Tilburg University

Towards Understanding Life Cycle Savings of Boundedly Rational Agents

Binswanger, J.

Publication date:
2010

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Binswanger, J. (2010). *Towards Understanding Life Cycle Savings of Boundedly Rational Agents: A Model with Feasibility Goals (Revision of DP 2008-14)*. (CentER Discussion Paper; Vol. 2010-138). Economics.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

No. 2010-138

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

By Johannes Binswanger

December 2010

This is a revised version of CentER Discussion Paper
No. 2008-14

February 2008

ISSN 0924-7815

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

Johannes Binswanger[†]

December 9, 2010

Abstract

This paper develops a new life cycle model that aims to describe the savings and asset allocation choices of boundedly rational agents. In this model, agents make forward-looking decisions without the requirement of anticipating their actual future decisions. Instead, agents pursue two simple so-called feasibility goals. The first goal entails assuring feasibility of a certain minimum consumption level in a future worst-case scenario. The second concerns feasibility of a certain standard of living in a normal scenario. The feasibility goals framework represents a particularly natural alternative to the standard rational benchmark model. The framework is highly tractable and parsimonious. In particular, it is able to explain important empirical patterns of asset allocation that are puzzling from the point of view of existing models.

Key words: Behavioral economics, bounded rationality, contingent planning, feasibility goals, life cycle saving, portfolio choice.

JEL classification: D81, D91.

*I am grateful to Lans Bovenberg, Simon Gächter, David Laibson, Wieland Müller, Manuel Oechslin, Ward Romp, Karl Scholz, Eytan Sheshinski, Steve Zeldes and seminar participants at the Austrian National Bank, at CeDex (University of Nottingham), the ENABLE Symposium on Economics and Psychology, Netspar, the 10th Annual Research Conference of the Dutch National Bank, and Tilburg University for helpful comments. All errors are my own.

[†]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 version of the economic life cycle model. According to this model, individuals maximize the discounted sum of expected per-period utilities, while being exposed to non-insurable labor income shocks and while facing borrowing constraints. This model has been partially successful in that predicted consumption and savings profiles over the life cycle come relatively close to the respective empirical observations. 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 particular, the model predicts that a much higher proportion of the population will own stocks than is actually observed. 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, a methodological weakness of the modern life cycle model is its assumption of unlimited rationality (Thaler, 1994). Several findings in the literature suggest that this assumption may indeed be problematic and cannot easily be justified by standard arguments such as learning. 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. 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 to make optimal decisions (in the sense of the stan-

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

²See Section 7 for a further discussion.

dard model), since learning may not even asymptotically lead to optimal behavior.

The main reason why decision making is so complex in the case of the modern life cycle model (henceforth called the standard model) is the latter's requirement of anticipating one's future actual decisions by means of a contingent plan. The model assumes that agents maximize the discounted sum of expected per-period utilities from consumption. Suppose that an agent makes a saving 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!³ 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 such contingent planning may greatly exceed the capabilities of boundedly rational agents.⁴

This paper takes it as given that it is desirable to develop 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-savings choices at least equally well.

The approach I pursue is to cut out the main source of complexity in the standard model, namely the requirement of contingent planning. In particular, I build on the framework of *feasibility goals* that is developed in Binswanger (forthcoming). In this framework, agents care about the *feasibility* of certain consumption levels in the future, i.e. that certain choices are contained in future budget sets. However, agents do not determine

³See Binswanger (2008) for a more thorough discussion.

⁴The importance of this has often been downplayed by referring to the argument that agents do not actually need to solve this complicated planning 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 the 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. Thus, in the domain of life cycle saving, people may rather be seen as novice players.

in advance which bundle to choose out of future budget sets. Put differently, agents are forward-looking in the sense that when they make choices about current consumption, they take into account the impact on future budget sets. However, *actual* future consumption choices are left completely open. Under these circumstances, the objective function in the standard model would not be well-defined. However, this paper shows that agents may still make prudent and forward-looking decisions. For obvious reasons, the model developed in this paper is dubbed feasibility goals (FG) model.

One major difficulty encountered when developing a new bounded-rationality life cycle model is that, at first sight, there seem to be many possible ways how to deviate from the standard model. Thus, while many may agree that the standard model has substantial weaknesses, economists have been hesitant to propose new models of bounded rationality, for fear that any particular model they would propose may be seen as arbitrary since not backed by a set of elegant axioms.

This issue has to be taken very seriously. However, the FG model represents a particularly natural way how to deviate from the standard model. This statement is based on the following argument. Suppose that we agree that the assumption that agents anticipate their future decisions in a precise way is unrealistic because of bounded rationality. This precludes contingent planning and hence the standard model. However, if we still find it desirable to assume that agents make forward-looking decisions, then we may make the less demanding assumption that agents are concerned about how their current decisions affect their future choice sets. In other words, agents may not want to make current choices that would mean to forgo feasibility of certain future choices that they may find desirable. Such a framework corresponds exactly to the FG approach put forward in this paper.⁵

This paper develops a simple and parsimonious model where agents have two feasibility goals that induce two different motives of saving. The first is to insure that future

⁵Any other approach would most likely rely on rule-of-thumb behavior. I do not follow the latter approach since it is desirable to first explore whether we may identify a useful model that deviates less drastically from the standard model than a rule-of-thumb approach would.

standards of living may not fall too far below the current one under unfavorable circumstances. The second is a desire to achieve a high *expected* standard of living in the future. These particular goals are specified in order to induce a particularly simple form of a risk-return trade-off. The presence of such a trade-off is a necessary requirement for any sensible model of life cycle saving and asset allocation.

While I will carefully discuss the mathematical form of the simple optimization programs that determine behavior in case of the FG model, it will ultimately not be solved analytically. The reason is straightforward. Any specification of a life cycle setup that is sufficiently realistic to be of practical relevance is necessarily too involved to allow for an insightful analytical solution. It entails taking into account transitory and permanent shocks to non-financial income and a multitude of periods. While an analytical solution of the FG model would in principle be feasible (unlike in the case of the standard model), the resulting expressions would be too cumbersome to offer any interesting insights. For this reason, the ultimate aim of this paper is to provide calibrations of the model for a representative scenario. This is done in a way such that the model's predictions can be compared to the predictions of existing models as well as to empirical estimates.

