Tilburg University

A Simple Bounded-Rationality Life Cycle Model

Binswanger, J.

Publication date:
2008

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

Citation for published version (APA):

Binswanger, J. (2008). *A Simple Bounded-Rationality Life Cycle Model*. (CentER Discussion Paper; Vol. 2008-13). Macroeconomics.

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. 2008–13

A SIMPLE BOUNDED-RATIONALITY LIFE CYCLE MODEL

By Johannes Binswanger

February 2008

ISSN 0924-7815

A Simple Bounded-Rationality Life Cycle Model*

Johannes Binswanger[†]

January 10, 2008

Abstract

Life cycle saving decisions belong to the most complex financial decisions that we are faced with in our life. Psychologists have found that when making complex decisions people use short-cuts in the form of minimum requirements for particular attribute categories of choice options. This paper presents a new simple life cycle model where agents do invoke such minimum requirements. The model is highly tractable and parsimonious. Calibrations show that it allows us to better understand important data on saving and asset allocation. It is shown that the model is much better able to explain these data than standard workhorse models even when generously controlling for subtle differences in the “degrees of freedom” between the new and existing models.

Key words: Asset allocation, behavioral economics, bounded rationality, life cycle saving, noncompensatory decision making, threshold goals.

JEL classification: D81, D91.

*I am very grateful to Roland Bénabou, Antoine Bommier, Michele Boldrin, Martin Browning, Peter Diamond, Josef Falkinger, Ernst Fehr, Reto Föllmi, Guido Friebel, Simon Gächter, Christian Gollier, David Laibson, George Loewenstein, Manuel Oechslin, James Poterba, Panu Poutvaara, Eytan Sheshinski, Daniel Schunk, Jean Tirole, Andreas Wagener, Richard Zeckhauser, Josef Zweimüller and seminar participants at GREMAQ, the University of Mannheim, the University of St. Gallen, the Stockholm School of Economics, the Stockholm University, and at Tilburg University for helpful comments, and to Karen Dynan for kindly providing information about PSID income quintiles. I am responsible for all remaining errors.

[†]CentER, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands; j.binswanger@uvt.nl.

1 Introduction

Some economic decisions, such as daily shopping, may be seen as truly simple. Others may be less simple, but often more difficult economic decisions are made frequently under similar circumstances. As a result, individuals develop a considerable amount of expertise through experience. In contrast, some economic decisions are truly difficult ones and are not made on a routinely basis. A prime example is life cycle saving and retirement preparation, which are the topic of this paper. It is hard to deny that such decisions are particularly complex, even with the help of financial experts.¹ Furthermore, such decisions are made relatively infrequently. And, while you get more or less immediate feedback on whether you made a suitable decision when having bought a new car or moved into a new apartment, there is no comparable feedback for life cycle saving decisions. Typically, you may get feedback only decades later. Clearly, this hinders learning.

The question of how people deal with complex decisions has been the subject of a large number of studies in psychology. Decision making usually entails choosing one out of several choice alternatives. Each alternative is characterized by a number of attributes. Psychologists have found that when facing complex decision tasks, people come up with minimum thresholds (deliberatively or rather “intuitively”) for at least some of the attribute categories. These minimum thresholds serve as criteria for making choices in that a decision maker disregards any alternative for which these minimum criteria are not met. This procedure greatly reduces the number of relevant alternatives and hence simplifies decision making considerably.² Such a procedure differs greatly from decision patterns under standard expected utility maximization. In the language of microeconomics, min-

¹Readers who believe that consulting expert advice would make life cycle saving a piece of cake are referred to www.consumerreports.org/cro/money. In particular, the summary of a report on financial advice reads as: “We also found that paying thousands of dollars to an independent financial professional doesn’t guarantee perfection either. The bottom line: You can’t expect to hand off your finances to someone and put your retirement on autopilot no matter how much you pay. As you’ll see from the experiences of three testers [...], financial planning is a time-consuming and sometimes nerve-wracking process. You can’t avoid spending hours on preparation, which should include checking on the adviser’s background and doing some soul-searching to figure out how to balance your current financial needs and desires with your retirement goals.” (See <http://www.consumerreports.org/cro/money/personal-investing/bargain-financial-advisers-206/overview/index.htm>.)

²See Section 2 for a brief summary of the psychological literature.

imum thresholds lead to discontinuous choice modes. In the language of psychologists, such decision making is called noncompensatory. The reason for the latter terminology is that if an alternative entails attribute values below a minimum threshold in some category, it is not possible to compensate individuals for this shortfall by sufficiently high values of attributes in other categories.

Since life cycle saving and retirement preparation are especially complex economic decisions, this psychological finding seems particularly relevant for understanding how people deal with these decisions. This paper will provide a new simple and parsimonious model of life cycle saving where decision making is characterized by the use of minimum thresholds which will be dubbed *threshold goals*.

The paper demonstrates that a very simple model, incorporating such threshold goals into a preference framework, has great power to explain cross-sectional patterns of saving and asset allocation choices. In particular, the threshold goal model is much better able to explain the cross-sectional variation of saving and asset allocation choices with income than standard workhorse models. This holds even when generously controlling for subtle differences in the degrees of freedom or “flexibility” between the threshold goal and alternative models.

It is a critical issue to identify the appropriate list of threshold goals that are likely to be relevant in the domain of life cycle saving. The strategy I pursue is to look for threshold goal elements that appear in the existing literature and are likely to be relevant for life cycle savings. A first threshold goal element that has commonly been incorporated into economic models of intertemporal decision making and asset pricing is habit formation (Campbell and Cochrane, 1999, Constantinides, 1990, Gomes and Michaelides, 2003). A habit level of consumption can also be seen as a reference point (Bowman et al., 1999). In the model of this paper it will be assumed that decision makers disregard all saving or asset allocation plans for which it is not guaranteed (or extremely likely) that future consumption exceeds a habit level. The difference between the new threshold goal model and existing habit formation models is discussed below.

The model includes a second threshold goal referring to current consumption (and, indirectly via habit formation, also to future consumption). Given habit formation, it

is assumed that the an agent desires to come as close as feasible to a certain “normal” standard of living. The idea is to capture, in a stylized manner, the finding in the literature that people seem to be concerned about how their standard of living compares to a reference consumption level (Easterlin (1974), (1995); see Clark et al. (forthcoming) for a recent survey). This reference level can be understood as the average standard of living in a reference group or in the society as a whole. It is also likely to correspond to generation-specific aspiration levels (Clark et al., forthcoming). The existence of such reference consumption levels explains why, for developed countries, the “average” happiness within a country increases only very little with economic growth.

Overall, the threshold goal model works as follows. Agents’ overall or residual goal is simply to maximize a weighted average of “utility” from current and *expected* future consumption. Such a risk-neutral overall objective function provides the simplest possible way to capture that agents are making a trade-off between current and future consumption and, in principle, like to have more of both. However, agents only consider to choose among decision alternatives for which the threshold goals are met. In particular, agents reject all alternative for which the habit level is not achieved. Among the remaining choices they reject all alternatives which do not come as close as feasible to achieving the normal standard of living.

There are two main differences to existing habit formation models. First, such models do not include the threshold goal of a normal standard of living. The latter is crucial for the main predictions of the threshold goal model, namely that both saving rates and the proportion of savings invested in equities increase with income, in a way that is consistent with empirical estimates but not easily obtainable from existing models.³ Second, the overall objective function in the threshold goal model is much simpler than for standard habit formation models since there is no element of risk aversion, in the sense of a concave utility function over risky prospects. Not only does this make the model particularly tractable but it also leads to much simpler patterns of behavior. Habit formation is the only source of risk aversion in the threshold goal model.

³See Sections 7 and 8.

To focus on the simplest and most parsimonious case, the paper considers a two-period life cycle model. In this model all decisions are made in the first period. In the second period the agent simply consumes whatever resources are available, depending on the state of the world that will materialize. Restricting the analysis to the two-period case allows to abstract from the intricate issue of how decision makers deal with the so-called “continuation problem,” i.e. whether and how they take into account today what decisions they will make at future knots of the decision tree. An investigation how boundedly rational decision makers may deal with this issue is beyond the scope of this paper and is addressed in Binswanger (2007a). Second, in a two-period framework it is particularly straightforward to compare calibrations of the new threshold goal model to calibrations of existing models.

The idea that economic decision making is heavily influenced by threshold goals goes at least back to Simon (1955). However, to my knowledge, this idea has not formally been applied to the study of life cycle decision making, so far. Existing behavioral life cycle models include the mental accounting model of Shefrin and Thaler (1988), the hyperbolic discounting model (e.g. Laibson, 1997, Laibson et al., 1998), and the loss aversion model of Bowman et al. (1999). These models do not include any noncompensatory elements of decision making. Gabaix and Laibson (2000, 2005) analyze abstract decision tasks where agents are to choose an initial “action” which triggers a chain of probabilistic consequences. Gabaix and Laibson develop models of how boundedly rational agents may process information on the probabilistic consequences of a given choice by using simple algorithms. In contrast, the threshold goal model is not a model of information processing. Rather, it describes how individuals are making choices by using threshold goals that refer to attributes of their choice alternatives.

