



Understanding the Heterogeneity of Savings and Asset Allocation: A Behavioral-Economics Perspective

Johannes Binswanger

► To cite this version:

Johannes Binswanger. Understanding the Heterogeneity of Savings and Asset Allocation: A Behavioral-Economics Perspective. *Journal of Economic Behavior and Organization*, 2010, 76 (2), pp.296. 10.1016/j.jebo.2010.06.010 . hal-00870188

HAL Id: hal-00870188

<https://hal.science/hal-00870188>

Submitted on 6 Oct 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Accepted Manuscript

Title: Understanding the Heterogeneity of Savings and Asset Allocation: A Behavioral-Economics Perspective

Author: Johannes Binswanger

PII: S0167-2681(10)00126-5
DOI: doi:10.1016/j.jebo.2010.06.010
Reference: JEBO 2567

To appear in: *Journal of Economic Behavior & Organization*

Received date: 12-8-2008
Revised date: 29-6-2010
Accepted date: 30-6-2010

Please cite this article as: Binswanger, J., Understanding the Heterogeneity of Savings and Asset Allocation: A Behavioral-Economics Perspective, *Journal of Economic Behavior and Organization* (2010), doi:10.1016/j.jebo.2010.06.010

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



Understanding the Heterogeneity of Savings and Asset Allocation: A Behavioral-Economics Perspective*

Johannes Binswanger[†]

Abstract

Why do saving rates and the proportion of savings invested in stocks substantially increase with income? This paper sheds light on this question from the perspective of a new bounded-rationality life cycle model. It has been shown elsewhere that this model is better than existing ones at predicting how savings and asset allocation choices vary over the life cycle. In this paper it is shown that a very simple version of the same model also allows for a better understanding of the variation in savings and asset allocation choices across income groups.

Key words: Behavioral economics, bounded rationality, income heterogeneity, life cycle saving, portfolio choice, stock market participation.

JEL classification: D81, D91.

*I am very grateful to Roland Bénabou, Antoine Bommier, Martin Browning, Peter Diamond, Josef Falkinger, Ernst Fehr, Reto Föllmi, Guido Friebe, Simon Gächter, Christian Gollier, David Laibson, Wieland Müller, Manuel Oechlin, 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.

Understanding the Heterogeneity of Savings and Asset Allocation: A Behavioral-Economics Perspective*

Johannes Binswanger[†]

June 29, 2010

Abstract

Why do saving rates and the proportion of savings invested in stocks substantially increase with income? This paper sheds light on this question from the perspective of a new bounded-rationality life cycle model. It has been shown elsewhere that this model is better than existing ones at predicting how savings and asset allocation choices vary over the life cycle. In this paper it is shown that a very simple version of the same model also allows for a better understanding of the variation in savings and asset allocation choices across income groups.

Key words: Behavioral economics, bounded rationality, income heterogeneity, life cycle saving, portfolio choice, stock market participation.

JEL classification: D81, D91.

*I am very grateful to Roland Bénabou, Antoine Bommier, Martin Browning, Peter Diamond, Josef Falkinger, Ernst Fehr, Reto Föllmi, Guido Friebel, Simon Gächter, Christian Gollier, David Laibson, Wieland Müller, Manuel Oechslin, James Poterba, Panu Poutvaara, Eytan Sheshinski, Daniel Schunk, Jean Tirole, Andreas Wagener, Richard Zeckhauser, Josef Zweimüller and seminar participants at GRE-MAQ, 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

Empirically, saving rates increase substantially with income. For the U.S., saving rates are below ten percent for the lowest income quintile, while they exceed 20 percent for the fifth quintile (Dynan et al., 2004).¹ Similarly, the proportion of savings invested in stocks increases strongly with income. Typically, the bottom two quintiles invest between zero and ten percent of their financial wealth in stocks. In contrast, the top quintile invests about 60 percent in stocks.²

The question why saving and investment behavior differ so much across income groups is of interest because it relates directly to the issue of adequate retirement preparation. Does adequate retirement preparation differ across income groups? Or do members of some income groups save and invest in an inadequate way? Any judgment about adequate retirement preparation behavior requires an understanding of individuals' saving motives and how they may differ across income groups. One major aim of this paper is to shed light on this issue.

In principle, differences in saving rates and equity shares between income groups can arise from three sources. First, the nature of the budget constraint may differ between income groups due to social welfare programs. Second, different income groups may differ in financial literacy and this may lead to differences in savings and asset allocation. Third, preferences may be non-homothetic, i.e. income expansion paths for saving rates and equity shares may be non-linear.

Differing budget constraints have been put forward as an explanation for differing *saving* rates by Hubbard, Skinner and Zeldes (1995). They demonstrate that lower-income households have less incentives to save since they are more likely to depend on welfare programs in future bad states of the world. When it comes to *asset allocation*, the effect of budget constraints is less clear. In fact, if low-income earners are more likely to depend on welfare, this may induce them to take *more* risk in the form of a higher equity

¹These numbers include various forms of retirement savings.

