consumption profile decreases after the age of 50. In contrast, predicted consumption profiles peak only after retirement (see Cocco et al., 2005). This remains true for many deviations from the standard framework (Cocco et al., 2005; Gomes and Michaelides, 2003; Polkovnichenko, 2007).[20]

Turning to equity holdings, the performance of existing models is quite disappointing. Under standard constant relative risk aversion (CRRA) preferences, and taking into account liquidity constraints and a labor income process such as (5), the prediction is that stock market participation is universal. In addition, equity shares are predicted to be much higher than typically observed in the data. Specifically, equity shares are predicted to be close to one until the age of about 40. Thereafter, they are predicted to decrease until retirement where they reach a level of about 50 percent on average (Cocco et al., 2005). This contrasts with the empirical picture, according to which a large part of the young do not participate in the stock market. Furthermore, equity shares are roughly constant at 40 percent for those participating in the stock market; or, if there are non-negligible cohort effects, equity shares may increase for those participating in the market.

The prediction that equity shares are far higher than empirical estimates over a large part of the life cycle is very robust across existing models. Very high equity shares occur even when taking into account the possibility of very low realizations of labor income or endogenous borrowing constraints (Cocco et al., 2005), fixed costs of stock market participation and Epstein-Zin preferences (Gomes and Michaelides, 2005) or habit formation (Gomes and Michaelides, 2003, Polkovnichenko, 2007). The conclusion from this is that the search for models that come closer to explain the data remains an issue.

---

[20]Predicted consumption profiles come closer to the data in a model with endogenous borrowing constraints (Cocco et al., 2005). However, equity shares are not well explained within such a setup. (See below for equity shares.)

7 Calibration Results

This section presents the results from simulating the FG model. (See the Appendix on how to solve the model computationally.) A first subsection presents the main results for the baseline parameter values. A second subsection discusses predictions with respect to a variation in permanent income. A third subsection discusses a host of alternative specifications. For all these alternative specifications only one parameter deviates from the baseline value at a time.

7.1 Results for the Baseline Specification

Figures 2 and 3 present the results for the baseline case. The solid line in Figure 2 refers to consumption. The dashed line represents non-financial income. The dotted line represents total income including financial income, which is endogenous since it depends on the agent's savings and asset allocation choices. Note that the horizontal axis is restricted to ages of working life, as the model is less suited for explaining decisions during retirement.

Consumption tracks income very closely until the age of 35. Given that individuals cannot borrow against their future income the amount of liquid funds that are available to them is low during early ages. At the beginning of the twenties, current liquidity is not sufficient to finance a consumption level above $\bar{c}$ while, at the same time, meeting the insurance goal. Thus, optimal choices are determined by step 1 of the FG algorithm, i.e. individuals' principal goal is to come close to a normal standard of living $\bar{c}$. The only motive to save is to achieve the insurance goal.

At the age of 27 the standard of living $\bar{c}$ is reached, such that from then on optimal decisions are determined by step 3 of the FG algorithm. This means that the accumulation goal is triggered. As a result, individuals start to participate in the stock market (see Figure 3). However, consumption is still close to $\bar{c}$, such that the marginal value of an increase in consumption is still relatively high. This keeps total savings at a relatively low level until the age of 35. After the age of 35 savings start to increase substantially. While the gap between non-financial income and consumption narrows after the age of 45, the gap between consumption and total income,

including income from bond and stock holdings, widens until retirement. Overall, it is immediately evident from Figure 2 that pursuing the insurance and accumulation goals is sufficient for achieving a very smooth consumption profile.

Comparing the consumption profile in Figure 2 to the simulation results for the standard model in Cocco et al. (2005)[21] reveals that the FG model comes astonishingly close to the predictions of the standard model. For both models savings are slightly higher at the very beginning where young savers accumulate a small "buffer stock." Savings are then very low until the age of 35. In both models the gap between consumption and non-financial income widens from the age of 35 on, peaks around 45 and becomes then very small again. In the case of the standard model consumption crosses non-financial income at the age of 55. For the FG model consumption stays slightly lower than non-financial income, at least in the case of the baseline calibrations.

Figure 3 presents the predicted profile for equity shares.[22] Equity shares are defined as the ratio of total stock holdings to total financial assets, i.e. $s_t/\left(b_t + s_t\right)$. As already mentioned, agents are predicted to stay out of the stock market until the age of 26, since optimal choices are determined by step 1 of the FG algorithm. As discussed in Section 4, step 1 of the FG algorithm implies that stock investments are zero. Individuals start to participate in the stock market as soon as the accumulation goal is triggered at the age of 27. Thereafter, equity shares increase up to 50 percent at the age of 45. From there on they stay roughly constant until retirement.

The reason for the increase in equity shares between the ages of 27 and 45 is that the downside risk associated with future income declines substantially over time. Since the effects of a sequence of low realizations of the permanent income shock accumulate over time (see (6)), the uncertainty about future income and hence its downside risk increase with the number of periods that lie between the current and a particular future period. In this sense, a 25-year old faces more future income risk than a 45-year old. To put it simply, more things can still go

---

[21]See their Figure 3(A).

[22]The raw outcomes of the simulations for equity shares are slightly wiggling. The equity shares in Figure 3 and all following figures are smoothed using a moving average smoothing algorithm.

wrong for the 25-year old than for the 45-year old. The former simply has more opportunities to pick a long chain of low realizations of the permanent income shock. It is precisely this effect of decreasing future income uncertainty which leads equity shares to increase until the age of 45. Remember that the simulations have been run for the case where, ex post, no income shocks occur, i.e. the shocks always take on their expected values.