As already mentioned, Binswanger (2007a) presents a multi-period model which, in the static case of two periods, coincides with the threshold goal model of this paper. Binswanger (2007a) is concerned with how boundedly rational agents may deal with the challenging problem that their satisfaction today may depend on choices in the more distant future, which choices may depend on choices made in intermediate periods between now and the more distant future etc. The idea put forward there is that agents may have

simple feasibility goals for future consumption along a few scenarios, such as a worst-case and a “normal” scenario. The idea of feasibility goals is intrinsically related to a multi-period framework.

The rest of the paper is organized as follows. Section 2 briefly summarizes the relevant psychological literature. Section 3 introduces the threshold goal model formally. Section 4 presents the setup regarding the economic environment. Section 5 provides the analytical solution of the model. Section 6 discusses how the model can be calibrated. Section 7 presents calibration results for a set of baseline parameter values as well as a host of alternative specifications. Section 8 compares the predictive power of the model to other existing models that are commonly used to study life cycle savings. Section 9 concludes.

2 A Summary of the Psychological Literature

In psychology there has been a large amount of experimental research investigating whether people are using noncompensatory decision rules and whether they are more prone to do so when a decision task becomes more complex. In a typical experiment people are asked to choose, say, an apartment, given information about several attribute categories such as size, distance from working place etc. The experiments are designed in a way such that the experimenter can observe the information that a subject has been using when making her decision.⁴ It is the pattern of information processing that reveals whether a decision maker has been using compensatory or noncompensatory decision strategies. In order to make a decision in the compensatory mode a decision maker has to evaluate a larger amount of information. In particular, she has to check all attribute values of a given option, since she may always be willing to accept being compensated for a low value in one category by a sufficiently high value in another. In contrast, in the noncompensatory mode a decision maker will start by checking the values of the attributes for the categories for which she has a minimum requirement in mind. This will lead to an outright rejection of many alternatives, such that no further information is investigated for these alternatives.

⁴The two methods used are verbal protocol analysis and information boards (see Ford et al., 1989).

Ford et al. (1989) provide a large survey summarizing 45 studies that investigate decision modes. They put their conclusion as follows (p. 105).⁵

Two substantive conclusions are derive from the review of the empirical evidence. First, from an overall perspective, the results firmly demonstrate the extensive use of noncompensatory strategies in completing the problem task. In most studies, the nonlinear [= noncompensatory] use of cues was the dominant mode of decision making identified by the researchers. Compensatory strategies [...] were used when the number of alternatives and dimensions were small or after a number of task stimuli were eliminated from consideration. [...]

Focusing more specifically on the determinants of strategy use, a second substantive conclusion is that the most robust findings associated with the set of studies reviewed is that task complexity is strongly related to strategy use. The research indicates that as the number of dimensions and alternatives increase, there is a greater likelihood that the decision maker will use nonlinear strategies. This finding has been found with different types of samples (students, consumers, professional staffs), different types of tasks (housing decisions, consumer product decisions, evaluations of others), and across both verbal protocol analysis and information board approaches to [decision] process tracing.

The quote highlights that the observation that people are switching to noncompensatory decision modes when making complex choices has the status of a stylized fact. So far, there has been very little attempt to incorporate this stylized fact into analyses of economic decision making.

3 The Threshold Goal Model

Formally, the threshold goal model can be understood as a preference model where preferences are represented by the list of threshold goals, plus a residual goal that determines

⁵Interested readers are also referred to Beach and Mitchell (1978), Bröder (2003), Einhorn (1971), Newell and Simon (1972), Payne (1976), and Payne et al. (1993).

behavior overall when all threshold goals are achieved. For simplicity, I will refer to the entries in this list just as “goals.” Individual behavior is determined by maximization of a goal i in the list, subject to achieving the maximum of all goals with an index smaller than i , and subject to budget constraints. The particular goal i is determined as the goal with the highest index, such that achieving the maximum of all goals with a lower index is feasible. Generally speaking, the higher an individual’s income, the higher is the number of goals for which the maximum can be achieved.

Specifically, LLA preferences are represented by the list

$$\Lambda = \{I(\tilde{c}_2^{min} \geq \alpha c_1), \min[c_1, \bar{c}], E\tilde{c}_2(c_1 - \bar{c})^\gamma\}, \quad (1)$$

containing three goals. The first entry in the list specifies the threshold goal of achieving a habit level of consumption in all future circumstances. In particular, the decision maker will never choose a plan for which it is not assured that at least a fraction $\alpha > 0$ of first-period consumption c_1 can be consumed in the second period. $\tilde{c}_2^{min}$ denotes minimum second-period consumption. I denotes the indicator function. It takes on a value of one whenever $\tilde{c}_2^{min} \geq \alpha c_1$, and of zero otherwise. The first goal is thus maximized if $\tilde{c}_2^{min} \geq \alpha c_1$.⁶ Clearly, achieving the first goal in (1) is only feasible if there exists a financial asset with strictly positive returns in all states of the world. In the following it will be assumed that such an asset exists, which provides a useful simplification. $\tilde{c}_2^{min}$ should not be understood too literally. It is rather intended to represent a level of future consumption that can be achieved with a very high probability, e.g. 99 percent.

Consider next the second goal. It takes on a value of c_1 or $\bar{c}$, whichever is smaller. It expresses the desire to achieve a normal standard of living $\bar{c} \geq 0$ and represents a further threshold goal. For simplicity, $\bar{c}$ is treated as exogenous. Combining the first and the second goal in (1), it follows that an individual maximizes c_1 s.t. $\tilde{c}_2^{min} \geq \alpha c_1$ and the budget constraints, as long as income is sufficiently low such that the resulting optimal choice of c_1 does not exceed the threshold goal $\bar{c}$.

⁶It should be noted that, if (1) is literally understood as a preference model, an increase in c_1 may make an individual worse off, holding everything else constant. This is an intrinsic property of habit formation models.

If income is sufficiently high, such that fully achieving the first two goals is feasible, behavior is determined by the third goal, which represents the residual goal. It specifies nonsatiation in c_1 and the expected level of second-period consumption $E\tilde{c}_2$ (E stands for the mathematical expectation operator), as well as a complementarity between c_1 and $E\tilde{c}_2$. The degree of complementarity is determined by the parameter $\gamma > 0$. Beyond this, the third goal entails a desire for “speculative” savings in the sense of a preference for a high value of Ec_2 . Subtracting $\bar{c}$ from c_1 assures that c_1 represents a strictly normal good at all income levels. Note that the marginal rate of substitution between c_1 and $E\tilde{c}_2$ is always well defined for $c_1 > \bar{c}$. It is only when this inequality holds that the third goal is actually relevant. It should be emphasized that the desire to reach a habit threshold provides the only source of risk aversion in the model since the third goal entails risk neutrality. This provides the simplest way to capture a basic risk-return trade-off.

Inspection of (1) shows that c_1 and $\tilde{c}_2$ enter preferences in a somewhat asymmetric way. First, the evaluation of the random consumption stream $\tilde{c}_2$ is split up into the separate evaluation of $\tilde{c}_2^{min}$ and of $E\tilde{c}_2$. In contrast, no such splitting occurs for c_1 because it is deterministic. Second, $\tilde{c}_2$ is evaluated relative to the habit level αc_1 . In contrast, c_1 is not evaluated against any initial value of a habit stock variable. Rather, c_1 is evaluated relative to the reference point $\bar{c}$. Note that $\tilde{c}_2$ is implicitly also evaluated against a reference point $\alpha\bar{c}$. The reason is that the first goal in (1) implies that an expansion of c_1 towards $\bar{c}$ only has value if at the same time $\tilde{c}_2^{min}$ is expanded towards $\alpha\bar{c}$. In this sense, a more symmetric model specification could be achieved by specifying the third goal as $(E\tilde{c}_2 - \alpha\bar{c})(c_1 - \bar{c})^\gamma$. This would not lead to any qualitatively different predictions. The mentioned asymmetries all stem from an attempt to keep the model as simple and tractable as possible.⁷

A question that naturally may come up is whether there is any compelling logic behind the order of the goals in (1). Clearly, it must be that the two threshold goals precede the non-saturable overall goal of maximizing $E\tilde{c}_2(c_1 - \bar{c})^\gamma$. Given this, suppose that the order of the first and second threshold goal would be reversed. This would mean that a decision

⁷Beyond that, the asymmetries might nevertheless raise worries about the possibility to extend preferences (1) to a multi-period setting. It is shown in Binswanger (2007a) that a natural such extension with an isomorphic parameterization does indeed exist.

maker would first want to achieve a normal standard of living during working age before she would care in any way at all about her future standard of living. Thus, choices would be fully myopic for a certain range of lower incomes, which seems rather odd. Overall, I view the order of arguments as stated in (1) as the only plausible one.