However, the model can be solved analytically in an insightful way for some interesting benchmark cases. These include a three-period and a stylized infinite horizon setup. Analytical solutions for these two benchmark cases are contained in a separate paper (Binswanger, 2008). They highlight the mechanisms at work and show that the qualitative predictions of the model are highly robust.

When calibrating the FG 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 during their entire life. Finally, conditional on stock market participation, equity shares increase with permanent income.

The paper is organized as follows. Section 2 discusses how this paper relates to the existing literature. Section 3 introduces the FG model. Section 4 shows how behavior is determined in case of the FG model. Section 5 discusses some methodological issues. Section 6 explains how the model is calibrated. Section 7 provides a brief overview of some stylized facts about how consumption, savings, and asset allocation vary over the life cycle. It also contrasts these facts with the predictions of existing models. Section 8 presents the calibration results for a baseline calibration of the FG model. Section 9 provides results for alternative specifications. Section 10 concludes. Some details concerning the calibration of the model are contained in an appendix.

2 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 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 tree. 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 contingent planning models has been made by Pemberton (1993). In his model, an agent making a current decision simply decides what share of her available resources to spend now, and what share she will carry forward into the future in general, without specifying any particular future allocation. In contrast, in case of the FG model agents plan in a more specific way since they want to assure feasibility of certain consumption levels at particular dates and in particular states of the world in the future.

The feasibility goal concept is developed in Binswanger (forthcoming). In Binswanger

(2008), I derive analytical solutions for the FG model of this paper for certain useful benchmark cases. These include a three-period setup as well as a stylized infinite-horizon setup. It is shown that the FG model is highly tractable compared to existing models. The analytical results are complementary to the simulations contained in this paper since they allow for a better understanding of the mechanisms behind the simulation results. Binswanger (2010) analyzes the variation in savings and asset allocation choices across income groups using a model which may be seen as a static version of the FG model of this paper. In Binswanger (2007), I discuss the implications income heterogeneity for pension risk management, using this static version of the model.

3 The FG 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 working, 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 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 and captures the nature of financial life cycle decisions in a stylized way.⁷

Since it is the concept of feasibility that underlies the model to be developed, it is natural to start with a discussion of intertemporal budget constraints. 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*

⁶See Section 9 for a specification of the model where the size of the household changes over time.

⁷An 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.

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 $t + i$, $i = 1, 2, \dots, T - t$, 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 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

$$b_{t|t+i} \geq 0, \quad s_{t|t+i} \geq 0 \quad (1)$$

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

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.⁸ 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 $r > r^{s \min} \geq 0$, where $r^{s \min} \equiv \min r^s$.

At each age t the agent earns an exogenous and possibly risky stream of non-financial income Y_t . This income may be exposed to transitory and permanent shocks and, in expectation, follow a typical hump-shaped pattern over the life cycle. Thus, the process for Y_t is generally not *iid*. For the time being, there is no need specify this process any further. I will make specific assumptions when it comes to the calibration of the model (see Section 6 and the Appendix). Denote future non-financial income conditional on

⁸This assumption is not crucial for the behavior of the FG model and could be dropped.

information at age t by $Y_{t|t+i}$. The intertemporal budget constraint is then given by

$$c_{t|t+i} + b_{t|t+i} + s_{t|t+i} = Y_{t|t+i} + b_{t|t+i-1}r + s_{t|t+i-1}r^s \equiv X_{t|t+i} \quad (2)$$

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 intertemporal 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 two feasibility goals that are at the heart of the FG model. To this end, it proves convenient to introduce a more compact notation. Let $\mathbf{z}_{t|t+i} := (c_{t|t+i}, b_{t|t+i}, s_{t|t+i})$, $0 \leq i \leq T - t$ denote a choice that is feasible at age $t + i$ from the perspective of t , and denote the set of all feasible choices at age $t + i$ from the perspective of t by

$$\mathbf{B}_{t|t+i}(\mathbf{z}_{t|t+i-1}, Y_{t|t+i}) := \{\mathbf{z}_{t|t+i} \in \mathfrak{R}_+^3 : c_{t|t+i} + b_{t|t+i} + s_{t|t+i} \leq X_{t|t+i}(\mathbf{z}_{t|t+i-1}, Y_{t|t+i})\}.$$

where $X_{t|t+i}$ is defined as in (2).

The first goal is dubbed insurance goal and relates to a precautionary motive of saving. It is 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*

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

and

$$(\alpha_{i+1} c_t, b_{t|t+i+1}^*, s_{t|t+i+1}^*) \in \mathbf{B}_{t|t+i+1}$$

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 i periods ahead in the future in all circumstances. This includes the worst-case scenario, which corresponds to a sequence of lowest possible values for stock returns and non-financial incomes.⁹ In principle, the sequence $\{\alpha_i\}$ is unrestricted and could even be non-monotonic (except that the elements are to be nonnegative).¹⁰ Note 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. Here, this is guaranteed by the existence of a risk-free asset.

The insurance goal captures in a basic way the desire of assuring that the current standard of living will not deteriorate too much in the future, even in unfavorable states of the world. Clearly, the formulation of the insurance goal reminds of existing models of habit formation. However, as will be discussed below in Section 5.1 and 5.3, there are major differences between the FG model and existing models of habit formation.

The second goal determining behavior in the FG model is dubbed 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.¹¹ 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 fact that the accumulation goal refers to a normal scenario while the insurance goal implicitly refers to a worst-case scenario introduces a very simple risk-return trade-off.

The accumulation goal captures the desire of enjoying an increasing standard of living,

⁹This should not be taken too literally. The worst-case scenario should be understood as corresponding to, say, the first percentile of the respective random variables.

¹⁰When 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” since $\alpha_{i+1} < \alpha_i$. See Section 6 for a more detailed discussion.

¹¹An 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 9 for a different specification.

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).

At this point, it is crucial to realize that a standard of living that will be feasible from some future age $T^* > t$ on generally depends not only on choices made at age t , but also on savings and asset allocation choices between t and T^* . This is exactly what induces the necessity for contingent planning in existing models of intertemporal decision making. It is the *raison d’être* of the FG model that the decision maker does not engage in contingent planning. This effectively entails the assumption that the decision maker does not know her future decisions. In particular, she does not anticipate what consumption and investment choices she will make between t and T^* in the normal scenario.