Whether equity shares are flat, increasing or decreasing after the age of 45 turns out to depend on the size of the lower bound of the medical expenditure shock $\underline{M}$. The medical expenditure shock has almost no influence on equity shares prior to the age of 45 (compare Figures 3, 5 and 7, drawn for $\underline{M}$-values of .6, .5 and .75, respectively). The reason is that it lies too far ahead in the future. Remember that agents do not want to let their future feasible standards of living fall below $\alpha^i c_t$, where i represents the distance between the current and a particular future period. In early ages retirement and hence the date of the realization of the medical expenditure shock lie far ahead in the future, such that α^i is very low. When agents get closer to retirement, the uncertainty of future labor income still decreases, as described in the last paragraph. This tends to raise equity shares further. However, α^{R-t} increases as agents get closer to retirement, at which the medical expenditure shock materializes. This makes agents more prudent and tends to decrease equity shares.

The net effect on equity shares may then be positive, negative or neutral. It happens that it is neutral in the baseline case (see Figure 3). In the case where $\underline{M}$ amounts to .5 (instead of .6), equity shares decrease after the age of 45 (see Figure 5). If $\underline{M}$ is equal to .75 then equity shares continue to increase after the age of 45 (see Figure 7). Figures 4 and 6 show that lower (higher) medical expenditure uncertainty leads to lower (higher) savings, without affecting the overall shape of the consumption profile.

In the case of asset allocation there is no close coincidence between the predictions of the FG and the standard model at all, unlike in the case of consumption profiles. For the FG model (in the baseline case), equity shares increase from zero to 50 percent until the age of 45 and then stay at this level until retirement. In contrast, in the standard model equity shares amount

25

to 100 percent until about the age of 40, where they start to decrease gradually to 50 percent, which is reached upon retirement.

While the equity share profiles that are predicted by the FG model do not directly coincide with the preferred estimations of Ameriks and Zeldes (2004), they nevertheless come much closer than the predictions of the standard model. First, the FG model is consistent with the fact that many do not join the stock market from the very beginning.[23] Second, with respect to equity shares conditional on stock market participation, the FG model is able to predict levels that are realistically low compared to the standard model. Furthermore, the model is consistent with a flat profile from the age of 45 on.

The favorable performance of the FG model raises the question whether this model just offers some additional degrees of freedom that are absent in the standard model. The FG model may indeed be seen as slightly more flexible than the standard model with CRRA preferences, which underlies much of the analysis in Cocco et al. (2005). However, what is crucial in this respect is that existing preference models have great difficulties to generate low overall equity exposures even when the flexibility of the setup is greatly increased, e.g. by assuming Epstein-Zin preferences, habit formation, by assuming heterogeneous preferences, by including fixed costs of stock market participation, endogenous borrowing constraints, by assuming the possibility of a disastrous labor income shock, among other possibilities.[24]

I conclude from this that the favorable performance of the FG model is not due to a particularly high degree of flexibility. Rather, it is likely to stem from the fact that it is more in line with an intuitive kind of reasoning that drives behavior of boundedly rational beings. In particular, according to the model, the young should reason that they face a lot of career uncertainty. They should wait and engage in the stock market more seriously only when this career uncer-

---

[23]In fact, they may not enter the stock market at all. As has been mentioned previously and will be discussed further in the next subsection, the FG model further predicts that if income is sufficiently low over the entire life cycle then agents will never enter the stock market. The reason is that choices are then always determined by step 1 of the FG algorithm. In contrast, in case of the standard model, the level of permanent income does not have any influence on stock market participation and equity shares, since these preferences are linearly homogeneous.

[24]See Cocco et al. (2005), Gomes and Michaelides (2003, 2005), and Polkovnichenko (2007).

tainty has sufficiently been reduced. This logic is highly intuitive and hence plausibly guiding the behavior of boundedly rational life cycle savers. However, it does not at all correspond to the logic of the standard model. There equity shares are very high until the age of about 40. The reason is the following. For the relevant class of expected utility preferences it turns out that the prospect of facing an increasing labor income profile with a high probability provides a substitute for risk-free savings. Hence, the young should invest all savings in stocks (see Cocco et al., 2005, for an illuminating discussion). According to personal experience of the author, even most economists are rather surprised by this result. So it is not completely unexpected that normal people don't seem to take the logic of the standard model as an intuitive guideline for their investment behavior.

7.2 The Variation of Stock Market Participation and Equity Shares with Permanent Income

Figure 8 shows how the FG model's predictions with respect to asset allocation differ for various levels of permanent income. The solid line corresponds to the baseline case. The dotted line corresponds to a case where non-financial income is 50 percent lower than in the baseline case at any age. Similarly, the dashed line correspond to a case where non-financial income is 50 percent higher than in the baseline case at any given age.

It follows from Figure 8 that the FG model is able to predict the stock holding patterns that are apparent in Table 2. In particular, if income is 50 percent below the baseline level, agents will never participate in the stock market since the accumulation goal is never triggered. Conditional on stock market participation, equity shares increase with a higher permanent income.

The ability of the model to explain that stock market participation and equity shares increase with income depends crucially on the assumption that $\bar{c}$ is greater than zero. For $\bar{c}$ equal to zero, these choices would not react to proportional increases in permanent income. Furthermore, equity shares would always be strictly positive (see Figure 10). Remember that expected utility models with CRRA as well as Epstein-Zin preference are not able to explain any variation of

27

savings or asset allocation choices with permanent income, since these preferences are linearly homogeneous.

7.3   Results for Alternative Model Specifications

Figures 9 and 10 refer to the case where $\bar{c} = 0$ (whereas all other parameter are set to baseline values). This is certainly the most parsimonious specification of the model. In this case only step 3 of the FG algorithm is relevant. Furthermore, the model is specified by basically only two free parameters: α^*, and β (apart from T^*). For $\bar{c} = 0$, consumption is slightly lower than in the baseline case up to the age of about 50 and higher thereafter, since agents have accumulated more wealth. Equity shares are substantially higher than in the baseline case. The young participate in the stock market right from the beginning, although their equity shares are very low. Thus, if there were any substantial costs for participating in the stock market, the young would probably stay out of the market even for $\bar{c} = 0$.