²See Section 6 below.

share. Furthermore, both saving rates and equity shares continue to increase for higher income quintiles where a potential dependence on welfare seems less likely.³ Thus, this pattern cannot easily be explain by referring to differing budget constraints. The same is true for fixed absolute costs of stock market participation. These may explain why individuals with low income stay out of the stock market. However, such fixed costs do again not necessarily imply that equity shares would increase with income among those participating in the stock market.

Bernheim (1998) and Lusardi and Mitchell (2007a, 2007b) provide evidence that financial literacy is correlated with the level of education. Moreover, financial literacy is correlated with wealth accumulation and stock market participation. To the degree that the level of education is positively correlated with income, this suggests that financial literacy may explain why savings and the likelihood of stock market participation increase with income. While the evidence is strong that a higher level of financial literacy leads to higher absolute levels of savings – and thus higher wealth accumulation – there is no clear evidence that it leads to higher saving *rates*. Furthermore, the variation of equity shares with income has a similar shape for all education groups with at least a high school degree.⁴ The latter observation is not compatible with the view that only financial literacy, and not income itself, leads to equity shares raising with income.

In sum, differing budget constraints and differing levels of financial literacy are very likely to explain part of the variation of saving rates and equity shares over income. However, they cannot fully account for the pattern in the data. Thus, non-homothetic preferences are also likely to be part of the explanation for why saving rate and equity share profiles are jointly increasing with income. This paper puts forward a new model of non-homothetic preferences.

The analysis of this paper is based on a new bounded-rationality life cycle model, dubbed feasibility-goals (FG) model, which has been recently developed (Binswanger, 2008). The theoretical innovation of this model is that it cuts out the main source of

³See Dynan et al. (2004) and Section 6 of this paper.

⁴See Section 6 and footnote 27 below.

behavioral complexity in the standard model: the requirement of contingent planning, i.e. the need for anticipating one's future actions in each state of nature. The FG model eliminates this requirement while maintaining the assumption that agents are forward-looking and prudent rather than myopic. Agents achieve such behavior by means of focusing on so-called feasibility goals. These represent levels of consumption that agents want to be feasible in certain states of nature. Importantly with regard to this paper, the structure of the FG model is such that it implies a very simple pattern of non-homothetic behavior that is governed by one single parameter.

In Binswanger (2008) it is demonstrated by means of calibrations that the FG model, while being very tractable and parsimonious, is particularly well-suited to explaining how saving and asset allocation behavior vary *over the life cycle* for a given agent. There, the analysis is carried out within a fully-fledged multiperiod setup with transitory and permanent labor income shocks.

The aim of this paper is to show analytically and numerically that a very simple two-period version of the FG model is well-suited to better understanding how savings and asset allocation vary across different income groups. I will refer to the two-period version of the FG model simply as the goal model.⁵

Due to its high transparency and conceptual simplicity, the analysis of a two-period setup is particularly insightful and provides a natural first step before moving to more complex setups to study the heterogeneity of savings and asset allocation behavior. A two-period setup represents the simplest case in which the joint decision of saving and asset allocation can be studied in a meaningful way. The setup is intended to represent a life cycle framework where agents work and prepare for retirement in the first period and live from the returns on their savings in the second. The two-period setup allows for concentrating on the variation of savings and asset allocation across different levels of lifetime (permanent) incomes and is easily accessible to an analytical examination. This is useful for developing insight into the basic mechanisms at work. Besides a theoretical

⁵The aspect of “feasibility” is vacuous in a two-period environment.

analysis, I also provide simulation results that allow for a broad judgment about how well the model captures the key patterns in the data from a quantitative point of view.

The analysis includes a comparison of the predictions of the goal model with the predictions of a subclass of hyperbolic absolute risk aversion (HARA) preferences as well as with existing habit formation preferences. While the goal model explains the data very well, HARA and existing habit preferences have difficulties in accounting for the patterns observed in the data. In particular, for plausible parameters, one obtains very high and – in the case of HARA preferences – *decreasing* instead of increasing saving rates. The performance of the goal model remains superior even if the HARA and habit models are endowed with an additional large number of degrees of freedom by assuming that their parameters may differ across income groups. In this way, it is demonstrated that the success of the goal model is not due to a particular flexibility that may allow it to explain just any pattern in the data.

The variation of saving rates across different income groups has been carefully analyzed by Dynan et al. (2004). I am not aware of any comprehensive analysis of how savings and asset allocation jointly vary with lifetime income. This paper also contributes to the behavioral-economics literature on life cycle saving. 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), as well as the loss aversion model of Bowman et al. (1999). To the best of my knowledge, the behavioral-economics literature has not attempted to explain the variation of saving and asset allocation choices across income groups, so far.

The rest of the paper is organized as follows. Section 2 introduces the goal model. Section 3 discusses the psychological underpinning of the model. Section 4 states the assumptions made about the economic environment. Section 5 presents the analytical solution of the model. In Section 6, I calibrate the model and discuss its quantitative predictions. In Section 7, I provide a numerical analysis of some extensions of the goal model. Section 8 compares the predictions of the goal model to predictions obtained for HARA and existing habit formation preferences. Finally, Section 9 concludes.