How can a forward-looking individual then nevertheless make sure of not forgoing feasibility of a certain standard of living from period T^* on in a normal scenario? There is basically one possibility, entailing that the agent ensures feasibility of a particular *hypothetical* accumulation scheme that, in turn, leads to feasibility of a particular standard of living from age T^* on. Making current decisions such as not to forgo feasibility of this accumulation scheme implies not forgoing feasibility of a particular standard of living from age T^* on.

The hypothetical accumulation scheme that underlies the definition of the accumulation goal 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 starting the hypothetical accumulation scheme.

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, i.e. $r_{t+1}^s = Er^s$, $Y_{t+1} = E_t Y_{t+1}$. It is now assumed that the hypothetical accumulation scenario entails replicating the period- t decision. Thus, $c_{t|t+1}^n = c_t$, where the superscript n indicates the normal scenario. Furthermore, the agent chooses again the portfolio that maximizes

the expected return, subject to achieving the insurance goal. The latter is determined conditional on the event that the normal scenario has prevailed until period $t+1$ but may not prevail any further. Thus, the agent projects a prudent portfolio choice. Iterating an identical decision process forward in time until age T^* and projecting that $r_{t+i}^s = Er^s$, $Y_{t+i} = E_t Y_{t+i}$, $1 \leq i \leq T^* - t$, one ends up with a hypothetically feasible level of resources $X_{t|T^*}^n$ at age T^* .

The accumulation goal is then defined as $X_{t|T^*}^n / (T^* - t)$. The aim of this normalization is to convert $X_{t|T^*}^n$ into 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 , its value should be understood as a function of c_t . An 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 $\mathbf{I}_{t|t+i}^n(c_{t|t+i})$ the set of all portfolios $(b_{t|t+i}, s_{t|t+i})$ that are feasible such that the insurance goal is satisfied for $(c_{t|t+i}, b_{t|t+i}, s_{t|t+i}) \in \mathbf{B}_{t|t+i}^n$, given that the normal scenario has prevailed between t and $t+i$. (All superscripts n refer to the normal scenario.) The set $\mathbf{I}_{t|t+i}^n$ 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 $c_{t|t+i}$, define $(b'_{t|t+i}, s'_{t|t+i})$ as the element of $\mathbf{I}_{t|t+i}^n(c_{t|t+i})$ 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 $\mathbf{I}_{t|t+i}^n$ is nonempty such a portfolio exists. It is assumed that, given $c_{t|t+i}$, the agent considers to choose $(b'_{t|t+i}, s'_{t|t+i})$ for the hypothetical accumulation scenario at age $t+i$. (It will become clear from the FG algorithm discussed below that $\mathbf{I}_{t|t+i}^n$ will always be nonempty along the “optimal” trajectory.) Furthermore, the agent projects a choice of $c_{t|t+i} = c_t$. The hypothetical wealth accumulation scheme is determined by

$$X_{t|t+i+1}^n = E_t Y_{t+i+1} + b'_{t|t+i} r + s'_{t|t+i} E r^s. \quad (3)$$

The precise definition of the accumulation goal is then as follows.

Definition 2 Take $c_t \in \mathbf{B}_t$ as given. If all sets $\mathbf{I}_{t|t+i}^n(c_t)$, $0 \leq i \leq T^* - t - 1$, are nonempty, then the value of the accumulation goal, denoted by $a_{t|T^*}$, is equal to $X_{t|T^*}^n / (T - T^*)$. Otherwise $a_{t|T^*} = 0$.

Note that, according to this definition, $a_{t|T^*}$ is well-defined for all values of $c_t \in \mathbf{B}_t$. Overall, behavior in the FG model is determined according to the following algorithm dubbed FG algorithm.

Definition 3 (FG algorithm) The FG algorithm entails 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. If the resulting value for c_t falls short of a target level $\bar{c} \geq 0$, then the agent implements this choice. Otherwise, the optimal choice is determined by step 2.

Step 2. Choose the element of $\mathbf{B}_t$ which maximizes

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

where $0 < \beta < 1$.

Step 1 just means maximizing c_t , subject to meeting the insurance goal. The parameter $\bar{c}$ represents a “normal standard of living.” The economic rationale of this parameter is to account for the fact that people may be concerned about how their standard of living compares to a reference consumption level (see Clark et al., 2008). This reference level may reflect the average standard of living in the economy.¹² From a mathematical point of view, this parameter makes decision making non-homothetic. For instance, it implies that saving rates and the fraction of savings allocated to stocks increase with lifetime income (see Section 5.2 for a further discussion).

The objective function (4) in step 2 specifies the way in which the agent solves the trade-off between her goal of enjoying a high level of current consumption and of the

¹²This suggests that $\bar{c}$ is likely to increase over time (see Binswanger, 2008, for a further discussion). In this paper, $\bar{c}$ is simply treated as exogenous.

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 current consumption. The weight depends on the distance between current age and the future date T^* , which is the age the accumulation goal refers to. A further discussion of the objective function (4) as well as of the question whether the outcome of the FG algorithm is always well defined is postponed to the next section.

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, irrespective of her earlier choices and plans. This feature of the model reflects an element of procedural rationality. The question whether this leads to time-inconsistent behavior is discussed in Section 5.1.

4 Decision Making under the FG Model

This section states the precise mathematical form of the decision problems that determine optimal behavior under the FG model. This will allow for a better understanding of the mechanisms of the model. Ultimately, the model will not be solved analytically, for the reasons mentioned in the introduction. A full analytical solution of the model for the case of three periods and for an infinite horizon can be found in Binswanger (2008).

The FG model is an optimizing model like any standard model of intertemporal choice. However, the decision problems that agents solve are conceptually simpler than in the case of the standard model. Most notably, the decision problems are such that there is no need to engage in contingent planning.