Figures 11 and 12 refer to the case where $\beta = .94$. This is the best fit value for β obtained in Binswanger (2006) for the cross-sectional setting considered there (see Section 5). All other parameters are again equal to their baseline levels. Naturally, consumption is slightly higher in this case at later ages since agents accumulate more wealth. This stems from the fact that equity shares are generally higher than in the baseline case, such that agents accumulate more wealth.

Figure 13 and 14 show simulations for a specification where $\bar{c}$ is indexed to family size. In particular, it is assumed that the household unit consists of one adult person until the age of 24. At the age of 25 a second adult person with the same age joins the household unit. The couple is assumed to get a first child at the age of 30 and a second one at the age of 33. Children leave home at the age of 18, i.e. when the parents are aged 48 and 51, respectively. I use the equivalence scales reported in Fernández-Villaverde and Krueger (2007) to index $\bar{c}$ to the family size.[25] Specifically, I assume that the baseline value of 27,000 year-2001 U.S. dollars refers to two adults and one kid. Using equivalence scales, a profile for $\bar{c}$ is calculated

---

[25]I follow the authors using the mean equivalence scales reported in their Table 1.

for the family history as mentioned above. This profile is then used for the simulations shown in Figures 13 and 14. The consumption profile resembles closely the baseline case. Since $\bar{c}$ is now lower during the twenties than in the baseline case, stock market participation occurs from the beginning on. Furthermore, due to the hump-shaped profile of $\bar{c}$, equity shares increase beyond the age of 45.

For the baseline case it is assumed that $T^* = 1/2\,t + 1/2\,T$. Figures 15 and 16 show simulations for the case where $T^* = 2/3\,t + 1/3\,T$. Thus, T^* moves closer to t. As a result, consumption becomes slightly hump-shaped. Equity shares are also hump-shaped in this case.

Overall, the consumption profiles resemble each other very closely across all specifications. Furthermore, equity shares are always very low if not zero in the twenties and increase thereafter. The precise shapes of the equity share profiles depend on the details of the specification.

8 Conclusion

This paper provides a new life cycle framework in which agents do not engage in full contingent planning but pursue two simple feasibility goals. The paper shows that pursuing such goals is sufficient for obtaining very smooth consumption profiles. Predicted consumption profiles are very similar in shape to the profiles for the standard expected-utility case. In contrast, the pursuance of feasibility goals leads to a much lower equity exposure than in the standard model. In particular, the feasibility goals (FG) model is consistent with the fact that low-income earners and many young do not participate in the stock market. Those who do participate choose relatively low equity shares. As a result, the FG model is better than existing preference models at explaining the empirical patterns of stock market participation and the variation of equity shares over the life cycle.

9 Appendix: The Computational Solution of the FG Model

For each age t the simulation starts with step 1 of the FG algorithm. As explained in Section 4, step 1 implies $s_t = 0$. Furthermore, the requirement of meeting the insurance goal implies for each value of c_t a corresponding level of bond investments. From a computational point of view, it is easiest to find this level of bond investments using a backward calculation. Specifically, set the terminal value $\hat{b}_T$ in the recursion $\hat{b}_{t+i} = \max\left\{\left(\hat{b}_{t+i+1} + \alpha_{i+1}c_t - \min_t Y_{t+i+1}\right)/r,\ 0\right\}$, $0 \leq i \leq T - t - 1$ to zero. The required level of bond investments for a given level of c_t is then obtained by solving this recursive equation backward for $\hat{b}_t$. The choice that solves the step-1 problem of the FG algorithm is determined by the fact that $c_t + \hat{b}_t(c_t) = X_t$, where X_t denotes the total level of resources that are available at age t. Simple algorithms can be used to find the optimal level of c_t.

According to step 2 of the FG algorithm, the optimal choice is the outcome of step 1 whenever the resulting value of c_t falls short of $\bar{c}$. Otherwise, the step-3 problem has to be solved. In order to do so, start with a particular level of c_t. In order to evaluate the function (4), the value of the accumulation goal $a_{t|T^*}$ has to be determined. Remember that the accumulation goal is defined as the result of a hypothetical accumulation from age t to age T^* such that, in each period, consumption is kept constant at c_t and a portfolio is chosen for which the value of stocks is maximal, under the constraint the the insurance goal has to be met.

Given a starting value of c_t, savings are equal to $X_t - c_t$. The first step is to find the maximum value of stocks s_t such that $b_t + s_t = X_t - c_t$ and the insurance goal is met. Although not very efficient, the simplest algorithm would solve this problem by starting with the highest possible value $s_t = X_t - c_t$. To check whether this choice allows to meet the insurance goal it is sufficient to check whether this goal can be met if, from the next period on, all savings are invested in bonds, since minimum stock returns are smaller than bond returns. Furthermore, consumption is set to the habit level in all following periods. Check thus whether the insurance goal can be met by setting consumption in the next period to $\alpha_1 c_t$ and bond investments in the next period to $\min_t X_{t+1} - \alpha_1 c_t$, and, iterating further, setting $c_{t+i} = \alpha_i c_t$, $b_{t+i} = \min_t X_{t+i} - \alpha_i c_{t+i}$.

The insurance goal is feasible if all values b_{t+i} are nonnegative for $i = 1, 2, \ldots, T - t - 1$ and $X_T \geq \alpha_{T-t} c_t$. If this holds for an initial value $s_t = X_t - c_t$, then this and $b_t = 0$ is the optimal investment choice at age t for the chosen initial value of c_t. Otherwise repeat the same procedure for gradually smaller values of s_t and higher values of b_t, until the highest value of s_t is found for which the insurance goal can be met. (Possibly, the insurance goal cannot be met for any portfolio, in which case the accumulation goal takes on a value of zero.) In order to continue the calculation of the accumulation goal, still for a given initial value of c_t, one has to iterate over time the process of finding the highest level of stock investments that is compatible with meeting the insurance goal, taking into account equation (3). This will finally lead to a value of the accumulation goal for a given starting level of c_t. The efficiency of the algorithm outlined here may be improved by slightly refining its baseline version.