2 The Goal Model

This section introduces the two-period version of the FG model in a self-contained way. The fully-fledged version of this model has been developed in Binswanger (2008). As already mentioned, within the two-period framework of this paper this model will simply be dubbed goal model. This model should be understood as a preference model. However, in contrast to the standard case, preferences are represented by a list of goals that a decision maker aims to achieve. The goals are ordered in a hierarchical way, such that a goal i is triggered only if all goals with an index lower than i are already fully achieved. This sequential activation of goals represents the main novel element of the model. The goals should be understood as functions of the decision variables. A discussion of the psychological underpinning of the model is postponed to the next section since it comes more natural once the model has already been outlined.

It may appear from the fact that preferences are represented by a hierarchically ordered list of goals that this hierarchical order directly implies a multitude of non-homothetic elements. As a result, the model would have ample latitude in accounting for heterogeneity in savings and asset allocation behavior. As will become clear later, this is, however, not the case. In fact, there will be one *single* parameter that governs the degree of non-homothetic behavior of the model.

There are a total of three goals in the model. The first goal is dubbed insurance goal. It entails assuring that second-period (random) consumption \tilde{c}_2 does not fall below a specific fraction α of first-period consumption c_1 with a very high probability. I consider the limit case where this probability equals one. Therefore, the insurance goal is met if and only if $\min \tilde{c}_2 \geq \alpha c_1$. In principle, the parameter α can take on any nonnegative value, but it is natural to imagine that it is smaller than one. The economic rationale of the first goal is the desire to avoid extreme downside risks in future consumption. Naturally, this goal relates to habit formation or loss aversion.⁶

⁶Note that achieving the insurance goal is always feasible whenever there exists a financial asset with strictly positive minimum returns. To see this note that the ratio $\min \tilde{c}_2/c_1$ can then be made arbitrarily large by choosing a value of c_1 arbitrarily close to zero.

The second goal represents a desire to achieve what I dub a normal standard of living \bar{c} in the first period. By the logic of the hierarchy of goals, this also entails the desire to achieve a minimum consumption of at least $\alpha\bar{c}$ in the second period. The parameter \bar{c} is again assumed to be nonnegative. It is a crucial parameter, as it is exactly this (and only this) parameter that makes decision making non-homothetic. It is most suitable to think of this parameter as the real value of expenditures associated with a “typical” or “normal” standard of living.

The economic rationale of \bar{c} is that it captures, in a stylized manner, the finding that people seem to be concerned about how their standard of living compares to a reference consumption level.⁷ This reference level may reflect the average standard of living in a reference group or in the society as a whole.⁸

The final goal is a non-satiation goal ensuring that agents do not reach an overall saturation point. It is given by the Cobb-Douglas-type function $E\tilde{c}_2(c_1 - \bar{c})^\gamma$ and entails the desire to achieve jointly a high value of current “excess” consumption and expected future consumption. The parameter γ is assumed to be strictly positive and determines the preference weight of current consumption. Subtracting \bar{c} from c_1 ensures that c_1 is always a strictly normal good for all levels of income. Furthermore, it implies that the motive of achieving a high level of $E\tilde{c}_2$ becomes gradually more important as c_1 gradually increases beyond \bar{c} . The fact that the first goal entails infinite risk aversion (since individuals will never risk to fall below the “habit level” αc_1) while there is risk neutrality in the third goal gives rise to a particularly simple risk-return trade-off.

Overall, preferences in the goal model are represented by the list

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

containing the three goals introduced above. I denotes the indicator function. Agents

⁷See Easterlin (1974, 1995), Clark et al. (2008), and the discussion in the next section.

⁸This suggests that \bar{c} may increase over time. Taking into account that \bar{c} underlies a time trend does not add any crucial insights to the analysis and I will thus treat \bar{c} as exogenous and constant in the main part of the analysis. In Appendix B, I discuss how the model can be extended such that \bar{c} raises endogenously over time, whereas all the main predictions of the goal model are preserved.

maximize a particular goal i , subject to achieving the maximum of all goals with an index lower than i .⁹ The particular goal i that represents the objective function is determined by the size of the budget (see Section 5). The first goal can always be fully achieved if there is an asset with a strictly positive minimum return. Therefore, the objective function that determines an agent's behavior will always correspond to either $i = 2$ or $i = 3$. If $\bar{c} = 0$, then we will always have $i = 3$ and the model is fully homothetic.^{10, 11}

From simple inspection of (1) it is not immediately clear whether choices based on the maximization on those preferences are always well defined. However, the analysis in Section 5 will make this sufficiently clear, such that there is no need to explore this issue in greater detail here. It may be useful, however, to 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.

As will be discussed in the next section, one important psychological underpinning of the model is bounded rationality. A particular property of the goal model that relates to bounded rationality is the fact that, concerning the future, individuals take into account only two states of nature (out of many possible ones). The first goal relates to a future worst-case scenario.¹² It entails the desire not to consume less than a fraction α of current consumption in case that this worst-case scenario arises. The term $E\tilde{c}_2$ in the third goal can be understood as relating to a future “normal” scenario. This normal

⁹The maximum of the first goal is achieved if $I(\tilde{c}_2^{min} \geq \alpha c_1) = 1$, the maximum of the second goal is achieved if $c_1 \geq \bar{c}$.