To begin with, consider behavior in a particular period t in case that it is determined by step 1 of the FG algorithm. The optimization problem is then to maximize c_t subject to assuring that $\alpha_i c_t$ remains feasible in all circumstances in future periods $t + i$, $i = 1, \dots, T - t$. 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 were fulfilled, stocks could be substituted by bonds and, at the same time, c_t could be increased, in a way that the insurance goal would still be fulfilled.¹³

It is instructive to consider the step-1 problem in the absence of any borrowing constraints. The present value of future “habit” consumption levels is given by $\sum_{i=1}^{T-t} \alpha_i c_t / r^i$. An agent’s savings are then determined by how much this amount exceeds the present value of future minimum income. The latter amounts to $\sum_{i=1}^{T-t} Y_{t|t+i}^{min} / r^i$, where $Y_{t|t+i}^{min}$ denotes the minimum of Y_{t+i} , given information in period t . Thus, assuming that there are no (binding) borrowing constraints, the decision maker would maximize c_t subject to $b_t = \sum_{i=1}^{T-t} \alpha_i c_t / r^i - \sum_{i=1}^{T-t} Y_{t|t+i}^{min} / r^i$ and $b_t + c_t = X_t$.

In the case of borrowing constraints, the algebra is slightly more involved. Given a level of current consumption c_t , the level of bonds necessary to achieve the insurance goal is then given by the solution to the recursive equation

$$b_{t|t+i} = \max \left\{ (b_{t|t+i+1} + \alpha_{i+1} c_t - Y_{t|t+i+1}^{min}) / r, 0 \right\},$$

with $b_{t|T} = 0$.

If X_t is sufficiently high, then the outcome of step 1 for c_t exceeds $\bar{c}$. In this case, behavior is determined by step 2 of the FG algorithm, which triggers the accumulation goal. The latter specifies how a boundedly rational agent aims for a particular future standard of living in a normal scenario without engaging in contingent planning. Since the accumulation goal is a function of c_t , start by fixing a particular level of this variable. (The optimal level of c_t will be discussed later.) Savings are thus given by $X_t - c_t$. The hypothetical accumulation scheme that determines the accumulation goal entails identifying the highest level of s_t , given $s_t + b_t = X_t - c_t$, that is compatible with meeting the insurance goal.

A crucial insight is that the highest level of s_t that is consistent with achieving the insurance goal is found by assuming that, *from the next period on*, all savings will be

¹³To see this more clearly, consider the case where feasibility of $\alpha_i c_t$ would have to be achieved entirely through savings. In the case of stocks, this would require a level of savings of $\alpha_i c_t / (r^{s \min})^i$. In the case of bonds, it would require savings of $\alpha_i c_t / r^i < \alpha_i c_t / (r^{s \min})^i$

invested in bonds, should it turn out that a worst-case scenario materializes. The reason is straightforward. This will simply minimize the amount of resources required to achieve the insurance goal because $r > r^{s\min}$. Let us again disregard non-negativity constraints for b_t and s_t for a moment. In this case, the agent has to make sure of entering the next period with a level of resources of at least $\sum_{i=1}^{T-t} \alpha_i c_t / r^{i-1} - \sum_{i=1}^{T-t} Y_{t|t+i} / r^{i-1}$. Thus, the highest level of s_t compatible with the insurance goal is then determined by the equation $s_t r^{s\min} + (X_t - c_t - s_t) r = \sum_{i=1}^{T-t} \alpha_i c_t / r^{i-1} - \sum_{i=1}^{T-t} Y_{t|t+i} / r^{i-1}$, where $X_t - c_t - s_t$ is just equal to b_t and $s_t r^{s\min} + (X_t - c_t - s_t) r$ represents the lowest possible return on the agent's portfolio. Clearly, the resulting solution for s_t is only meaningful if it is nonnegative. Otherwise, the optimal value of s_t is zero. An important observation is that if c_t is very high then no nonnegative value of b_t and s_t may fulfill the above equation.¹⁴ According to Definition 2, the accumulation goal would take on a value of zero in this case.

Lemma 1 below states how b_t and s_t are determined taking into account non-negativity constraints for both b_t and s_t .

Lemma 1 *Given a particular level of c_t , asset allocation at age t is determined by the program*

$$\begin{aligned}
\max_{b_t, s_t} \quad & s_t \quad s.t. \\
& b_t r + s_t r^{s\min} \geq \hat{b}_{t|t+1}(c_t) + \alpha_1 c_t - Y_{t|t+1}^{min}, \\
& c_t + b_t + s_t = X_t, \\
& b_t, s_t \geq 0,
\end{aligned} \tag{5}$$

where $\hat{b}_{t|t+1}$ is determined by the recursive equation

$$\hat{b}_{t|t+j} = \max \left\{ \left(\hat{b}_{t|t+j+1} + \alpha_{j+1} c_t - Y_{t+j+1}^{min} \right) / r, 0 \right\} \tag{6}$$

¹⁴To see this, note that in the extreme case where the agent would eat up the entire present value of (minimum) lifetime resources, there would be clearly no way how the agent could ever achieve the insurance goal.

and where $\hat{b}_{t|T} \equiv 0$; whenever a solution to (5) exists. (Otherwise, we have $a_{t|T^*} = 0$.)

Given an optimal choice of c_t (see below), the *actual* portfolio choice in period t is determined by program (5). This portfolio allocation also represents the first step of the hypothetical accumulation scenario that determines the value of the accumulation goal.

We next have to consider portfolio allocation in period $t + 1$, according to the hypothetical accumulation scheme. However, for this and all following periods between t and T^* , the hypothetical portfolio allocation problem is completely isomorphic to the actual one in period t . More precisely, the agent hypothetically projects ending up in period $t + 1$ with a level of resources $X_{t+1}^n = b_t r + s_t E r^s + E_t Y_{t+1}$. If $T^* > t + 1$, i.e. the accumulation goal refers to a future period more than one period ahead, then the hypothetical accumulation process is continued. In $t + 1$, consumption is again set to c_t and savings are thus $X_{t+1}^n - c_t$. The agent hypothetically projects choosing a portfolio that maximizes $s_{t|t+1}$, conditional on meeting the insurance goal. Thus, the portfolio is again determined as stated in Lemma 1, except that indexes have to be interchanged (see program (9) in Proposition 1 below). This yields a hypothetical portfolio $(b_{t|t+1}^n, s_{t|t+1}^n)$ and leads to a level of resources in period $t + 2$ of $X_{t+2}^n = b_{t|t+1}^n r + s_{t|t+1}^n E r^s + E_t Y_{t+2}$. This hypothetical accumulation process is iterated further until period T^* . Note that in all periods $t + i$, $0 \leq i \leq T^* - t - 1$, the hypothetical portfolio choice is such that the insurance goal is achieved even in the case that the normal scenario may not prevail any more from period $t + i + 1$ on (again see program (9) in Proposition 1 below). In period T^* , the agent ends up with a level of resources $X_{t|T^*}^n$. The accumulation goal is then given by $a_{t|T^*} = X_{t|T^*}^n / (T - T^*)$. This is true except in the case where the optimization problem of the form of (5) has no solution for a particular period $t + i$. In this case, we have $a_{t|T^*} = 0$, according to Definition 2.