Since the objective function (4) is strictly concave, simple algorithms can finally be used to find the optimal value of c_t.

References

Allen, Todd W. and Carroll, Christopher D. (2001), "Individual Learning about Consumption," *Macroeconomic Dynamics*, 5(2), pp. 255-71.

Ameriks, John, and Zeldes, Stephen P. (2004) "How Do Household Portfolio Shares Vary With Age?" Working Paper.

Binswanger, Johannes (2006) "A Simple Bounded-Rationality Life Cycle Model," Working Paper.

Bowman, David, Minehart, Deborah, and Rabin, Matthew (1999), "Loss Aversion in a Consumption-Savings Model," *Journal of Economic Behavior and Organization*, 38(2), pp.155-178.

Campbell, John Y., and Viceira, Luis M. (2002), *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press, Oxford.

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

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

Fernández-Villaverde, Jesús, and Krueger, Dirk (2007), "Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data," *Review of Economics and Statistics*, 89(3), pp. 552-565.

Gabaix, Xavier, and Laibson, David (2000), "A Boundedly Rational Decision Algorithm," *American Economic Review*, Papers and Proceedings, 90, pp. 433-438.

Gabaix, Xavier, and Laibson, David (2005), "Bounded Rationality and Directed Cognition," Working Paper.

Gollier, Christian (2001), "What Does the Classical Theory Have to Say about Household Portfolios?" In: Guiso, Luigi, Haliassos, Michael and Jappelli, Tullio (Eds.), *Household Portfolios*, MIT Press, Cambridge MA, pp. 27-54.

Gomes, Francisco, and Michaelides, Alexander (2003), "Portfolio Choice with Internal

Habit Formation: A Life-Cycle Model with Uninsurable Labor Income Risk," *Review of Economic Dynamics*, 6, pp. 729-766.

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

Hubbard, R. Glenn, Skinner Jonathan, and Zeldes, Stephen P. (1995), "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103(2), pp. 360-399.

Keynes, John Maynard (1936), *The General Theory of Employment, Interest and Money,* MacMillian, London.

Laibson, David (1997), "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112(2), pp. 443-477.

Laibson, David, Repetto, Andrea, and Tobacman, Jeremy (1998), "Self-Control and Saving for Retirement," *Brookings Papers on Economic Activity*, 1998(1), pp. 91-196.

Lettau, Martin, and Uhlig, Harald (1999), "Rules of Thumb versus Dynamic Programming," *American Economic Review*, 89(1), pp. 148-174.

Pemberton, James (1993), "Attainable Non-Optimality or Unattainable Optimality: A New Approach to Stochastic Life Cycle Problems," *The Economic Journal*, 103(416), pp. 1-20.

Polkovnichenko, Valery (2007), "Life-Cycle Portfolio Choice with Additive Habit Formation Preferences and Uninsurable Labor Income Risk," *Review of Financial Studies*, 20(1), pp. 83-124.

Scholz, John Karl, Seshadri Ananth, and Khitatrakun Surachai (2006), "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy*, 114(4), pp. 607-643.

Shefrin, Hersh M., And Thaler, Richard H. (1988), "The Behavioral Life-Cycle Hypothesis," *Economic Inquiry*, 26(4), pp. 609-643.

Thaler, Richard H. (1994), "Psychology and Savings Policies," *American Economic Review*, Papers and Proceedings, 84(2), pp. 186-192.

33

Table 1: The baseline specification

T_0	21
R	66
T	86
T^*	$(t+T)/2$
r	1.02
Er^s	1.06
$r^{s\,\min}$	0.74
$\underline{U}$	0.97
$\underline{V}$	0.81
ρ	0.6
$\underline{M}$	0.6
$\{\alpha_i\}$	0.96^i
β	0.9
$\bar{c}$	$27,000^{\,a)}$

$^{a)}$Year-2001 U.S. dollars.

Table 2: Median equity shares for age-income cells

Income quintiles	Age					
	<30	30-39	40-49	50-59	60-69	70+
First	0	0	0	0	0	0
Second	0	0	0	0	0	0
Third	0.28	0.36	0.36	0.26	0.26	0.36
Fourth	0.46	0.56	0.48	0.49	0.44	0.41
Fifth	0.65	0.54	0.60	0.57	0.56	0.55

NOTE. Source is the 2001 wave of the U.S. Survey of Consumer Finances.