4 The Economic Environment

The analysis focuses on a stylized setup where the first period corresponds to working life and the second period to retirement. Individuals earn a positive exogenous income in the first and active period of their life. In contrast, income is zero in the second period of life. Thus, the model does not specify Social Security income or other sources of non-private retirement income. The reason for this is that the model will later be calibrated and compared to empirical estimates of saving rates that take into account non-private sources of retirement income (see Section 7).

The stylized time setup of the model is best understood as follows. Suppose that an agent starts working at age W , retires at age R and dies with certainty at age D . One may now conceive that in each year during working life an agent is preparing for $(D - R) / (R - W)$ years during retirement. More precisely, the household may be thought to prepare during the year where she is aged W for her old age consumption at ages from R to $R + (D - R) / (R - W)$; and during the year where she is aged $W + 1$ she may prepare for her old age consumption at ages from $R + (D - R) / (R - W)$ to $R + 2(D - R) / (R - W)$. The average time horizon for this retirement preparation amounts then to $(D - W) / 2$ years. In this case, the household solves $R - W$ identical life cycle problems, which have the form of the problem analyzed in this paper.

Individuals face two decisions. They have to decide how much to save and how to allocate their savings between stocks and bonds. The returns of both financial assets are exogenously given. For simplicity, housing is omitted, as in a large body of existing work on life cycle saving.⁸

⁸See e.g. Carroll (1997), Campbell and Viceira (2002), Cocco et al. (2005).

Denote bond holdings by b , stock holdings by s , and first-period income by w . Furthermore, denote *gross* return of bonds by $\tilde{x}$. Bond returns are possibly risky, as indicated by the tilde. Gross stock returns are denoted by $\tilde{y}$. I will make the following assumption about return distributions.

Assumption 1 (i) $\tilde{x} \in [\underline{x}, \bar{x}]$, (ii) $\tilde{y} \in [\underline{y}, \bar{y}]$, (iii) $0 \leq \underline{y} < \underline{x}$ and $E\tilde{y} > E\tilde{x}$, (iv) $\Pr[\tilde{x} = \underline{x}, \tilde{y} = \underline{y}] > 0$.

Part (i) and (ii) specify the support of bond and stock returns. Part (iii) states that stocks have a larger downside risk than bonds. But the larger risk is compensated by a higher expected return. Finally, by (iv) it is assumed that there is a strictly positive probability that both assets realize their minimum returns. This rules out that investing in stocks allows for increasing the minimum return of a portfolio with respect to a pure bond portfolio.

Budget constraints are given by

$$c_1 + b + s = w, \quad (2)$$

$$\tilde{c}_2 = bx + s\tilde{y}. \quad (3)$$

Furthermore, I assume that

$$b \geq 0, s \geq 0 \quad (4)$$

must hold, i.e. there is no borrowing and short selling. This is again a usual assumption in the literature (see e.g. Cocco et al., 2005).

5 An Analytical Solution of the Model

Optimal choices have to be analyzed separately for low and high income levels. Specifically, fully achieving the first two goals in (1) will only be feasible if w is sufficiently high. (Fully achieving the first goal is always feasible since $\underline{x} > 0$, which follows from Assumption 1.) Denote the threshold value of w for which the first two goals are just

fully achievable by w^{crit} . For $w \leq w^{crit}$ an individual solves the program

$$\begin{aligned}
& \max_{c_1, b, s} c_1 \quad s.t. \\
& \underline{b}\underline{x} + s\underline{y} \geq \alpha c_1, \\
& c_1 + b + s = w \\
& c_1 \geq 0, \quad b \geq 0, \quad s \geq 0,
\end{aligned} \tag{5}$$

where (3) and Assumption 1 have been used for the first constraint.

For $w > w^{crit}$ an individual maximizes the third goal in (1), subject to fully achieving the first two goals. As any monotonic transformation of an objective function does not change the outcome of an optimization problem, this program may be stated as

$$\begin{aligned}
& \max_{c_1, b, s} \gamma \ln(c_1 - \bar{c}) + \ln(bE\tilde{x} + sE\tilde{y}) \quad s.t. \\
& \underline{b}\underline{x} + s\underline{y} \geq \alpha c_1, \\
& c_1 \geq \bar{c}, \\
& c_1 + b + s = w \\
& c_1 \geq 0, \quad b \geq 0, \quad s \geq 0.
\end{aligned} \tag{6}$$

The rest of this section deals with the solution to the two programs (5) and (6), respectively. Let us begin with problem (5). For this program, optimal stock holdings are zero. Suppose that stock holdings were positive, instead, and consider whether an individual can do better by substituting all stock holdings by bonds. Since $\underline{x} > \underline{y}$, habit consumption can then be maintained with less resources. Hence, the individual can afford a higher level of c_1 without jeopardizing the habit goal, thus realizing a higher value of the objective function in (5). It follows from this (and the budget and non-negativity constraints) that $c_1 = \frac{\underline{x}}{\alpha + \underline{x}}w$ and $b = \frac{\alpha}{\alpha + \underline{x}}w$. It remains to determine the range of incomes for which program (5) applies. This will be the case if and only if the resulting choice of c_1 does not exceed $\bar{c}$, i.e. if and only if $\frac{\underline{x}}{\alpha + \underline{x}}w \leq \bar{c}$. This is equivalent to $w \leq \frac{\alpha + \underline{x}}{\underline{x}}\bar{c}$. This yields the following result.

Proposition 1 *If $w \leq w^{crit} \equiv \frac{\alpha+\underline{x}}{\underline{x}}\bar{c}$, then $c_1 = \frac{\underline{x}}{\alpha+\underline{x}}w$, $b = \frac{\alpha}{\alpha+\underline{x}}w$, $s = 0$.*

Proof. See the text above. ■

Proposition 1 is of interest since it provides a very simple solution to the so-called stock market participation puzzle. Empirically, a substantial fraction of households do not participate in the stock market in any country of the world. In particular, households with a relatively lower income are much less likely to participate in the stock market (see Section 7). In contrast, expected utility models predict that all households should participate in the stock market if the equity premium is positive, at least in the absence of fixed costs of stock market participation (see e.g. Gollier, 2001). Typically, predicted equity shares are quite substantial for reasonable calibrations of expected utility as well as more general models (Cocco et al. 2005, Gomes and Michaelides, 2005, see also Section 8). The stock market participation puzzle refers to the fact that this is so much at variance with the data. According to Proposition 1, the threshold goal model is able to explain stock market non-participation for lower levels of income (as well as high levels of γ .)

According to the threshold goal model, earners of a lower income save exclusively for precautionary reasons. Their only saving motive is to guarantee the threshold goal of a habit consumption level, as this is a more basic goal than the desire to savor a high expected old age consumption. Evidently, maintaining habit consumption is achieved most efficiently by only investing in the asset with the lowest downside risk, i.e. bonds. This explanation of the stock market participation puzzle is certainly not very challenging from an intellectual point of view (nor is the usual explanation of the puzzle by means of fixed costs). However, this does not preclude that threshold goals do play an important role in deciding whether or not to enter the stock market.

Although the threshold goal model is close in spirit to existing habit formation models, the latter are not able to explain the stock market participation puzzle (Gomes and Michaelides, 2003). What is crucial for the predictive success of the threshold goal model with respect to the participation puzzle is the strict sequential priority order of goals. Individuals for which the second goal of an average standard of living is not yet achieved are not at all concerned about the third goal. In contrast, there is only one holistic

objective function for existing habit formation models. For these, all agents share the same *full* range of “goals,” so to speak. As a result, there are always both a precautionary and a speculative motive of saving.

I next characterize the solution to problem (6) which is relevant for $w > w^{crit}$. There are three possible types of solutions. First, in some cases it may be optimal to choose $b > 0$, $s = 0$. Second, it may occur that $b > 0$, $s > 0$ or, finally, $b = 0$, $s > 0$. It will turn out that all three cases can occur, depending on the values of w and/or γ . For the statement of the results it is convenient to define $\Omega \equiv (\alpha + \underline{x})E\tilde{y} - (\alpha + \underline{y})E\tilde{x}$. Furthermore, $\bar{\gamma} \equiv \frac{\underline{x}\Omega}{\alpha(\underline{x}-\underline{y})E\tilde{x}}$ and $\bar{w} \equiv \frac{(\alpha+\underline{x})\Omega}{\underline{x}\Omega - \alpha\gamma(\underline{x}-\underline{y})E\tilde{x}}\bar{c}$.