¹⁰Inspection 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 $E\tilde{c}_2$. 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 $\alpha \bar{c}$. 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.

¹¹It 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 models including habit formation and loss aversion.

¹²This should be understood in pragmatic terms. The “worst” case may correspond, say, to the first percentile of the consumption distribution.

scenario is defined as a state of nature where all random variables take on their expected values. The fact that, in the model, individuals disregard all information regarding future consumption that relates to states of nature other than the worst-case or normal scenario may make it particularly suitable for the description of the behavior of boundedly rational agents.

The focus on only two possible states of nature excludes, however, that individuals may explicitly exhibit any form of volatility aversion. To see this, consider a given savings/investment plan and hold fixed the payoff of this plan for the worst-case and the normal scenario, i.e. hold fixed $\min \tilde{c}_2$ and $E\tilde{c}_2$. Suppose now that a shock occurs that increases the volatility of \tilde{c}_2 (only) across all other states of nature. According to the goal model, this shock would not affect an individual's savings and investment behavior.

At first sight, the absence of volatility aversion may seem a rather restrictive feature that may limit the predictive power of the goal model. However, in practice, this may be of limited importance. A shock that affects only the volatility of the consumption distribution without affecting its downside risk seems rather artificial. For most distributions that may realistically describe the distribution of payoffs from savings, such as the log-normal distribution, an increase in the variance also increases the downside risk.¹³ Thus, boundedly rational agents may benefit from this by focusing only on the downside risk of future consumption since, implicitly, this will mean also taking into account its volatility. Section 7 provides evidence on this point by presenting simulations that explicitly take volatility aversion into account.

An interesting question is whether the preference model (1) should be seen as particularly “flexible,” given that it consists of three separate goals. At first sight, one may get the impression that the model consists, in fact, of a collection of three different utility functions. However, these three utility functions would be so stiff and stylized that they barely merit this name. This notwithstanding, Section 8 will address the issue of flexibility more formally from a quantitative point of view.

¹³For instance, an increase in the variance is associated with a decrease of, say, the first percentile which, in the model, would be treated as the literal minimum.

3 The Psychological Underpinning of the Goal Model

It is useful to start the discussion of the psychological underpinning of the goal model with highlighting the main difference to existing models. Consider, for instance, an *existing* model of habit formation. As it is the case for the goal model, also existing models of habit formation are characterized by multiple goals. In each period, there is a habit level that an individual aims to achieve.¹⁴ Beyond this, an individual values a high consumption level in the current period and values a high expected utility of a risky consumption profile, in excess of the habit level, in the future. However, unlike in the case of the goal model, existing models of habit formation are characterized by the fact that all goals are always active *simultaneously*. Thus, a decision maker pursues the habit levels, a high level of current consumption and of expected utility of future consumption *at the same time*. In contrast, the goal model is characterized by a *sequential activation of goals* according to a priority order.

The most important psychological underpinning of this sequential activation of goals is the common observation that some human needs and desires are more basic than others. In other words, it is natural to order needs in a hierarchical way. Desires with a lower priority are only pursued once the desires with a higher priority are satisfied (see Foellmi and Zweimüller, 2006, 2008, for an illuminating discussion). In economics, this idea has first been put forward by Engel (1857), based on his own empirical research. In psychology, this has famously been propagated by Maslow (1943). As noted by Foellmi and Zweimüller, the idea of hierarchical needs can be traced back at least to Plato.

The first goal in the goal model specifies the desire to avoid extreme downside risks in future consumption streams. Therefore, it is directly related to the goal of adequate survival. In principle, this goal could also be captured as the desire to achieve an absolute minimum standard of living that allows for a level of calories intake that ensures survival. However, the evidence is in favor for that individuals aim to avoid downside risk relative to their current standard of living since there is habit formation (Duesenberry, 1949;

¹⁴See (11) and (12) in Section 8 for a formal statement of existing habit formation preferences.

Constantinides, 1990; Campbell and Cochrane, 1999).

The second goal in the goal model specifies the desire to achieve a normal standard of living. It can be understood as the desire to “keep up with the Joneses”, in other words, to consume what others have around you (again Duesenberry, 1949; Easterlin, 1974, 1995; Clark et al., 2008). There is evidence that this desire has a “hard-wired” component (Fliessbach et al., 2007) and it is thus a powerful driver of behavior. This notwithstanding, it is plausible to assume, as the goal model does, that avoiding large downside risks in future consumption is still a more basic goal than keeping up with the Joneses and that the former has a higher priority whenever there is a conflict between the two goals.

The first and second goal in the model are satiable in that they can be fully achieved with a specific finite level of resources. If we believe that humans are always endowed with a goal to long for, a further *non-satiable* goal is required to close the model. The model specifies this residual non-satiable goal as the general goal to “become rich,” in the sense of longing for a high level of expected future consumption. This – relatively unspecific – goal has the lowest priority and is activated only once the other two goals are satiated. The goal model uses the simplest possible mathematical specification for this goal by neglecting any type of risk aversion.¹⁵

From a technical point of view, optimal choices in a model with only satiable goals may not be unique. The reason is that, if the budget set is sufficiently large, there may be many choices that satisfy the set of satiable goals. If a researcher’s aim is to make sharp predictions and to test them, this is an undesirable property. This provides a further rationale for the non-satiable and residual nature of the third goal.