I next turn to some more technical issues. Let us return once more to Lemma 1 and program (5). Solving the budget constraint for b_t and substituting it into the insurance goal constraint, we get

$$(r - r^{s \min}) s_t \leq X_t r - \hat{b}_{t+1} (c_t) - (\alpha_1 + r) c_t + Y_{t|t+1}^{min}. \quad (7)$$

Suppose that c_t becomes sufficiently large. Clearly, $\hat{b}_{t+1}$ is an increasing function of c_t (see (6)). Thus, the right-hand side of (7) eventually becomes negative. This would imply that $s_t < 0$. (The agent would like to borrow at the rate $r^{s \min}$.) Since this is not feasible, we must have $s_t = 0$ for high levels of c_t , assuming that a solution to (5) still exists.

In the opposite case where c_t is very low it is the constraint $b_t \geq 0$ which becomes binding. To see this assume that $c_t = 0$. Thus, (7) would imply that $s_t \leq (X_t r + Y_{t|t+1}^{\min}) / (r - r^{s \min})$. Suppose that this would hold with equality, which would lead to the unconstrained maximum value of s_t . This would imply that $s_t > X_t$ and hence $b_t < 0$. As a result, the borrowing constraint would be violated. It follows that for low values of c_t the borrowing constraint may be binding.

From (6), $\hat{b}_{t+1}$ is an increasing function of c_t , as already mentioned. Furthermore, $\hat{b}_{t+1}$ is continuous and piecewise linear in c_t . From this and the fact that either the constraint $s_t \geq 0$ or $b_t \geq 0$ may be binding, it follows that both b_t and s_t are also continuous and piecewise linear functions of c_t , provided that a solution of (5) exists. This does not only hold for b_t and s_t , but also for all $b_{t|t+i}^n, s_{t|t+i}^n, i = 1, 2, \dots, T^*$ (see program (9) in Proposition 1 below). As a result, $a_{t|T^*}$ must also be a continuous and piecewise linear function of c_t in the domain of c_t where a solution to the portfolio allocation problems exists. Otherwise, we have $a_{t|T^*} = 0$.

Clearly, $a_{t|T^*}$ must be a decreasing function of c_t . To see this, note first that a higher level of c_t means a lower level of saving in each period from t to T^* . Furthermore, a higher c_t leads to a (weak) increase of $\hat{b}_{t|t+i}, i \geq 1$. In case that neither of the constraints $b_{t|t+i} \geq 0, s_{t|t+i} \geq 0$ are binding, this leads to a higher value of $b_{t|t+i}$ and a lower value of $s_{t|t+i}$. This follows from the fact that the insurance goal constraint is binding in this case and that $r > r^{s \min}$.¹⁵ Overall, we can conclude that $a_{t|T^*}$ is indeed a decreasing function of c_t .

Since $a_{t|T^*}$ is piecewise linear and decreasing in c_t , the objective function (4) in step 2 of the FG algorithm is well defined. In particular, it is strictly convex and the outcome

¹⁵Suppose that the insurance goal would not be binding. In this case the agent could choose a higher level of s_t without violating the insurance goal constraint.

of step 2 is thus unique. Overall, behavior under step 2 is determined according to the following proposition.

Proposition 1

1. Under step 2 of the FG algorithm, c_t is determined by the program

$$\max_{c_t} \ln(c_t - \bar{c}) + \beta^{T^*-t} \ln(a_{t|T^*}(c_t)). \quad (8)$$

Furthermore, $a_{t|T^*} = X_{t|T^*}^n / (T - T^*)$. $X_{t|T^*}^n$ is determined according to the accumulation equation $X_{t|t+i}^n = b_{t|t+i-1}^n r + s_{t|t+i-1}^n E r + E_t Y_{t+i}$, $i = 1, \dots, T^* - t$. The portfolios $(b_{t|t+i}^n, s_{t|t+i}^n)$, $i = 0, 1, \dots, T^* - t - 1$, are determined by the program

$$\begin{aligned} \max_{b_{t|t+i}, s_{t|t+i}} \quad & s_{t|t+i} \quad s. t. \\ & b_{t|t+i} r + s_{t|t+i} r^{s \min} \geq \hat{b}_{t|t+i+1}(c_t) + \alpha_1 c_t - Y_{t|t+i+1}^{n \min}, \\ & c_t + b_{t|t+i} + s_{t|t+i} = X_{t|t+i}^n, \\ & b_{t|t+i}, s_{t|t+i} \geq 0, \end{aligned} \quad (9)$$

where $Y_{t|t+i+1}^{n \min}$ denotes the minimum of Y_{t+i+1} conditional on the normal scenario prevailing until $t + i$, where $X_{t|t}^n \equiv X_t$, and $\hat{b}_{t|t+i+1}$ is defined as in Lemma 1. If no feasible solution to (9) exists for at least one $i = 0, 1, \dots, T^* - t - 1$, then $a_{t|T^*} = 0$ (see Definition 2).

2. The actual portfolio choice in period t is determined according to (9) for $i = 0$.

Although the optimization programs determining behavior in case of the FG model may appear somewhat unusual to economists used to work with the standard model, it should be reemphasized that behavior under this model is conceptually much simpler than in case of the standard model. This is evident in the case that behavior is determined by step 1 of the FG algorithm. Also, in case of step 2, the agent effectively only solves a two-period decision problem in each period, according to (8). He trades off current consumption against feasibility of a certain standard of living from some future period on in case that the normal scenario evolves. The level of this feasible standard of living is determined by a simple hypothetical accumulation process which the agent makes use of to engage

in forward-looking behavior in the absence of contingent planning. The portfolio choice underlying the hypothetical accumulation scheme is determined by maximizing expected returns subject to making a prudent choice that allows the agent to meet the insurance goal.