Proposition 2 *Assume $w > w^{crit}$. If $w \leq \bar{w}$ or $\gamma \geq \bar{\gamma}$, then $c_1 = \frac{\underline{x}}{\alpha+\underline{x}}w$, $b = \frac{\alpha}{\alpha+\underline{x}}w$, $s = 0$.*

Proof. See Appendix. ■

Note that $\bar{w} > w^{crit}$ for $\gamma < \bar{\gamma}$. To see this, verify that we would have $\bar{w} = w^{crit}$ for $\gamma = 0$. But $\bar{w}$ is a strictly increasing function of γ for $\gamma < \bar{\gamma}$. This implies the strict inequality. Thus, Proposition 2 shows that surpassing the threshold income w^{crit} will not immediately lead to positive equity investments. The reason is that, locally around $\bar{c}$, the marginal value of an increase in c_1 is close to infinity. In contrast, the marginal value of an increase in $E\tilde{c}_2$ is finite, since the validity of the first constraint of program (6) implies that $E\tilde{c}_2$ sufficiently exceeds zero. This feature of the model assures that c_1 is always a strictly normal good. For similar reasons, we have $s = 0$ when γ is sufficiently high. Then, the relative value of a marginal increase in $E\tilde{c}_2$ is low at all levels of c_1 .

I turn next to the case where $b > 0$ and $s > 0$. Define $\bar{\bar{\gamma}} \equiv \frac{\underline{y}\Omega}{\alpha(\underline{x}-\underline{y})E\tilde{y}}$ and $\bar{\bar{w}} \equiv \frac{(\alpha+\underline{y})\Omega}{\underline{y}\Omega - \alpha\gamma(\underline{x}-\underline{y})E\tilde{y}}\bar{c}$. By direct calculations it can be verified that $\bar{\bar{\gamma}} < \bar{\gamma}$ and $\bar{\bar{w}} > \bar{w}$ for $\gamma < \bar{\gamma}$ (see Figure 1). This allows for stating the next result.

Proposition 3 *If $\bar{\gamma} \leq \gamma < \bar{\gamma}$ and $w > \bar{w}$, or if $\gamma < \bar{\gamma}$ and $\bar{w} < w < \bar{\bar{w}}$, then*

$$\begin{aligned} c_1 &= \frac{1}{1+\gamma} \bar{c} + \frac{\gamma}{(1+\gamma)\Omega} (\underline{x}E\tilde{y} - \underline{y}E\tilde{x}) w, \\ b &= \frac{1}{(1+\gamma)(\underline{x} - \underline{y})\Omega} \{ [\alpha\gamma(\underline{x} - \underline{y})E\tilde{y} - \underline{y}\Omega] w + (\alpha + \underline{y})\Omega\bar{c} \}, \\ s &= \frac{1}{(1+\gamma)(\underline{x} - \underline{y})\Omega} \{ [\underline{x}\Omega - \alpha\gamma(\underline{x} - \underline{y})E\tilde{x}] w - (\alpha + \underline{x})\Omega\bar{c} \}. \end{aligned}$$

Proof. See Appendix A. ■

Note that $\gamma > \bar{\gamma}$ implies that bond investments increase with income. Moreover, the expression for b is strictly positive whenever $\gamma \geq \bar{\gamma}$ or when $\gamma < \bar{\gamma}$ and $w < \bar{w}$. The inequality $\gamma < \bar{\gamma}$ implies that the coefficient of w in the formula for s is strictly positive. $w > \bar{w}$ implies that the overall expression for s is strictly positive.

Proposition 3 shows that individuals invest in both assets, stocks and bonds, whenever γ lies in an intermediate range and income is sufficiently high. To understand this note that when γ is very high then the relative preference weight of $E\tilde{c}_2$ is so low that individuals will never want to invest in stocks (see Proposition 2). Conversely, when γ is very low, then the marginal value of an increase in c_1 becomes very low whenever it sufficiently exceeds $\bar{c}$. For such low values of γ individuals will spend a large share of their income on $E\tilde{c}_2$ and thus on s . As a consequence, c_1 is rather low and habit consumption can eventually be met entirely by equity savings. For intermediate values of γ and incomes that allow financing a level of c_1 which sufficiently exceeds $\bar{c}$, the marginal values of c_1 and $E\tilde{c}_2$ are sufficiently balanced. Hence, s is not so high, and c_1 not so low, that habit consumption can be met entirely by stock investments. Thus, individuals invest in both stocks and bonds. A second case to which Proposition 3 applies is a low γ and an income level which is not too high. Here, too, the marginal values of increases in c_1 and $E\tilde{c}_2$ are sufficiently balanced such that individuals invest in both assets.

With regard to empirical data (see Section 7), it is an important implication of Proposition 3 that both the saving rate $(b + s)/w$ and the equity share $s/(b + s)$ are predicted to increase with income under the conditions of Proposition 3.⁹ The intuition for these

⁹For the saving rate this follows from the fact that the ratio c/w is decreasing in w . For the equity

results follows from the fact that seeking to achieve a high level of $E\tilde{c}_2$ in the third goal in 1 becomes more and more important with a higher income. This stems from the fact that a higher income allows to finance a higher level of c_1 , which moves further and further apart from $\bar{c}$. As a result, the marginal value of an increase of c_1 becomes lower and lower with respect to the marginal value of an increase in $E\tilde{c}_2$. This leads to a higher optimal value of both savings and the proportion of savings invested in equities. Note that the assumption of $\bar{c} > 0$ is vital for this result. In the case of $\bar{c} = 0$ saving rates as well as equity shares would be constant.

I turn now to the final case where $b = 0$ and $s > 0$. It is convenient to define $\bar{\bar{w}} \equiv \frac{\alpha+y}{y-\alpha\gamma}\bar{c}$. It can be verified that $\bar{\bar{w}} > \bar{w}$ for $0 < \gamma < \frac{y}{\alpha}$. Furthermore, it can be checked that $\bar{\gamma} > \frac{y}{\alpha}$ (again, see Figure 1). The last proposition, dealing with the case where the relative marginal value of an increase in c_1 is low, reads as follows.

Proposition 4 *If $\gamma < \bar{\gamma}$ and $w \geq \bar{w}$, then $b = 0$. Moreover,*
(i) if $\frac{y}{\alpha} \leq \gamma < \bar{\gamma}$ and $w \geq \bar{w}$, or if $\gamma < \frac{y}{\alpha}$ and $\bar{w} \leq w \leq \bar{\bar{w}}$, then

$$\begin{aligned} c_1 &= \frac{\underline{y}}{\alpha + \underline{y}}w, \\ s &= \frac{\alpha}{\alpha + \underline{y}}w. \end{aligned}$$

(ii) If $\gamma < \frac{y}{\alpha}$ and $w > \bar{\bar{w}}$, then

$$\begin{aligned} c_1 &= \frac{\gamma}{1 + \gamma}w + \frac{1}{1 + \gamma}\bar{c}, \\ s &= \frac{1}{1 + \gamma}(w - \bar{c}). \end{aligned}$$

Proof. See Appendix A. ■

It is straightforward to show that under the conditions of Proposition 4(i), the first constraint in (6) is binding, whereas this is not the case under the conditions of (ii). Figure 1 integrates the results of Propositions 1 to 4 into a coherent picture, depicting the different portfolio regimes in the γ/w -space. It should be noted that the different

share this can be shown by calculating the derivative $d[s/(b+s)]/dw$.

threshold values for w , i.e. $\bar{w}$, $\bar{\bar{w}}$ and $\bar{\bar{\bar{w}}}$, can all be understood as hyperbolic functions of γ . In particular, they exhibit jump discontinuities at $\bar{\gamma}$, $\bar{\bar{\gamma}}$ and $\frac{y}{\alpha}$, respectively. Moreover, they are strictly positive and increasing to the left of their jump discontinuities and negative (and thus not of any interest) to their right. Note also that for $\gamma = 0$ we have $\bar{\bar{w}} = \bar{\bar{\bar{w}}} > \bar{w} = w^{crit}$. Furthermore, as mentioned previously, $\frac{y}{\alpha} < \bar{\bar{\gamma}} < \bar{\gamma}$, $\bar{w} < \bar{\bar{w}}$ for $\gamma < \bar{\bar{\gamma}}$, and $\bar{\bar{w}} < \bar{\bar{\bar{w}}}$ for $0 < \gamma < \frac{y}{\alpha}$.

6 Calibrating the Model

This section discusses how the threshold model is calibrated, such that it is possible to obtain quantitative predictions for saving rates and equity shares. These will then be compared to empirical data for the U.S. Moreover, Section 8 will compare the calibration results to the predictions of two often-used existing preference models.

With respect to the time setup of the model, I assume that an agent starts his life cycle saving program at the age of 26, enters retirement at the age of 66 and dies with certainty at the age of 86. This setup takes into account that retirement typically lasts for fewer years than working life. The setup abstracts from mortality risk. This may be seen as appropriate in an environment where agents have access to an explicit or implicit form of annuities. Using the argument set out in Section 4, the average time horizon for an agent's retirement preparation problem is then 30 years.