The hierarchy of goals also relates to the psychological phenomenon of limited attention and information processing capabilities, in other words to bounded rationality.¹⁶ Clearly, the act of deciding upon one’s life cycle savings is a cognitively demanding task. There is

¹⁵This is not a crucial assumption, however. See Section 7 for an inclusion of volatility aversion. The term $E\tilde{c}_2$ in the third goal could also be replaced by $Eu(\tilde{c}_2)$ where u represents a standard (Bernoulli) utility function. However, the model would lose its analytical tractability, as a result.

¹⁶See Pashler (1998) for a comprehensive survey on the psychology of attention.

overwhelming evidence that, if individuals face cognitively demanding choice tasks, they tend to focus on a very limited number of goals and only information is processed that relates to these goals (Ford et al., 1989, Payne et al., 1993, Markman and Brendl, 2000). Once the goals with the highest priority have been achieved, cognitive resources are freed for processing information that relates to further goals.

In sum, there are two main psychological underpinnings of the basic structure of the goal model: hierarchy of desires; and limited attention. Both make the case for a sequential activation of goals. In contrast, existing models with multiple goals, such as existing habit formation or Stone-Geary preference models, entail that all goals are pursued simultaneously. One main purpose of this paper is to demonstrate that the hypothesis of sequential goal activation may be a powerful approach to understand typical patterns of life cycle saving and asset allocation.

4 The Economic Environment

Consistent with the idea that the two-period model of this paper represents a simple model of retirement preparation, agents earn an exogenous income w in the first period while they earn no non-financial income in the second period.¹⁷

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.¹⁸

Denote bond holdings by b and stock holdings by s . Furthermore, denote the *gross* returns on bonds by \tilde{x} . Bond returns are possibly risky, as indicated by the tilde. Gross stock returns are denoted by \tilde{y} . I make the following assumptions about return distribu-

¹⁷Thus, the model does not account for income from Social Security. In Section 6, the model will be calibrated and compared to empirical estimates of saving rates that take into account various sources of retirement income including Social Security. This assures a proper comparison of model predictions with the data.

¹⁸See e.g. Carroll (1997), Campbell and Viceira (2002), Cocco et al. (2005), Polkovnichenko (2007).

tions.

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 in comparison to a pure bond portfolio.

Budget constraints are given by

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

$$\tilde{c}_2 = b\tilde{x} + s\tilde{y}, \quad (3)$$

for the first and second period, respectively. Furthermore, I assume that

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

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

5 Solving the Model

It follows from the discussion in Section 2 that (if $\bar{c} > 0$) 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 $\underline{x} > 0$ 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

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

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 the objective function does not change the outcome of an optimization problem, this program may be stated as

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

The rest of this section deals with the solution of 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 could do better by substituting all stock holdings by bonds. Since $\underline{x} > \underline{y}$, the habit consumption αc_1 could then be maintained with less resources. Hence, the individual could afford a higher level of c_1 without jeopardizing the insurance 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, households with a relatively lower income are much less likely to participate in the stock market (see Section 6 and Binswanger, 2008). This corresponds exactly to the prediction of Proposition 1. 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 substantial transaction costs (see e.g. Gollier, 2001). This observation gives rise to the so-called stock market participation puzzle.

The logic of the goal model's explanation of the puzzle follows immediately from the derivation of Proposition 1. According to the goal model, earners of a lower income save exclusively for precautionary reasons. Their only saving motive is to achieve the insurance goal, as this is a more basic goal than the desire to savor a high expected old age consumption. Given the assumptions of the model, this result is not surprising. However, this does not preclude that the insurance goal may be a key motive of saving for households with a relatively low lifetime income.

It should be mentioned that zero or very low equity shares at low levels of income can also be explained by existing models with additive internal habit formation (Polkovnichenko, 2007). Unfortunately, this comes at the costs of the prediction of high saving rates for low incomes that initially *decrease* with income. In fact, saving rates approach the limit of one when moving arbitrarily close to the level of resources that just allows for meeting habit levels. In contrast, according to Proposition 1, the goal model predicts a saving rate of $\frac{\alpha}{\alpha+\underline{x}}$ which is always strictly between zero and one. A further problem of existing habit formation models is that optimal choices may not be defined for low endowment levels.