5 Some Methodological Observations

This section discusses whether the FG model leads to dynamically inconsistent behavior, the role of the model's parameters as well as the differences between the FG and existing models of habit formation.

5.1 Dynamic Consistency or Inconsistency?

An important issue to be discussed is whether behavior under the FG model is dynamically inconsistent. At first sight, the answer seems to be dynamic inconsistency. One may put forward the following argument. According to the agent's decision calculus in period t , the level of current consumption c_t is crucial for how the agent engages in forward-looking behavior. This is most salient in case of the hypothetical accumulation scheme underlying the determination of $a_{t|T^*}$. The agent "projects" to consume c_t in all periods between t and T^* . Furthermore, he aims to achieve the insurance goal in all these periods, which is again defined relative to c_t . However, once period t is passed, c_t does not enter the decision calculus any more.

It is instructive to contrast this with existing models of habit formation, where agents anticipate that c_t will affect their future choices. Suppose that there is a habit state variable H_t with $H_t = \delta c_{t-1} + (1 - \delta) H_{t-1}$, where δ is a parameter that takes on a value between zero and one. Assume, for the sake of the argument, that the objective function is $E_t \sum_{i=0}^{T-t} D^i u(c_{t+i} - H_{t+i})$ for all t , where D represents a discount factor and u is a von Neumann-Morgenstern utility function with the usual properties. Clearly, in this case c_t influences not only decisions in period t , but it has also a direct influence on decisions in later periods $t+i$ via the state variable H_{t+i} . This differs from the logic of the FG model

and suggests that the latter leads to dynamically inconsistent behavior.

This argument is flawed, however, since it confounds two conceptually different sets of variables that influence agents' decision calculus in case of the FG model. First, there is the variable c_t . Trivially, this represents an actually chosen level of consumption in period t . However, all variables referring to the future, such as the habit levels $\alpha_i c_t$ determining the insurance goal or the variables related to the accumulation process determining $a_t|_{T^*}$, are not variables that are linked to actual future choices. Rather, these are *planning variables* belonging to a *set of potential future choices that are feasible, given a particular current decision*. If we treat the FG algorithm as a representation of preferences, then we may say that agents have preferences over both currently realized consumption and over future feasibility sets. Clearly, the latter are a function of c_t . Therefore, c_t plays a dual role. First, it generates direct utility in period t . Second, it affects future feasibility sets. We may not conclude from the fact that c_t is not of direct relevance any more in later periods that decision making is dynamically inconsistent. The reason that c_t does not show up any more in later periods is that feasibility sets in later periods are not a direct function of c_t any more. (They depend on c_t only indirectly via intertemporal budget constraints.) This insight is most essential for understanding the philosophy of the FG model. The conclusion is that the standard notion of dynamic consistency is not meaningful in case of the FG model. Since there are no contingent plans ex ante, such plans cannot be violated ex post.

5.2 The Role of the Model's Parameters

The role of the parameters α_i , $i = 1, \dots, T - t$ should have become sufficiently transparent from Section 4. Increasing the level of α_i means increasing the minimal level of consumption that an agent wants to be feasible in a future worst-case scenario i periods ahead in the future. Thus, this entails a higher level of risk aversion and prudence. This will generally increase savings and decrease equity shares. The role of the parameter β becomes clear from (8). A higher β entails a higher weight of the accumulation goal. This

increases the propensity to save as well as the propensity to engage in risky savings.

The parameter whose role may be least clear at first sight is $\bar{c}$. This parameter plays two roles. First, it determines the critical level of X_t that is required to trigger step 2 of the FG algorithm. Second, $\bar{c} > 0$ entails saving rates and equity shares that increase with X_t , given that step 2 is activated. To see this, consider the case where X_t is such that it is only slightly higher than the critical level that triggers step 2 of the FG algorithm. In this case, the chosen level of c_t is close to $\bar{c}$. As a result, the marginal propensity to consume is very high. When X_t increases, c_t will exceed $\bar{c}$ by more and more. As a result, the marginal propensity to consume decreases and the desire to achieve a high level of the accumulation goal becomes stronger. This, in turn, increases the propensity to save and to invest in high-risk/high-return assets, i.e. stocks.

The parameter $\bar{c}$ is not linked in any particular way to the concept of feasibility goals. The model is well defined even for $\bar{c} = 0$. The aim of this parameter is to take into account that an agent may compare his consumption level to a reference point that reflects a normal standard of living in his social environment.

The role of the three model parameters becomes much clearer if the model is solved analytically for certain benchmark cases. Such an analysis is carried out in Binswanger (2008) and the interested reader is referred to that companion paper.

5.3 A Comparison to Existing Models of Habit Formation

Since, on the face of it, the FG model resembles existing models of habit formation, it is in order to clearly state the main difference between both models. The most crucial difference is that any existing habit formation model requires contingent planning. The essential element of the FG model is that agents engage in coarse planning and are concerned about future feasibility sets as discussed in Section 5.1. In particular, the role of the “habit levels” $\alpha_i c_t$ is to define future feasibility sets that agents want to achieve. Given that agents are assumed not to anticipate any actual future decisions, any future feasibility set can only be defined or marked by referring to current decision variables. Habit levels

such as $\alpha_i c_t$ provide a particularly simple way of specifying future feasibility sets that are useful for determining precautionary savings. Put differently, in the FG model the habit levels $\alpha_i c_t$ provide a “planning device” for boundedly rational agents.

6 Calibrating the Model

For the calibration of the model, I consider a setup where an agent starts working at age 21, retires at the beginning of age 66, and dies with certainty after age 85. The model is simulated at an annual frequency. The calibration analysis is restricted to working life since the FG model as presented here does not take into account mortality risk nor bequest motives. As a result, it is less suited to explaining behavior during retirement.

The parameters of the FG model are the sequence $\{\alpha_i\}$, β , $\bar{c}$, and also 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.) In this sense, the specification of the sequence $\{\alpha_i\}$ captures a particular element of discounting.

The baseline value for α^* is set to .96, which is chosen as follows. Binswanger (2010) provides a cross-sectional analysis of savings and asset allocation based on a static model corresponding to the FG model in the degenerate case where there are only two periods. Thus, 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 (2010) this static 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 .27.¹⁶ In order to impose discipline

¹⁶This number should not be interpreted as implying that the consumption level that individuals

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 (2010) amounts to 27,000 year-2001 U.S. dollars. The parameter β is set to .9.¹⁷

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. It represents thus an agnostic specification. (See Section 9 for an alternative specification.)