The model will first be calibrated for a set of baseline parameter values that are chosen as being plausible on a priori grounds. A straightforward idea for the determination of α is to look at U.S. Social Security replacement rates. The latter range between 30 percent for high income levels and 60 percent for low income levels (Munnell and Soto, 2005). The underlying reason for this monotonic decline in income is that households with a higher income level are supposed to engage in private retirement preparation. Thus, Social Security replacement rates for higher income levels should be seen as lower than minimally acceptable. In contrast, the Social Security replacement rates for low income levels are set in such a way that they lead to a minimally acceptable standard of living during retirement in the absence of private retirement preparation. Thus, 60

percent seems an a priori natural benchmark value for α . However, this value needs to be adjusted to the time setup of the model. Implicitly, Social Security replacement rates compare expenditure levels of two periods with equal length, e.g. of two different years such as one during working life and one during retirement. However, the two periods in the model are not of equal length. To correct for this, the raw replacement rate of 60 percent has to be multiplied by the ratio of the length of retirement to the length of working life, i.e. by one half, according to the above assumptions. This leads to a baseline value for α of 0.30.

As $\bar{c}$ represents a normal standard of living, it would be natural to identify it with median consumption of the total U.S. population. There are, however, no good data available. Thus, I set $\bar{c}$ to the upper limit of the second income quintile for the year 2001, amounting to 33 thousand year-2001 U.S. dollars.¹⁰ ¹¹ The parameter γ can be interpreted as the inverse of a time discount factor. It is thus set to the inverse of 0.96 to the power of 30, as 0.96 represents a standard specification for an annualized discount factor.

Turning to the economic environment, it is necessary to numerically specify the distribution of bond and stock returns. In this respect, I adopt standard assumptions from the literature.¹² Specifically, annual gross returns of bonds are set to 1.02, assuming that they are risk-free. Annual net returns are thus equal to 2 percent. The distribution of stock returns is derived as follows. It is assumed that the mean and standard deviation of annual gross returns are equal to 1.06 and 0.157, respectively. I assume further that annual log-returns are distributed normally with parameters corresponding to a level mean and standard deviation of 1.06 and 0.157, respectively. This allows for a determination of the mean and standard deviation of 30-year level returns, assuming that stock returns are identically and independently distributed over time. I proceed by drawing 1,000,000 draws from the 30-year distribution, yielding a “data set” of stock returns. The minimum return is then set to the first percentile of the 30-year distribution, while the expected

¹⁰Source: U.S. Census Bureau, Historical Income Tables. See <http://www.census.gov/hhes/www/income/histinc/h01ar.html>

¹¹The year 2001 has been chosen since model predictions will be compared to empirical equity shares obtained from the 2001 wave of the Survey of Consumer Finances.

¹²See e.g. Campbell and Viceira (2002).

return is set to 1.06³⁰.

7 Calibration Results for the Threshold Goal Model

Figures 2 and 3 present the results for the baseline calibrations for saving rates and equity shares, respectively. The solid lines refer to model predictions while the dashed lines correspond to empirical estimates. Empirical saving rates are taken from Dynan et al. (2004) and represent separate estimates of saving rates for each income quintile.¹³ It is important to stress that these savings rates include prospective income from Social Security, which is converted into a notional saving rate and added to observed private savings rates. Furthermore, these saving rates include private pensions.¹⁴ Such all-inclusive saving rates provide the natural empirical benchmark for evaluating quantitative predictions that are obtained in a setup that abstracts from private and public pensions.

Equity shares are obtained from the 2001 wave of the Survey of Consumer Finances (SCF). I define brackets with a width of 20,000 year-2001 U.S. dollars that are symmetric around the median quintile income levels that refer to the five Dynan et al. estimates of saving rates. For each bracket I calculate the median equity share, defined as the median of the ratio of equity holdings to total financial assets. For the five percent of observations with no financial assets the equity share is set to zero. Equity holdings include indirect holdings through mutual funds and pensions.¹⁵

¹³The estimates for saving rates refer to ages between 40 and 49. The differences to savings rates for the 30-39 as well as for the 50-59 age group are quantitatively small (see Table 3 in Dynan et al., 2004).

¹⁴Specifically, saving rates are taken from the “Active+Pension” column in Table 3 of Dynan et al. (2004). I am grateful to Karen Dynan for kindly providing information about median incomes for the Panel Study of Income Dynamics.

¹⁵The dashed line in Figure 3 is based on all SCF observations with age below 65, which corresponds to the fact that the model aims to explain equity allocation choices during working life. It should be mentioned that the shape of the empirical equity curve is fairly similar for narrower age classes, such as ages from 30 to 39, 40 to 49 etc. Furthermore, similar curves are obtained for different SCF education classes, except for the lowest. Finally, it should be mentioned that the shape of the empirical curve in Figure 3 cannot be biased by time effects, since all data refer to the year 2001. Time effects can only influence the position of the curve. Whether cohort effects may affect the shape of the empirical equity curve is a more subtle issue. They do not seem a prime reason of concern in the context of this study, however, since, as already mentioned, the empirical equity curve exhibits a similar shape for different age classes.

As shown in Figures 2 and 3, both empirical saving rates and equity shares increase substantially with income. The threshold goal model is well able to explain the general shape of these curves. In particular, the model is able to explain non-participation in the stock market for the lowest income quintile. This follows from the fact that only the first two goals in the list (1) are active for the first income quintile under the baseline parameter values. It should again be stressed that the model explains non-participation in the stock market at low income levels in the absence of fixed costs of market participation. (See Section 5 for a discussion of this result.) In contrast, it is a well known fact that expected utility models generically predict stock market participation at all income levels in the absence of fixed costs of market participation. Moreover, predicted equity shares are quit high for low income levels for standard parameter values (see next Section).

As discussed in Section 5, the intuition for increasing saving rates and equity shares follows from the fact that seeking to achieve a high level of $E\tilde{c}_2$ in the third goal in (1) becomes more and more important with a higher income. This stems from the fact that a higher income allows to finance a higher level of c_1 which moves further and further apart from $\bar{c}$. As a result, the marginal value of an increase of c_1 becomes lower and lower with respect to the marginal value of an increase in $E\tilde{c}_2$. Remember from Section 5 that $\bar{c} > 0$ is crucial for the prediction of increasing saving rates and equity shares. For $\bar{c} = 0$ both curves would be flat.

In the next section the predictions of the threshold goal model will be compared to those of alternative existing models. It is therefore useful to introduce a formal measure of how closely a model matches the empirical estimates. I will focus on the square root of the mean squared error, averaged over both saving rates and equity shares, as a natural loss function. Define empirical saving rates for quintile i by $\hat{\sigma}_i$ and predicted saving rates by σ_{ij}^* . Similarly, define empirical and predicted equity shares by $\hat{\rho}_i$ and ρ_{ij}^* , respectively. The measure of fit for model j is then defined by

$$FIT_j = \left[0.5 \sum_{i=1}^5 (\sigma_{ij}^* - \hat{\sigma}_i)^2 + 0.5 \sum_{i=1}^5 (\rho_{ij}^* - \hat{\rho}_i)^2 \right]^{0.5}.$$

For the baseline specification of the threshold goal model the FIT value amounts to

0.1006.

In the rest of this section I will discuss calibration results for some alternative specifications. I first present three specifications where one parameter differs from the baseline value at a time. This allows for clearer insights into the “mechanics” of the model. At the end I will present a specification which minimizes the *FIT* value.

Figures 4 and 5 show simulations for $\alpha = 0.35$ instead of 0.30. According to the logic of Social Security replacement rates this would correspond to a replacement rate of 70 percent (see Section 6). Since the habit level αc_1 is higher in this case precautionary savings are higher. This shifts saving rates upwards and equity shares downwards with respect to the baseline case. The *FIT* value amounts to 0.0725.

Figures 6 and 7 present results for $\bar{c} = 25$ instead of 33. In this case c_1 departs sooner from $\bar{c}$. As a result, maximization of $E\tilde{c}_2$ is relatively more important in the third goal of (1). Consistent with this logic, saving rates and equity shares are higher than in the baseline case. The *FIT* value amounts to 0.1924.

Figures 8 and 9 present saving rates and equity shares for the case where γ is set to 0.94^{-30} instead of 0.96^{-30} . This leads to both lower savings and lower equity shares since c_1 has a higher weight in the third goal in (1). As a consequence, $E\tilde{c}_2$ has a lower relative weight. The *FIT* value amounts to 0.0702 in this case.