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,

$$\bar{\gamma} \equiv \frac{x\Omega}{\alpha(\underline{x}-\underline{y})E\tilde{x}} \text{ and } \bar{w} \equiv \frac{(\alpha+x)\Omega}{x\Omega-\alpha\gamma(\underline{x}-\underline{y})E\tilde{x}}\bar{c}.$$

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

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, in a local environment to the right of \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 $\bar{\gamma} \equiv \frac{y\Omega}{\alpha(\underline{x}-\underline{y})E\tilde{y}}$ and $\bar{w} \equiv \frac{(\alpha+y)\Omega}{y\Omega-\alpha\gamma(\underline{x}-\underline{y})E\tilde{y}}\bar{c}$. By direct calculations it can be verified that $\bar{\gamma} < \bar{\gamma}$ and $\bar{w} > \bar{w}$ 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 $\bar{w} < w < \bar{w}$, then

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

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. The inequality $w > \bar{w}$ 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 αc_1 can eventually be met only with stocks. 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, c_1 is not so low that habit consumption can be met only with stocks. 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 c_1 and $E\tilde{c}_2$ are sufficiently balanced such that individuals invest in both assets.

The key implication of Proposition 3 is that both the saving rate $(b + s)/w$ and the equity share $s/(b + s)$ are predicted to increase with w . For the saving rate this follows from the fact that the ratio c/w is decreasing in w . For the equity share this can be shown by calculating the derivative $d[s/(b + s)]/dw$. The intuition for these results follows from the fact that achieving 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 away from \bar{c} . As a result, the marginal value of an increase of c_1 becomes lower and lower relative to the marginal value of an increase in $E\tilde{c}_2$. This leads to both higher savings and a higher share of savings invested in stocks. Note that the assumption of $\bar{c} > 0$ is key for this result. With $\bar{c} = 0$ the model would be fully homothetic and saving rates and equity shares would be constant. (Note that we would have $w^{crit} = \bar{w} = \bar{\bar{w}} = 0$ in this case.) A non-zero value of this

parameter leads to a particularly simple way of non-homothetic decision making.¹⁹

The main aim of the goal model is to shed light on the determinants of the distribution of saving rates and equity shares across income groups. Nevertheless, it is of interest to ask whether the goal model can contribute to an understanding of why there has been a downward trend in *aggregate* saving rates since the seventies in the U.S.²⁰ One major change that occurred in the mid-sixties was the introduction of Medicare – a social insurance program providing health insurance for the elderly – and Medicaid – providing health insurance for the poor. In the model, the introduction of these programs can be captured as a decrease in the effective value of the parameter α . This means that these social insurance programs reduce the effective downside risk of old-age consumption which has the same effect as if a household had a lower aversion to downside risk, i.e. a lower value of α . As a result of a decrease in α , Proposition 1, 2, 3²¹ all predict a unambiguous increase in the consumption share c_1/w and hence a decrease in saving rates. Thus, while many other factors have surely been important, the goal model points out that the introduction of Medicare and Medicaid may have played a significant role for the decrease in national savings in the U.S.

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 $\bar{\gamma} > \frac{y}{\alpha}$ (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 $w \geq \bar{\bar{w}}$, then $b = 0$. Moreover,*

¹⁹If \bar{c} were to be seen as constant over a long-term time horizon, then, in the presence of economic growth, the goal model would predict (aggregate) saving rates to increase over time. However, as outlined in Appendix B, in the presence of economic growth it is appropriate to treat \bar{c} as endogenous such that, in a steady state, it would increase at the same rate as aggregate income. In this case, there is no time trend in the aggregate savings rate. Importantly, the cross-sectional predictions of Proposition 3 continue to be valid if \bar{c} evolves endogenously over time. In particular, savings rates and equity shares still increase with income within a given period. See Appendix B for details.

²⁰In the sixties, the average ratio of net savings to GDP was about 11 percent. Since the year 2000, it has been around 2 percent. (See National Income and Product Accounts, Table 5.1.)

²¹Note that Ω depends positively on α .

(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 showing the different portfolio regimes in the γ/w -space. It should be noted that the different 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 $\gamma = 0$ (and $\bar{c} > 0$) we have $\bar{\bar{w}} = \bar{\bar{\bar{w}}} > \bar{w} = w^{crit}$.

6 Calibrating the Model

6.1 Parameter Values for a Benchmark Calibration

In this section, I discuss how the goal model can be calibrated in order to develop insight into the predictions of the model from a quantitative point of view. Interest in the quantitative aspects of the model's predictions arises for three reasons. First, the theoretical analysis predicts that there are several portfolio regimes (see Figure 1). It is of interest

to see which regimes arise for plausible parameter values. Second, calibrating the model for plausible parameter values shows whether the quantitative predictions come close to matching the relevant empirical data, at least in a broad sense. In particular, I will compare the model's predictions to estimates of saving rates and equity shares for the U.S. Third, a quantitative analysis allows for investigating how well the goal model captures the patterns in the data in comparison to other models such as HARA or existing habit formation preferences.

A comparison of the goal model with existing ones within a two-period framework may not allow for a definite conclusion about which model will ultimately explain reality better. It is, in principle, conceivable that the goal model would explain the data better than, e.g., the HARA model within a two-period setup while the opposite would hold in a fully-fledged multiperiod setup. While theoretically possible, this seems rather unlikely in light of the analysis in Binswanger (2008) as well as in light of the results reported below. I thus view the two-period calibrations shown below as suggestive and as a useful preliminary step for a horse race between more fully-fledged versions of the goal and existing models, which I leave for future research.