Non-financial income is calibrated following standard assumptions that are outlined in the Appendix. In the absence of shocks, non-financial income would follow the hump-shaped profile in Figure 1. There are permanent and transitory multiplicative shocks to non-financial income. The expected value of both shocks is one and the minimum realization is assumed to be strictly positive (see Appendix). There is also a medical expenditure shock M that is realized upon retirement and is fully permanent. This means that the resources that are available for normal consumption are reduced by a factor $1 - M$. The expected value of M is one and its lower bound $\underline{M}$ is set to 60 percent in case of the baseline specification.

Real bond returns are assumed to be risk-free and to equal 2 percent per year. Stock returns are assumed to be *iid* with an expected annual return of 6 percent and a strictly positive minimum return (again see the Appendix).

would like to have at all costs during retirement would be only .27 of their working life consumption. See Binswanger (2010).

¹⁷The 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 Binswanger, 2008). 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 annual periods allows agents to rebalance their portfolio every year. It is well known that increases in flexibility increase the willingness to take risks 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 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.

In the interest of transparency, all simulations are run for the case where, ex post, non-financial income as well as stock returns always take on their expected values.

7 Empirical Facts and Predictions of Existing Models

Before turning to the calibration results it is useful to have a look at the empirical stylized facts regarding the variation of consumption/saving and asset allocation choices over the life cycle. These are contrasted with the predictions of existing life cycle models.

7.1 Empirical Stylized Facts

In a recent study, Fernández and Krueger (2007) explore how consumption expenditures vary over the life cycle. 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.

Concerning asset allocation behavior, it is instructive to look at how median and mean equity shares¹⁸ vary over age for data referring to a single year. In this way, the slope of the resulting age profile cannot be distorted by what is known as time effects. Table 1 contains median and mean equity shares for different age groups from the 2004-wave of the Survey of Consumer Finances (SCF). These equity shares contain direct as well as indirect equity holdings through mutual funds and retirement accounts. In order to get insight into the behavior of a typical person within a given age group it is particularly meaningful to look at the median numbers in the second column of Table 1. The figures clearly demonstrate that early in their career agents do typically not participate in the

¹⁸Equity shares are defined as the ratio of total equity holdings to total financial wealth.

stock market, which translates into an equity share of zero. Stock market entry occurs only after age 35. Overall, the profile is increasing except for the last age category. The maximum equity share amounts to roughly 40 percent during the second half of the fifties. The same pattern prevails for means, except that they are always strictly positive.^{19 20}

Table 2 presents evidence on the variation of equity shares over *income*, again based on the SCF 2004. The table reports median equity shares for particular age/income cells. 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 three upper quintiles do always participate. There is also rather strong evidence that equity shares increase with income among stock market participants. Note again that this pattern cannot be distorted by time effects since all data come from the SCF 2004. Furthermore, cohort effects may lead to a downward distortion in the slope of the equity share profile within a given column. However, they are unlikely to affect the profile within a given row.

7.2 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 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.

¹⁹One explanation for the pattern in Table 1 may be that younger workers would simply have less access to employer-sponsored pension accounts. However, it does not seem that this is responsible for this pattern (see Ameriks and Zeldes, 2004).

²⁰The slope of the true equity share profile may be understated by the figures in Table 1 if older cohorts are more skeptical towards equity investments than younger ones. It is also important to note that the joint level of the figures in Table 1 (not their value relative to each other) may be influenced by the particular year in which they have been observed. However, the respective numbers from the SCF 2001 (not shown) are very similar.

Such models successfully explain that consumption increases with age up to some peak level.²¹ Regarding 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 (10) in the Appendix, the prediction is that stock market participation is universal. In addition, equity shares are predicted to be much higher than observed in the data. Specifically, equity shares are predicted to be close to 100 percent 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 pattern discussed above.

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 arise 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.

8 Calibration Results

This section presents the results from simulating the FG model for a baseline specification. Results for several alternative specifications are presented in Section 9.

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

²¹However, the peak occurs too 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).

during retirement.

According to Figure 2, consumption tracks income very closely until the age of 35. 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. The only motive for saving is to achieve the insurance goal.

At the age of 27 the standard of living $\bar{c}$ is reached. From then on optimal decisions are determined by step 2 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 goal 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)²² 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 small again.

Figure 3 presents the predicted profile for equity shares.²³ Equity shares are defined as the ratio of total stock holdings to total financial assets, i.e. $s_t/(b_t + s_t)$. As already

²²See their Figure 3(A).

²³The raw outcomes of the simulations for equity shares are slightly wiggling. The equity shares in Figure 3 are smoothed using a moving average smoothing algorithm.

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. Individuals start to participate in the stock market 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. In particular, the effects of a sequence of low realizations of the permanent income shock accumulate over time (see equation (11) in the Appendix). Hence, the uncertainty about future income and its downside risk increases with the number of periods that lie between a current period t and retirement. In this sense, a 25-year old faces more future income risk than a 45-year old. More things can still go wrong for the 25-year old since he has more opportunities to pick a long chain of low realizations of the permanent income shock. It is precisely this effect 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.

After the age of 45 the medical expenditure shock $\underline{M}$ comes into play. 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^{R-t} is very low. When agents get closer to retirement, α^{R-t} increases. This makes agents more prudent and stops equity shares from increasing further.

The predictions of the FG model come fairly close to the empirical asset allocation patterns shown in Table 1. This is especially true in comparison to the standard model. For the standard model, equity shares amount to 100 percent until about the age of 40, where they start to decrease gradually to 50 percent, which is reached upon retirement (see Section 7).

Figure 4 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 4 that the FG model is consistent with 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. If income is 50 percent higher compared to the benchmark case, equity shares are higher as well. 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, as discussed in Section 5.2. It is noteworthy that standard workhorse models such as CRRA or Epstein-Zin preferences predict no variation of saving and asset allocation choices with permanent income since these preferences are homothetic.