Finally, Figures 10 and 11 correspond to a set of parameter values that are chosen such as to minimize the *FIT* value. The minimization has been carried out over the parameter space $\mathcal{A} \times \mathcal{C} \times \Gamma$, where the three sets correspond to values for α , $\bar{c}$ and γ , respectively, and $\mathcal{A} = \{0.25, 0.26, \dots, 0.40\}$, $\mathcal{C} = \{15, 16, \dots, 35\}$, $\Gamma = \{0.80^{-30}, 0.81^{-30}, \dots, 0.98^{-30}\}$. The resulting best fit values are $\alpha = 0.27$, $\bar{c} = 29$, $\gamma = 0.94^{-30}$. The value for α of 0.27 corresponds to a pension income replacement rate of 54 percent according to the logic of Social Security replacement rates. It should be stressed that these parameter values are fairly close to the baseline parameters and do not strike as unnatural, if only the value of α is somewhat low.¹⁶

Figures 10 and 11 show that the fit is practically perfect for equity shares and very

¹⁶See Frederick et al. (2002) for a list of empirical estimates for discount rates. A annual value of 0.94 for the discount rate does not seem low at all in light of their discussion.

good for saving rates. The formal *FIT* value is 0.0226. This fit is so favorable that it naturally invokes some skepticism. A first issue that comes to mind is that any model with three free parameters may lead to a very favorable fit, if the fit is optimized over parameter values. Second, it may be the case that the threshold goal model has some hidden “degrees of freedom” in that the third preference goal in (1) is only triggered if income is sufficiently high. Therefore, saving rates and equity shares for lower income quintiles could be matched independently of saving and asset allocation choices for higher income quintiles.

An immediate objection against the last argument is that it is only the first quintile for which the third goal in (1) is not activated under the best-fit parameter values. However, the fit of the model is also very favorable for just the upper four income quintiles where always all goals are active. Nevertheless, the question of whether the fit obtained in Figures 10 and 11 is a true achievement of the threshold goal model or whether it may be obtained for any model with similar degrees of freedom can only be addressed by explicitly engaging in a similar exercise for alternative preference models. This is the topic of the next section. There, I will present simulations for constant relative risk aversion preferences as well as hyperbolic absolute risk aversion preferences. The conclusion will be that the threshold goal model performs much better than alternative existing models, even when generously controlling for possible differences in the degrees of freedom to the benefit of existing models.

8 Saving and Asset Allocation under CRRA and HARA preferences

The two-period constant relative risk aversion (CRRA) model is given by

$$U(c_1, \tilde{c}_2) = \frac{1}{1-\eta} c_1^{1-\eta} + \frac{\beta}{1-\eta} E[\tilde{c}_2^{1-\eta}].$$

The parameter η represents the coefficient of relative risk aversion and β represents the time discount factor. E stands for the mathematical expectation operator. It is assumed

that these preferences are maximized subject to the budget constraints $c_1 + b + s + kI(s > 0) = w$ and $\tilde{c}_2 = bx + s\tilde{y}$, as well as subject to $b \geq 0$, $s \geq 0$. I represents the indicator function and k represents a fixed costs of stock market participation. It is a well-known fact that in the absence of such fixed costs stock market participation always occurs for any expected utility preferences. There are thus three parameters: η , β and k . Figures 12 and 13 present simulations for the CRRA model for a set of parameters values that many economists would view as baseline values. Specifically, η is set to 3 and β to 0.96³⁰. Furthermore, k is set to a quarter of a monthly income. More precisely, if w_i denotes the median annual income for a particular income quintile i , then k is set to $0.25w_i/12$ for income quintile i . In my view this represents an upper bound of plausible values for fixed costs of stock market participation.

Figures 12 and 13 show that the fit obtained for this parameter specification is mediocre, at best. This is particularly worrisome in light of the fact that the CRRA model, also in its two-period version, is heavily used for developing basic insights about life cycle saving and retirement preparation. Both saving rates and equity shares are predicted to be constant. Saving rates are excessively high. Furthermore, although the fixed entry cost of stock market participation is high, it is not sufficient for preventing stock market entry for the first income quintile.

I turn next to a calibration that optimizes the fit for the CRRA model. Since this model, as stated above, has three parameters, its degrees of freedom can be seen as comparable to the threshold goal model with its three parameters α , $\bar{c}$ and γ . However, it can be argued that for the latter preferences differ between lower and higher income levels, since the third goal in (1) may not be active for low levels of income. Thus, referring to the fact that there are only three parameters one may understate the model's true flexibility in matching empirical data. Therefore, it may not be fair to directly compare the fit of the threshold goal to the normal version of the CRRA model.¹⁷ I consider therefore a

¹⁷Remember, however, that for all but the lowest income quintile all three goals of the threshold goal model are active for almost all specifications discussed in the last section. Thus, for the upper four quintiles the number of degrees of freedom for the threshold goal model is well captured by the number of parameters. The fit of the threshold goal model is much better than for any parameter values of the baseline CRRA model even when the comparison is restricted to the upper four quintiles. The reason is that the predicted saving rates and equity shares are always flat for the CRRA model.

version of the CRRA model where η is allowed to take on a different value for *each* income quintile. Furthermore, the values of k may differ between income quintiles. In contrast, β is constrained to be identical for all income levels. Specifically, for each income quintile η may take on one of the values of the set $\{1, 2, \dots, 15\}$, while k may amount to either one quarter or one tenth of a month's income. The discount factor β may take on any of the values $\{0.80, 0.81, \dots, 0.98\}$ to the power of 30.

The best fit is obtained for a sequence $(1, 8, 4, 3, 2)$ of η -values over income quintiles. Furthermore, the best fit value for k is equal to one quarter of monthly income for the first quintile and may take on either value for the other quintiles. The optimal value for β amounts to 0.88^{30} . Figures 14 and 15 present the corresponding saving rates and equity shares. In light of the high degree of flexibility that has been invoked for optimizing the fit for the CRRA model the result is lackluster. The best-fit model is able to match equity shares quite well. However, this is only achieved thanks to the large jumps in the values for η . In contrast, the model is barely able to match empirical saving rates. The *FIT* value amounts to 0.0732, compared to 0.0226 for the threshold goal model.

I turn next to the hyperbolic absolute risk aversion (HARA) model. Specifically, I consider a basic version of the model where CRRA preferences are augmented by age-specific subsistence consumption levels, represented by the two parameters $\bar{c}_1, \bar{c}_2$ in

$$U(c_1, \tilde{c}_2) = \frac{1}{1-\eta} (c_1 - \bar{c}_1)^{1-\eta} + \frac{\beta}{1-\eta} E[(\tilde{c}_2 - \bar{c}_2)^{1-\eta}] .$$

It can be shown that non-zero values for $\bar{c}_1, \bar{c}_2$ lead to non-flat saving rate and equity share profiles.¹⁸ Since there are no well-established baseline values for these two parameters I directly consider a calibration that optimizes the fit to the empirical data. Since the effects of letting η vary between income levels has already been explored above, I consider now the effects of allowing $\bar{c}_1, \bar{c}_2$ to differ between income quintiles. Specifically, they may each take on a value between 0 and 15 and need not be equal to each other. The fixed cost k is again either equal to one tenth or one quarter of monthly income and may also differ between income levels. The discount factor β is restricted to be identical for all

¹⁸The profiles for saving rates can be either increasing or decreasing.

income levels and may take on any value in the set $\{0.80, 0.81, \dots, 0.98\}$ to the power of 30. Finally, η is simply set to 3. Note that this model has more degrees of freedom than the version of the CRRA model examined above, since three parameters may now vary with income.

The sequence of best-fit values for $\bar{c}_1$ and $\bar{c}_2$ are given by $(12, 15, 2, 1, 1)$ and $(0, 11, 2, 0, 0)$. The best fit value for k is a quarter of monthly income for the first quintile and one tenth otherwise. The best fit value for β amounts to 0.90³⁰. The simulation results are shown in Figures 16 and 17. The *FIT* value amounts to 0.0633, compared to 0.0226 for the threshold goal model.¹⁹

Evidently, the versions of the CRRA and HARA model considered above have by far more degrees freedom to match the data than the threshold goal model. Even when equipped with such an enormous degree of flexibility their ability to explain the cross-section of saving and asset allocation choices is quite weak. In light of the above it seems fair to conclude that the favorable fit of the threshold goal model, shown in Figures 10 and 11, is due to an intrinsic ability of this model to capture some of the important determinants of individuals' behavior. This fit cannot merely be explained by the flexibility of the threshold goal model offered by the fact that not all goals are activated at low income levels.²⁰

Some readers may worry that the favorable performance of the threshold goal model is limited to savings and asset allocation data and, within this domain, to two-period models. Concerning the first issue, I would not easily accept to perceive this as a problem even if this was the case, since saving and asset allocations represent important economic data that are highly relevant for individual welfare. Concerning the second issue it is

¹⁹At first sight, it may seem that simple existing models of habit formation would do a better job than HARA models explaining the cross-section of saving and asset allocation choices. From a mathematical point of view this can, however, not be the case since, for given predictions of any (simple two-period) habit formation model, it is always possible to find a HARA model with appropriate parameter values such that its predictions are identical to the given habit formation model.