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 beginning of age 86. This setup takes into account that retirement typically lasts for fewer years than working life. For simplicity, I do not take mortality risk into account. This may be seen as appropriate in an environment where agents have access to an explicit or implicit form of annuities.

It remains to determine the investment horizon associated with a typical "retirement preparation dollar." I follow the methodology of Binswanger (2007) by assuming that the length of the investment horizon is given by the distance between the mid-point of retirement (76) and the mid-point of working age (46), which amounts to 30 years.²²

The model will first be calibrated for a set of baseline parameter values that are cho-

²²Alternatively, imagine that during each of the 40 years during working life the agent prepares for half a year of retirement in a sequential manner. Thus, during age 26 the agent prepares for his retirement from age 66 to 66.5. During age 27, he prepares for retirement from age 66.5 to age 67 etc. The *average* time horizon for these "split" retirement saving problems is then again 30 years, up to rounding.

sen 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 60 percent for low income levels and 30 percent for high income levels (Munnell and Soto, 2005). The underlying reason for this monotonic decline in income is that households with a higher income 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 “consumption replacement” ratios. 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 temporal 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 year during working life and one year 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. This ratio amounts to 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.²³ 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. I set annual gross returns on bonds to 1.02, assuming that they are risk-free. Annual net returns are thus equal to 2 percent. The distribution

²³See U.S. Census Bureau, Historical Income Tables. The 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.

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.²⁴ 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. These assumptions allow for simulating 30-year level returns. I draw 1,000,000 such 30-year returns and set the minimum stock return to the first percentile of the resulting simulated 30-year distribution, amounting to 0.62. The expected return is, of course, 1.06 to the power of 30.

6.2 Empirical Savings and Equity Share Profiles

Empirical saving rates are taken from Dynan et al. (2004) and represent separate estimates of saving rates for each income quintile.²⁵ It is important to stress that these saving rates include prospective income from Social Security, which is converted into a notional saving rate and added to private saving rates. Furthermore, private saving rates also include private pensions. Such all-inclusive saving rates provide the natural empirical benchmark for evaluating model 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 incomes for each income quintile referring to the 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. Equity holdings include indirect holdings through mutual funds and pension accounts.²⁶ Empirical saving rates

²⁴These numbers as well as the assumption of bond returns of 1.02 are fairly standard. See e.g. Campbell and Viceira (2002, Chapter 7).

²⁵Specifically, saving rates are taken from the “Active+Pension” column in Table 3 of Dynan et al. (2004). The estimates for saving rates refer to ages between 40 and 49. I am grateful to Karen Dynan for kindly providing information about median quintile incomes for the Panel Study of Income Dynamics.

²⁶There are five percent of observations for which financial assets have a value of zero. I either drop these observations or set their equity share to zero. Both procedures lead to a virtually identical profile of equity shares over income quintiles. Clearly, equity shares are also zero for households who do hold positive financial wealth but do not hold any equity. These households should, of course, not be dropped. Specifically, these households account for the zero equity share for the first income quintile in the second

and equity shares are shown in the top section in Table 1.²⁷

6.3 Calibration Results

The results for the baseline calibration are shown in the second section of Table 1 and in Figure 2 (saving rates) and 3 (equity shares). For the latter two figures, the solid lines refer to model predictions while the dashed lines correspond to empirical estimates. The goal model is fairly successful in capturing the general shape of the empirical profiles. It matches well the increasing profile of saving rates, the increasing profile of equity shares, as well as the fact that the equity share is zero for the first quintile. Relating to the portfolio regimes in Figure 1, the baseline calibration shows that it is the regimes corresponding to Propositions 1, 2, and 3 (in this order as income increases) that arise for empirically relevant parameter values. From the fact that equity shares are far below one we infer that the regime of Proposition 4 does not arise.

It is useful to have a summary measure at hand for how closely the model matches the empirical data. For this, I report the root of the mean squared prediction error which is defined as follows. Define empirical saving rates for quintile i by $\hat{\sigma}_i$ and predicted saving rates by σ_i^* . Similarly, define empirical and predicted equity shares by $\hat{\rho}_i$ and ρ_i^* , respectively. The root mean squared prediction error is then given by

$$RMSE = \left[0.5 \sum_{i=1}^5 (\sigma_{ij}^* - \hat{\sigma}_i)^2 + 0.5 \sum_{i=1}^5 (\rho_{ij}^* - \hat{\rho}_i)^2 \right]^{0.5}. \quad (7)$$

For brevity, I refer to this simply as the prediction error. As shown in Table 1, the

row of Table 1.

²⁷Consistent with the temporal setup of the model, equity shares in Table 1 refer only to ages less than 65. 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, the curve has a similar shape for different education classes. For those with only a high-school degree, equity shares are zero percent for the first and 63 percent for the fifth quintile. For the group with some college education, equity shares increase from 0 to 49 percent. For those with a college degree, they increase from 6 to 61 percent. Only the shape for the lowest education class with no high-school degree differs more substantially. For this latter group, equity shares range between zero and 14 percent.

prediction error amounts to 0.10 for the baseline specification.

I turn now to a sensitivity analysis where I consider the effect of varying one parameter at a time from the baseline level. This provides some insight into the “mechanics” of the model. I will then present a specification which optimizes the fit to the data.