Figure 1: Non-financial income in the absence of shocks
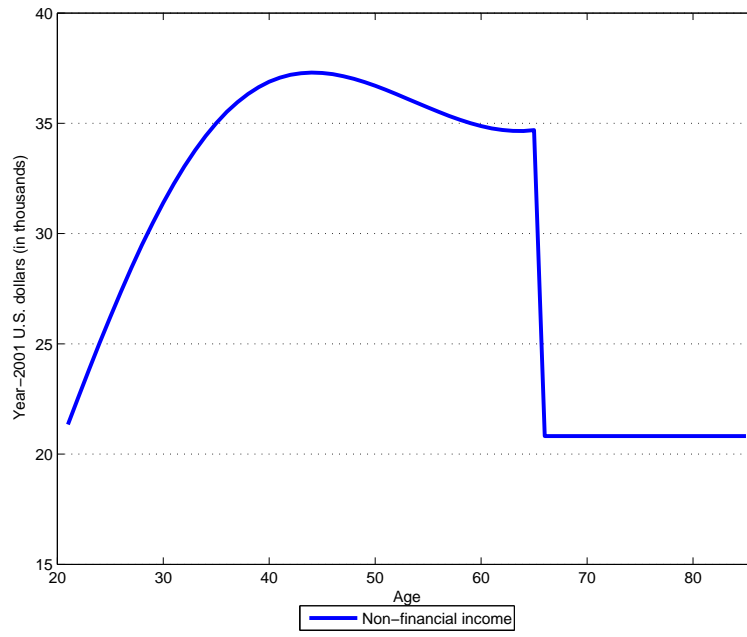
Figure 2: Consumption and income, baseline
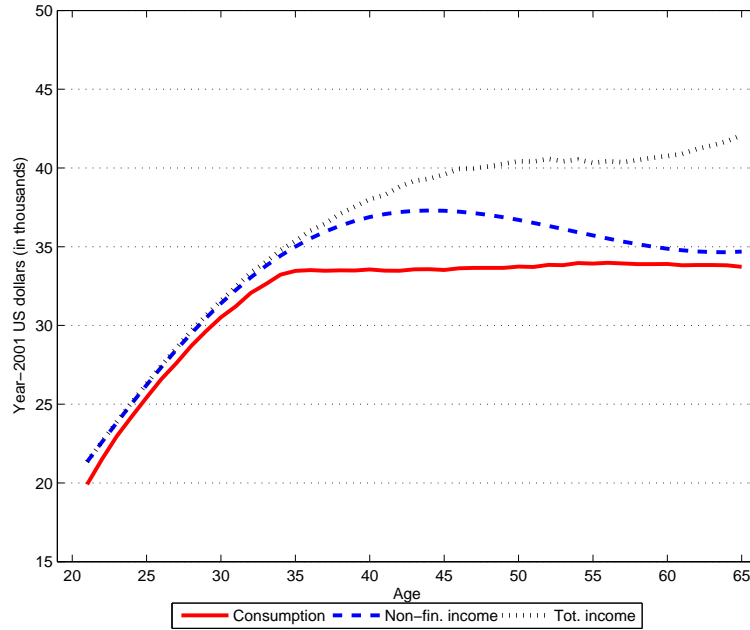
Figure 3: Equity shares, baseline
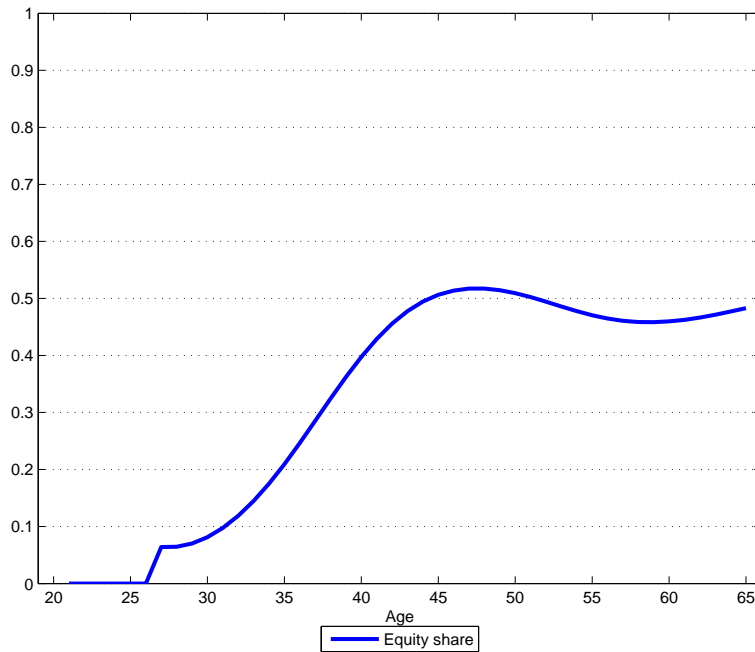
Figure 4: Consumption and income, higher medical expenditure uncertainty
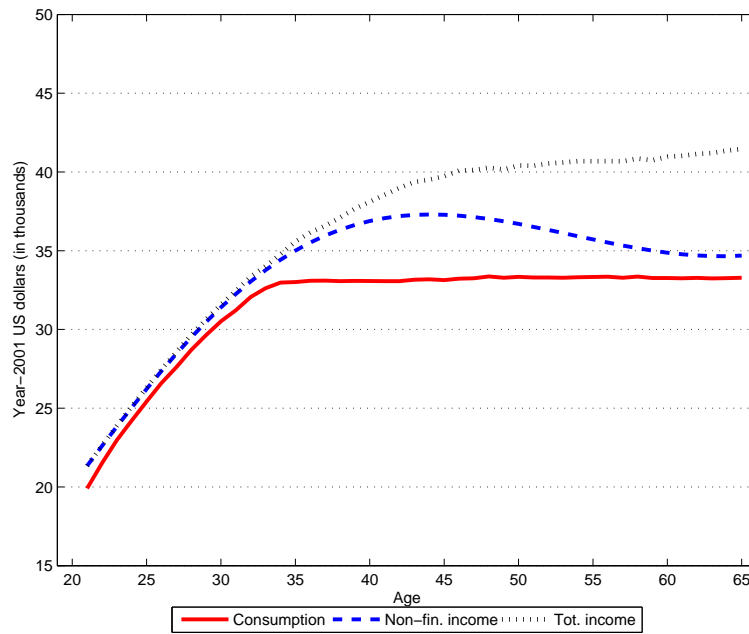
Figure 5: Equity shares, higher medical expenditure uncertainty
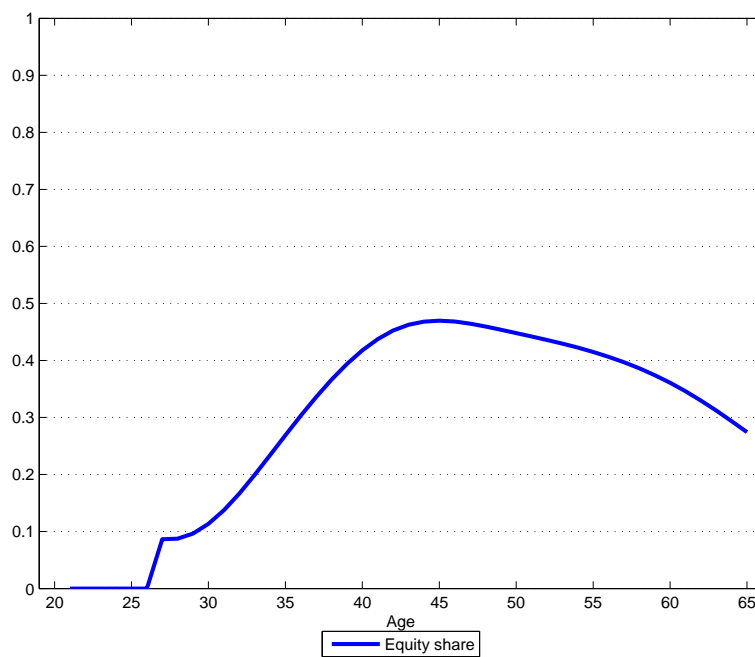
Figure 6: Consumption and income, lower medical expenditure uncertainty
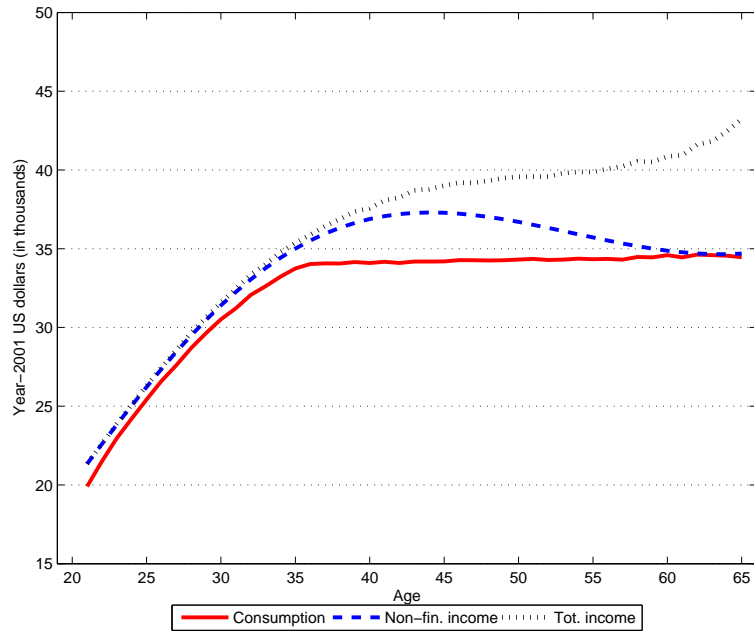
Figure 7: Equity shares, lower medical expenditure uncertainty
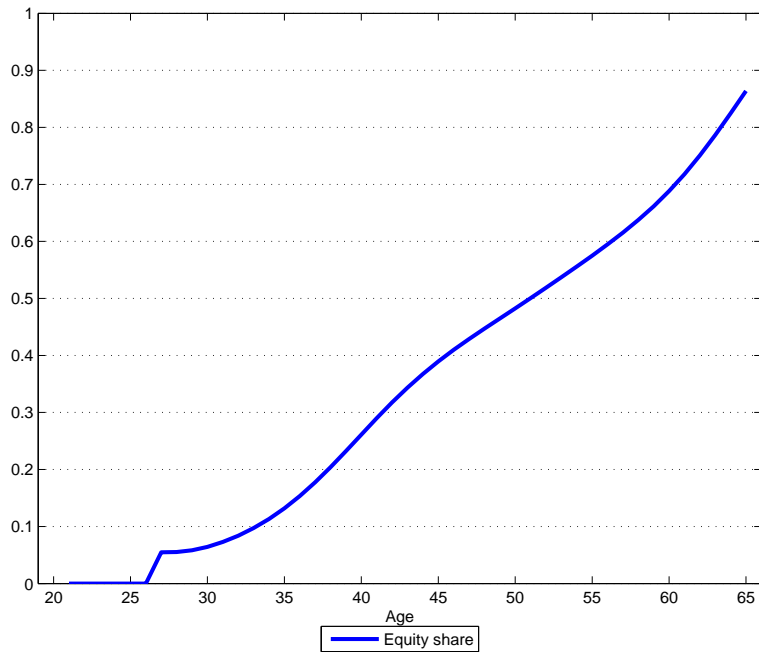
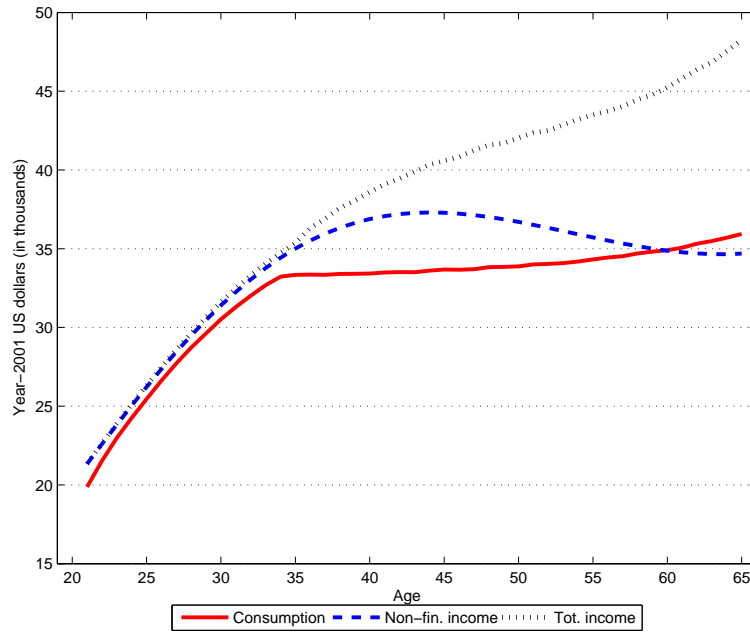Figure 8: Equity shares for various income levels
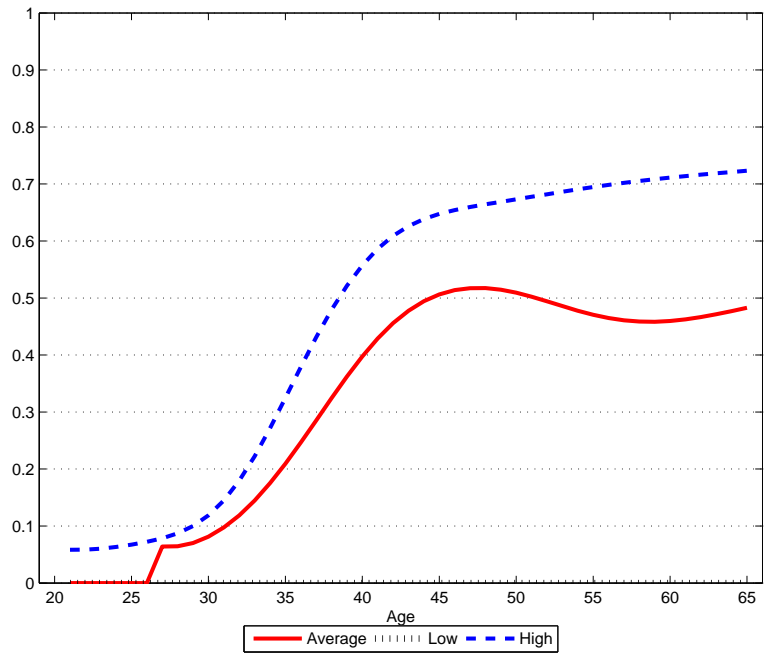
Figure 9: Consumption and income, $\bar{c} = 0$