The observation that the FG model is much better able than existing models to explain the empirical data on asset allocation raises the question whether the FG 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. However, what is crucial in this respect is that existing preference models have great difficulties in explaining the data even when the flexibility of the setup is greatly increased, e.g. by assuming Epstein-Zin preferences, habit formation, heterogeneous preferences, by including fixed costs of stock market participation, endogenous borrowing constraints, or by assuming a positive probability 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 high degree of flexibility. Rather, it is likely to stem from the fact that the logic of feasibility goals is intuitively compelling to boundedly rational agents. According to the model, the young may find that they face a high degree of career uncertainty. Thus, they

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

engage in the stock market only once this career uncertainty has sufficiently been reduced. This logic does not correspond at all to the standard model where equity shares are very high until the age of about 40. The reason is the following. It turns out that the prospect of facing an increasing labor income profile with a high probability provides a substitute for risk-free savings in case of the standard model. Hence, the young 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 people don't seem to take the logic of the standard model as an intuitive guideline for their investment behavior.

9 Alternative Specifications

This supplement presents calibrations for some alternative parameter values. In each case, only one parameter deviates from the baseline version at a time.

Figures 5 and 6 refer to a situation with a higher medical expenditure risk. In particular, $\underline{M}$ amounts to .5 (instead of .6). In this case, consumption is lower and therefore savings are higher than in the baseline case. Furthermore, equity shares decrease after the age of 45. Figures 7 and 8 show the consumption and equity share profiles for the case where medical expenditure risk is lower than in the baseline case. In particular, $\underline{M}$ amounts to .75. In this case, the consumption profile is shifted upwards. Furthermore, equity shares continue to increase after age 45.

Figures 9 and 10 refer to the case where $\bar{c} = 0$. This is certainly the most parsimonious specification of the model. In this case, only step 2 of the FG algorithm is relevant. Consumption is slightly lower than in the baseline case up to the age of about 50 and higher thereafter. The reason is that 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.

Figures 11 and 12 refer to the case where $\beta = .94$. This is the best-fit value for β obtained in Binswanger (2010) for the cross-sectional setting considered there (see Section

6 in the main paper). Naturally, consumption is slightly higher in this case at later ages since agents accumulate more wealth.

Figures 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 family size.²⁵ Specifically, I assume that the baseline value of 27,000 year-2001 U.S. dollars refers to two adults and one child. Using equivalence scales, a profile for $\bar{c}$ is calculated for the family history described 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.

10 Conclusion

This paper provides a new life cycle framework in which agents do not engage in contingent planning but pursue simple so-called feasibility goals. Such a feasibility goals framework represents a particularly natural way how to deviate from the standard rational benchmark

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

model.

The paper shows that pursuing such feasibility 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 case of the standard model. In particular, the 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.

Appendix: Calibration of Income Process and Financial Returns

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

$$Y_t = F_t P_t V_t, \quad (10)$$

for $T_0 \leq t < R$, where

$$P_t = P_{t-1} U_t. \quad (11)$$

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 $V_t \in [\underline{V}, \bar{V}]$ and $U_t \in [\underline{U}, \bar{U}]$. 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 (10), (11), as well as for the standard deviations of the logs of V_t and U_t .²⁶ 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. Expected values keep a value of one.²⁷ The resulting values for $\underline{U}$ and $\underline{V}$ 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

$$Y_t = \rho Y_{R-1} M_t, \quad (12)$$

²⁶See their Tables 1, 2, 3.

²⁷Note that the FG algorithm requires only specification of the minimum and expected values of all random variables involved.

The parameter ρ represents a “replacement rate” and is set to .6. $M_t \in [\underline{M}, \overline{M}]$ 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 $1 \leq i \leq T - R$. Furthermore, it is assumed that $E_t M_R = 1$ for $T_0 \leq t \leq R - 1$. 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.

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 .157.²⁸ 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 the original lognormal distribution, while keeping expected returns at 6 percent.²⁹ The gross minimum return amounts then to .74 per year.

²⁸These assumptions are fairly common. See e.g. Campbell and Viceira (2002).

²⁹Thus, 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.

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 (2007), "Risk Management of Pensions from the Perspective of Loss Aversion," *Journal of Public Economics*, 91, 641-667.
- Binswanger, Johannes (2008) "Life Cycle Saving: Insights from the Analytical Solution of a New Bounded-Rationality Model," Working Paper.
- Binswanger, Johannes (2010) "Understanding the Heterogeneity of Savings and Asset Allocation: A Behavioral-Economics Perspective," *Journal of Economic Behavior and Organization*, 76, 296-317.
- Binswanger, Johannes (forthcoming) "Dynamic Decision Making with Feasibility Goals: A Procedural-Rationality Approach," *Journal of Economic Behavior and Organization*, forthcoming.
- 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.
- Clark, Andrew E., Frijters, Paul, and Shields, Michael A. (2008), "Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles," *Journal of Economic Literature*, pp. 95-144.
- Cocco, João F., 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*, MacMillan, 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.

Table 1: Equity shares over the life cycle (in percent)

Age	Equity share	
	Median	Mean
<30	0	14.3
30-34	0	22.3
35-39	7.9	25.9
40-44	22.5	30.7
45-49	25.9	33.3
50-54	30.0	34.4
55-59	39.5	40.0
60-64	34.0	36.1

Source: Survey of Consumer Finances 2004.

Table 2: Median equity shares for age-income cells (in percent)

Age	Income quintiles				
	First	Second	Third	Fourth	Fifth
<30	0	0	17.6	37.9	34.6
30-34	0	0	26.0	36.5	48.0
35-39	0	0	32.8	40.7	56.9
40-44	0	0	22.9	46.5	53.3
45-49	0	0	17.6	46.7	52.9
50-54	0	0.1	8.4	47.3	51.0
55-59	0	17.0	33.2	43.3	57.5
60-64	0	13.0	24.7	44.8	50.0

Source: Survey of Consumer Finances 2004.

Figure 1: Non-financial income in the absence of shocks

Figure 2: Consumption and income

Figure 3: Equity shares

Figure 4: Equity shares for various income levels

Figure 5: Consumption and income, higher medical expenditure uncertainty

Figure 6: Equity shares, higher medical expenditure uncertainty

Figure 7: Consumption and income, lower medical expenditure uncertainty

Figure 8: Equity shares, lower medical expenditure uncertainty

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, $T^* = 2/3 t + 1/3 T$