²⁰Another existing model that is sometimes used for studying life cycle choices is the Epstein-Zin model. This model has not been explored here since it shares the property with the CRRA model that predicted savings and equity share profiles are flat. Thus, non-flat profiles can again only be obtained if one is willing to assume that preferences parameters differ between different income groups, which I view as unsatisfactory.

shown in Binswanger (2007a) that a multi-period extension of the threshold goal model also performs better than existing models.

9 Conclusion

This paper has introduced a new model for studying life cycle saving and asset allocation decisions. The defining property of this model is that it incorporates short-cuts that simplify these complex decisions. These short-cuts come in the form of so-called threshold goals.

It has been demonstrated that the new model explains the cross-section of empirical saving and asset allocation choices much better than existing standard models such as the constant relative risk aversion (CRRA) or the hyperbolic absolute risk aversion (HARA) model. Beyond this, the new model has some further advantages. First, as shown in Binswanger (2007a), it can be extended to a multi-period model. The multi-period generalization allows to explain further asset holding phenomena that are puzzling from the perspective of existing models. These are related to how stock market participation and equity shares vary with age. These phenomena can be explained for the same parameter values that have been shown to provide a good explanation of cross-sectional savings and asset allocation data in this paper. Second, the threshold goal model is very tractable. Third, it is highly transparent and therefore allows to gain some basic insights about life cycle saving and retirement preparation in a very straightforward way. These insights allow, for instance, to draw conclusions about how to compose a pension system out of a pay-as-you-go and funded components (Binswanger 2007b).

A major aim of this paper is to convince economists that simple alternatives to existing expected utility models do exist that incorporate realistic features of real-world decision making. It has been shown by means of an example that such models can be parsimonious and at the same time useful for understanding important empirical data. This seems thus a promising route for future research.

Appendix A: Proofs

Consider problem (6). Remember that for $w = w^{crit}$ we have exactly $c_1 = \bar{c}$ from Proposition 1. It follows that problem (6) is well defined if and only if $w > w^{crit}$. Note that the second constraint in (6), and thus also the nonnegativity constraint for c_1 , do not have to be explicitly taken into account. The reason is that at $\bar{c}$ the marginal value of an increase in c_1 is infinite. This guarantees that we will always have $c_1 > \bar{c}$. Using this and the budget constraints, the Lagrangian for program (6) can be written as

$$\mathcal{L} = \gamma \ln(w - b - s - \bar{c}) + \ln(bE\tilde{x} + sE\tilde{y}) + \lambda[b\underline{x} + s\underline{y} - \alpha(w - b - s)] + \mu b + \nu s, \quad (7)$$

where λ , μ and ν are multipliers. The first order Kuhn-Tucker conditions for this program are

$$\frac{E\tilde{x}}{bE\tilde{x} + sE\tilde{y}} - \frac{\gamma}{w - b - s - \bar{c}} + (\alpha + \underline{x})\lambda + \mu = 0, \quad (8)$$

$$\frac{E\tilde{y}}{bE\tilde{x} + sE\tilde{y}} - \frac{\gamma}{w - b - s - \bar{c}} + (\alpha + \underline{y})\lambda + \nu = 0, \quad (9)$$

$$\lambda \geq 0, \quad b\underline{x} + s\underline{y} - \alpha(w - b - s) \geq 0, \quad \lambda [b\underline{x} + s\underline{y} - \alpha(w - b - s)] = 0, \quad (10)$$

$$\mu \geq 0, \quad b \geq 0, \quad \mu b = 0, \quad (11)$$

$$\nu \geq 0, \quad s \geq 0, \quad \nu s = 0. \quad (12)$$

It can easily be checked that the objective function is strictly concave in b and s . Moreover, all constraints are linear. Thus, not only are the Kuhn-Tucker conditions (8)-(12) necessary and sufficient for an optimum, but optimal choices are also unique.

The proofs below will make use of the following definitions.

$$\begin{aligned}
\Omega &\equiv (\alpha + \underline{x})E\tilde{y} - (\alpha + \underline{y})E\tilde{x}, \\
\bar{\gamma} &\equiv \frac{\underline{x}\Omega}{\alpha(\underline{x} - \underline{y})E\tilde{x}}, \\
\bar{\bar{\gamma}} &\equiv \frac{\underline{y}\Omega}{\alpha(\underline{x} - \underline{y})E\tilde{y}}, \\
\bar{w} &\equiv \frac{(\alpha + \underline{x})\Omega}{\underline{x}\Omega - \alpha\gamma(\underline{x} - \underline{y})E\tilde{x}}\bar{c}, \\
\bar{\bar{w}} &\equiv \frac{(\alpha + \underline{y})\Omega}{\underline{y}\Omega - \alpha\gamma(\underline{x} - \underline{y})E\tilde{y}}\bar{c}, \\
\bar{\bar{\bar{w}}} &\equiv \frac{\alpha + \underline{y}}{\underline{y} - \alpha\gamma}\bar{c}.
\end{aligned}$$

Proof of Proposition 2

It has to be checked that the first order conditions (8)-(12) hold for some appropriate values of λ , μ , ν . In particular, consider

$$\begin{aligned}
\lambda &= \frac{1}{\alpha + \underline{x}} \left[\frac{\gamma}{c_1 - \bar{c}} - \frac{1}{b} \right], \\
\mu &= 0, \\
\nu &= \frac{1}{\alpha + \underline{x}} \left[(\underline{x} - \underline{y}) \frac{\gamma}{c_1 - \bar{c}} - \frac{\Omega}{bE\tilde{x}} \right].
\end{aligned}$$

It is straightforward to show that (8) and (9) hold for these multipliers (use $w - b - s = c_1$ and $s = 0$). Inserting for c_1 , b , s in the expression for λ and rearranging shows that $\lambda \geq 0$ if and only if

$$(\underline{x} - \alpha\gamma)w \leq (\alpha + \underline{x})\bar{c}. \quad (13)$$

Either $\gamma < \frac{\underline{x}}{\alpha}$ or $\gamma \geq \frac{\underline{x}}{\alpha}$. Consider the first case. Then (13) holds if $w \leq \frac{\alpha + \underline{x}}{\underline{x} - \alpha\gamma}\bar{c}$. Check that $\frac{\underline{x}}{\alpha} < \bar{\gamma}$. Thus, the conditions of Proposition 2 only hold if $w \leq \bar{w}$. Check further that $\bar{w} < \frac{\alpha + \underline{x}}{\underline{x} - \alpha\gamma}\bar{c}$. Thus, $w < \frac{\alpha + \underline{x}}{\underline{x} - \alpha\gamma}\bar{c}$ holds under the conditions of Proposition 2 for $\gamma < \frac{\underline{x}}{\alpha}$. In the case where $\gamma \geq \frac{\underline{x}}{\alpha}$, (13) always holds since the expression on the right-hand side

is strictly positive. From this and from the fact that $b\underline{x} = \alpha c_1$ and $s = 0$, it follows that (10) holds.

It remains to check (11) and (12). The first condition is obvious. Consider thus the latter. Clearly, $s \geq 0$ and $\nu s = 0$. Inserting the expressions for c_1 , b stated in Proposition 2 into the expression for ν above shows that $\nu \geq 0$ if and only if

$$[\underline{x}\Omega - \alpha\gamma(\underline{x} - \underline{y})E\tilde{x}]w \leq (\alpha + \underline{x})\Omega\bar{c}.$$

This inequality holds whenever $\gamma \geq \bar{\gamma}$. For $\gamma < \bar{\gamma}$ it holds whenever $w \leq \bar{w}$. We conclude that all first order conditions hold for the indicated values of the multipliers under the conditions of Proposition 2.

Proof of Proposition 3

It has to be checked that the first order conditions (8)-(12) hold for some appropriate values of λ , μ , ν . In particular, consider

$$\begin{aligned}\lambda &= \frac{E\tilde{y} - E\tilde{x}}{\underline{x} - \underline{y}} \frac{1}{bE\tilde{x} + sE\tilde{y}}, \\ \mu &= 0, \\ \nu &= 0.\end{aligned}$$

Inserting the values for c_1 , b , s , as they are stated in Proposition 3, and rearranging shows that (8) holds for the indicated values of the multipliers (use $w - b - s = c_1$). Inspection of (8) and (9) reveals that the latter condition holds whenever the former holds. To check this, solve (8) for $\frac{\gamma}{w-b-s-\bar{c}}$ and insert the resulting expression as well as the expression for λ into (9). Consider next (10). Clearly, $\lambda > 0$ from Assumption 1. Check by inserting for c_1 , b , s that $b\underline{x} + s\underline{y} = \alpha c_1$. Thus, (10) holds. Check next (11). Clearly, $\mu \geq 0$. Moreover, we have $b > 0$ if $\gamma \geq \bar{\gamma}$. For $\gamma < \bar{\gamma}$, $b > 0$ is implied by $w < \bar{w}$. Finally, consider (12). We trivially have $\nu \geq 0$ and $\nu s = 0$. Moreover, $\gamma < \bar{\gamma}$ and $w > \bar{w}$ imply $s > 0$. We conclude that all first order conditions hold for the indicated values of

the multipliers.