Figure 10: Equity shares, $\bar{c} = 0$
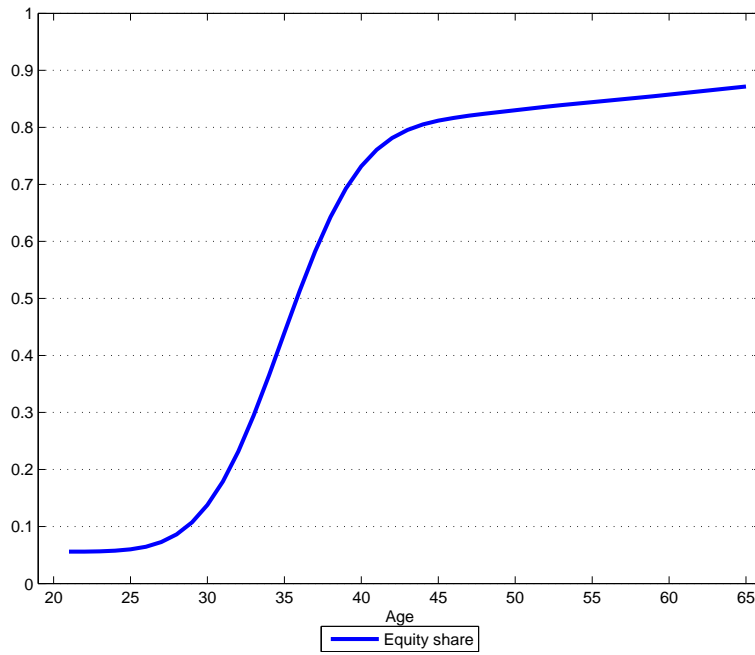
Figure 11: Consumption and income, $\beta = 0.94$
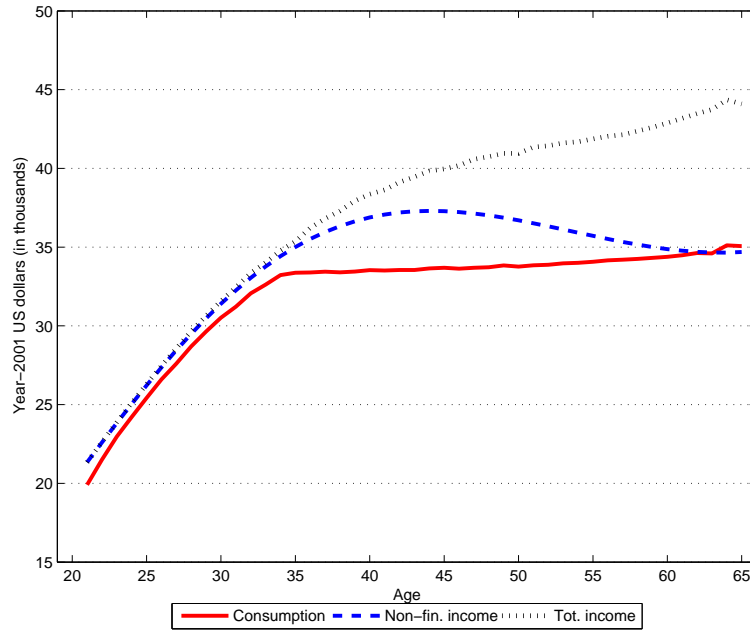
Figure 12: Equity shares, $\beta = 0.94$
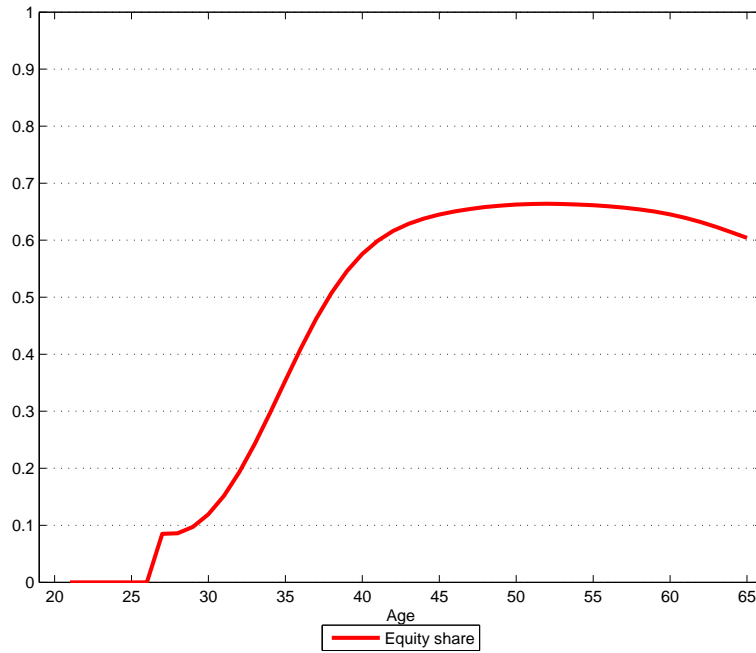
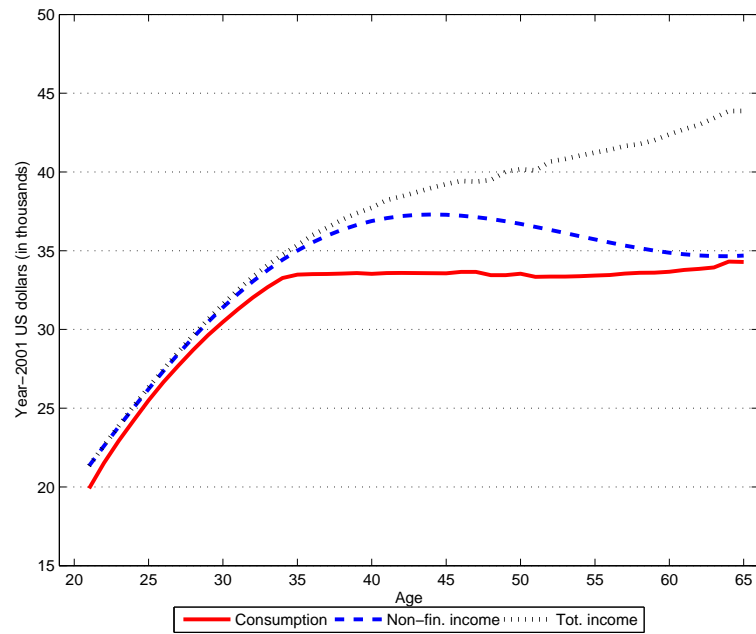Figure 13: Consumption and income, $\bar{c}$ indexed to household size

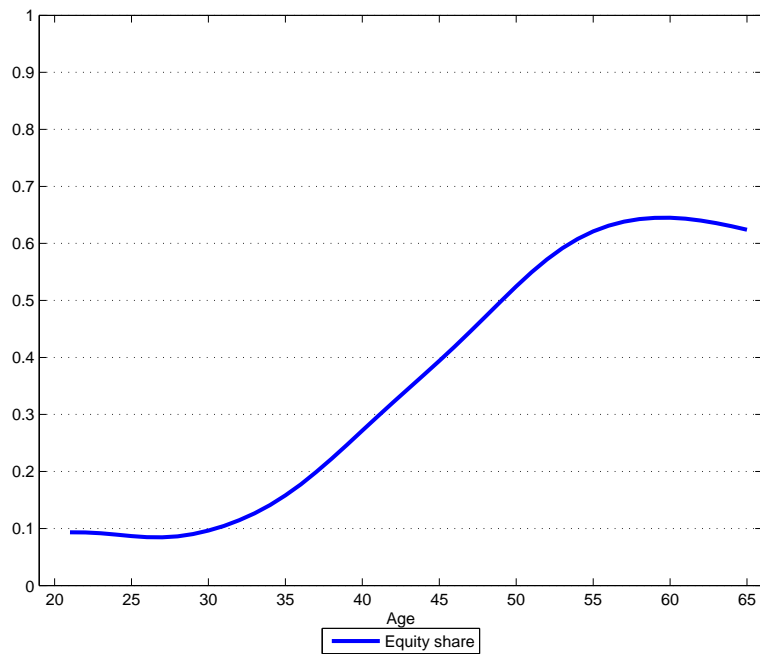Figure 14: Equity shares, $\bar{c}$ indexed to household size