Proof of Proposition 4

Proof of (i). Again, it has to be checked that the first order conditions (8)-(12) hold for $b = 0$ and some appropriate values of λ , μ , ν . In particular, consider

$$\begin{aligned}\lambda &= \frac{1}{\alpha + \underline{y}} \left[\frac{\gamma}{c_1 - \bar{c}} - \frac{1}{s} \right], \\ \mu &= \frac{1}{\alpha + \underline{y}} \left[\frac{\Omega}{sE\tilde{y}} - (\underline{x} - \underline{y}) \frac{\gamma}{c_1 - \bar{c}} \right], \\ \nu &= 0.\end{aligned}$$

It is straightforward to check that (8) and (9) hold (use $w - b - s = c_1$). Consider next (10). Insert for c_1 and s in the expression for λ . Rearranging shows that $\lambda \geq 0$ whenever

$$(\alpha\gamma - \underline{y})w \geq -(\alpha + \underline{y})\bar{c}. \quad (14)$$

Thus, $\lambda \geq 0$ if $\gamma \geq \frac{\underline{y}}{\alpha}$. For $\gamma < \frac{\underline{y}}{\alpha}$ we have $\lambda \geq 0$ whenever $w \leq \bar{\bar{w}}$. It follows from this and from $s\underline{y} = \alpha c_1$ and $b = 0$ that (10) holds.

Consider next (11). Insert for c_1 and s in the expression for μ . Rearranging shows that $\mu \geq 0$ whenever

$$[\underline{y}\Omega - \alpha\gamma(\underline{x} - \underline{y})E\tilde{y}] w \geq (\alpha + \underline{y})\Omega\bar{c}.$$

The inequalities $\gamma < \bar{\gamma}$ and $w \geq \bar{\bar{w}}$ imply $\mu \geq 0$. Since $b = 0$, we conclude that (11) holds. Finally, it is straightforward that (12) holds.

Proof of (ii). Set

$$\begin{aligned}\lambda &= 0, \\ \mu &= \frac{\gamma}{c_1 - \bar{c}} - \frac{E\tilde{x}}{sE\tilde{y}}, \\ \nu &= 0.\end{aligned}$$

Simple calculations show that (8) and (9) hold. (Insert for c_1 , b and s to check the latter

condition). Consider now (10). Inserting for c_1 , b , s shows that $b\underline{x} + s\underline{y} - \alpha c_1 \geq 0$ whenever

$$(\underline{y} - \alpha\gamma)w \geq (\alpha + \underline{y})\bar{c}.$$

This is implied by $\gamma < \frac{y}{\alpha}$ and $w > \bar{\bar{w}}$. (Note that the latter inequalities imply $\gamma < \bar{\gamma}$ and $w > \bar{\bar{w}}$.) This and $\lambda = 0$ imply (10). Consider next (11). Inserting for c_1 and s shows that $\mu \geq 0$ is implied by Assumption 1 and by the fact that $w > \bar{c}$. Since $b = 0$, it follows that (11) holds. Finally check (12), which holds in case that $s \geq 0$. This, again, is true because $w > \bar{c}$.

References

- Beach, Lee Roy, and Mitchell, Terence R. (1978), "A Contingency Model for the Selection of Decision Strategies," *Academy of Management Review*, 3, pp. 439-449.
- Binswanger, Johannes (2007a), "Life Cycle Decision Making without Backward Induction: A Model with Feasibility Goals," Working Paper.
- Binswanger, Johannes (2007b), "Risk Management of Pensions from the Perspective of Loss Aversion," *Journal of Public Economics*, 91, 641-667.
- 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.
- Bröder, Arndt (2003), "Decision Making with the 'Adaptive Toolbox': Influence of Environmental Structure, Intelligence, and Working Memory Load," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(4), pp. 611-625.
- Campbell, John Y., and Cochrane, John H. (1999), "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), pp. 205-51.
- 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. (forthcoming), "Relative Income, Happiness and Utility: An Explanation for the Easterlin Paradox and Other Puzzles," *Journal of Economic Literature*, forthcoming.
- Cocco, João F., Gomes, Francisco J., and Maenhout, Pascal J. (2005), "Con-

- sumption and Portfolio Choice over the Life Cycle,” *Review of Financial Studies*, 18(2), pp. 491-533.
- Constantinides, George M. (1990), “Habit Formation : A Resolution of the Equity Premium Puzzle,” *Journal of Political Economy*, 98(3), pp. 519-43.
- Dynan, Karen E., Skinner, Jonathan and Zeldes, Stephen P. (2004) “Do the Rich Save More?” *Journal of Political Economy*, 112(2), pp. 397-444.
- Easterlin, Richard A. (1974), “Does Economic Growth Improve the Human Lot? Some Empirical Evidence,” in: David Paul A., and Reder, Melvin W. (Eds.), *Nations and Households in Economic Growth: Essays in Honor of Moses Abramowitz*, New York: Academic Press.
- Easterlin, Richard A. (1995), “Will Raising the Incomes of All Increase the Happiness of All?” *Journal of Economic Behavior and Organization*, 27(1), pp. 35-47
- Einhorn, Hillel J. (1971), “The Use of Nonlinear, Noncompensatory Models as a Function of Task and Amount of Information,” *Organizational Behavior and Human Performance*, 6, pp. 1-27.
- Ford, Kevin J., Schmitt, Neal, Schechtman, Susan L., Hults, Brian M., and Doherty, Mary L. (1989), “Process Tracing Methods: Contributions, Problems, and Neglected Research Questions,” *Organizational Behavior and Human Decision Processes*, 43, pp. 75-117.
- Frederick, Shane, Loewenstein, George, and O’Donoghue, Ted (2002), “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 40(2), pp. 351-401.
- Gabaix, Xavier, and Laibson, David (2000), “A Boundedly Rational Decision Algorithm,” *American Economic Review*, 90(2), AEA Papers and Proceedings, pp. 433-438.
- Gabaix, Xavier, and Laibson, David (2005), “Bounded Rationality and Directed Cognition,” Working Paper.
- Gollier, Christian (2001), *The Economics of Risk and Time*, MIT Press, Cam-

- bridge, Mass.
- 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.
- 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.
- Munnell, Alicia H., and Soto, Mauricio (2005), "What Replacement Rates Do Households Actually Experience in Retirement?" Center for Retirement Research Working Papers, Boston College.
- Newell, Allen, and Simon, Herbert A., (1972), *Human Problem Solving*, Prentice-Hall, Englewood Cliffs, NJ.
- Payne, John W. (1976), "Task Complexity and Contingent Processing in Decision Making: An Information Search and Protocol Analysis, *Organizational Behavior and Human Performance*, 16, 366-387.
- Payne, John W., Bettman, James R., and Johnson, Eric J.(1993), *The Adaptive Decision Maker*, Cambridge University Press, Cambridge.
- Shefrin, Hersh M., and Thaler, Richard H. (1988), "The Behavioral Life-Cycle Hypothesis," *Economic Inquiry*, 26(4), pp. 609-43.
- Simon, Herbert A. (1955), "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69(1), pp. 99-118.

Figure 1: Optimal life cycle portfolio allocation: Theory

Figure 2: Saving rates for $\alpha = 0.3$, $\gamma = 0.96^{-30}$, $\bar{c} = 33$

Figure 3: Equity shares for $\alpha = 0.3$, $\gamma = 0.96^{-30}$, $\bar{c} = 33$

Figure 4: Saving rates for $\alpha = 0.35$, $\gamma = 0.96^{-30}$, $\bar{c} = 33$

Figure 5: Equity shares for $\alpha = 0.35$, $\gamma = 0.96^{-30}$, $\bar{c} = 33$

Figure 6: Saving rates for $\alpha = 0.3$, $\gamma = 0.96^{-30}$, $\bar{c} = 25$

Figure 7: Equity shares for $\alpha = 0.3$, $\gamma = 0.96^{-30}$, $\bar{c} = 25$

Figure 8: Saving rates for $\alpha = 0.3$, $\gamma = 0.94^{-30}$, $\bar{c} = 33$

Figure 9: Equity shares for $\alpha = 0.3$, $\gamma = 0.94^{-30}$, $\bar{c} = 33$

Figure 10: Saving rates for $\alpha = 0.27$, $\gamma = 0.94^{-30}$, $\bar{c} = 29$

Figure 11: Equity shares for $\alpha = 0.27$, $\gamma = 0.94^{-30}$, $\bar{c} = 29$

Figure 12: Saving rates for CRRA preferences, baseline

Figure 13: Equity shares for CRRA preferences, baseline

Figure 14: Saving rates for CRRA preferences, best fit

Figure 15: Equity shares for CRRA preferences, best fit

Figure 16: Saving rates for HARA preferences, best fit

Figure 17: Equity shares for HARA preferences, best fit