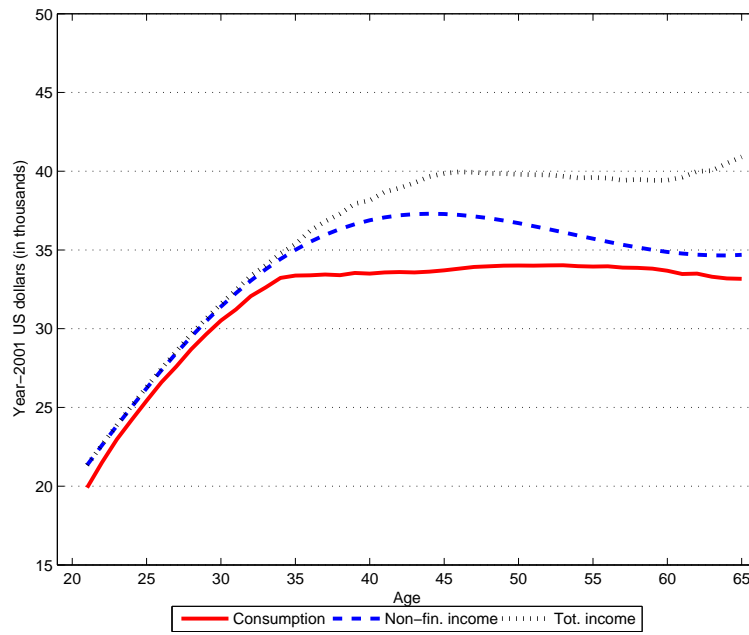Figure 15: Consumption and income, $T^* = 2/3\, t + 1/3\, T$

Figure 16: Equity shares, $T^* = 2/3\, t + 1/3\, T$