Under the label “high α ”, the third section in Table 1 shows 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 Subsection 6.1). In this case, the precautionary savings motive is more important than in the baseline case. This shifts saving rates upwards and equity shares downwards. Under the label “low \bar{c} ”, the fourth section in Table 1 presents results for $\bar{c} = 25$ instead of 33. In this case c_1 departs sooner from \bar{c} . As a result, maximization of $E\bar{c}_2$ 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. Under the label “high γ ”, the fifth section in Table 1 presents 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 relative weight in the third goal in (1). As a consequence, $E\bar{c}_2$ has a lower relative weight.

Figures 4 and 5, and the bottom section of Table 1 show calibrations for a set of parameter values that are chosen such as to minimize the prediction error, as defined by (7). Specifically, the minimization is carried out numerically 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. The minimization is carried out numerically and I set $\mathcal{A} = \{0.25, 0.26, \dots, 0.4\}$, $\mathcal{C} = \{15, 16, \dots, 35\}$, $\Gamma = \{0.80, 0.81, \dots, 0.98\}$.

The resulting best fit values are $\alpha = 0.27$, $\bar{c} = 29$, $\gamma = 0.94^{-30}$. On the parameter space $\mathcal{A} \times \mathcal{C} \times \Gamma$, these values provide the unique minimum for the prediction error. 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. These parameter values come fairly close to the baseline parameters and appear rather natural. Figures 4 and 5 show that the fit is practically perfect for equity shares and very good for saving rates. The prediction error is as low as 0.02.

7 Some Extensions

7.1 Volatility Aversion

In the baseline version of the goal model, the only source of risk aversion is the desire to avoid large downside risks in future consumption. This is captured by the first goal in (1). In contrast, there is no autonomous aversion to a high volatility of second period consumption, *holding fixed its downside risk*. I have argued in Section 2 that this feature should be seen as unproblematic since, in reality, an investment opportunity with a higher volatility is highly likely to be also associated with a higher downside risk. Thus, downside risk may serve as a sufficient statistic for both types of risks and there is no need to take volatility aversion into account explicitly.

In spite of this argument, it is of interest to see how volatility aversion can be incorporated into the goal model and how its numerical predictions are affected by this. In particular, I consider the following specification of the *third* goal in (1):

$$[E\tilde{c}_2 - \phi\sigma(\tilde{c}_2)](c_1 - \bar{c}), \quad (8)$$

where $\sigma(\tilde{c}_2)$ denotes the standard deviation of \tilde{c}_2 . The first two goals remain identical to the baseline version. While the model loses its analytical tractability with (8) as the third goal, it is straightforward to obtain numerical results.

Setting α , \bar{c} , γ to the best-fit values, it turns out that if ϕ is smaller than 5, volatility aversion has no effect on predicted choices. The reason is that downside risk aversion leads already to a portfolio choice that is associated with a rather low standard deviation of \tilde{c}_2 . The second section in Table 2 shows the predictions of the model, augmented by volatility aversion, for ϕ equal to 5. All other parameters are set to the best-fit values (see Table 1). The value for ϕ of 5 should be seen as rather high. To see this consider an individual with the prospect of consuming an expected annual amount of \$50,000. Suppose that the standard deviation would be equal to \$5,000. With ϕ equal to 5, this individual would be indifferent between this prospect and one with a safe annual consumption that is only

half as high, i.e. \$25,000. It is only that if volatility aversion is at least as high as this, that the predictions of the extended version of the model deviate from the ones for the best-fit specification. In fact, even for $\phi = 5$, the deviation is not very significant.

In sum, the main point of simulating the goal model with (8) as the third goal has been to show that, due to the fact that a higher downside risk tends to go together with a higher volatility, incorporating volatility aversion explicitly into the model does not alter its predictions, unless it is extraordinarily high.

7.2 Heterogeneous Values for α

In Section 6, the choice of the baseline value for α has been determined by the replacement rate of the U.S. Social Security system for low-income earners. I have assumed that the value of α is identical across income groups. There have been two reasons for this choice. First, a single value of α makes the model more parsimonious. Second, Social Security replacement rates for higher income quintiles may not be lower because higher-income earners exhibit a lower relative downside risk aversion. Rather, society may have agreed that these income groups would save also by themselves, while the lowest income group would not.

Nevertheless, it is of interest to see how the numerical predictions of the goal model are affected if α is assumed to be heterogeneous and set according to the Social Security replacement rates for each income group. The Social Security replacement rates are equal to .63, .49, .44, .39, .33, for the five quintiles, respectively (see Munnell and Soto, 2005, Table 10). As discussed in Section 6.1, these numbers need to be divided by two, in order to obtain quintile-specific values for α that are adjusted to the time-setup of the model. The other parameter values are set to the best-fit values. The simulation results are shown in the third section of Table 2 (third and fourth row). The specification captures the empirical pattern quite well, again. However, the prediction error of 0.10 is much larger than the one for the best-fit specification of 0.02. Overall, this leads to the conclusion that heterogeneous values of α according to the Social Security replacement

rates do not increase the predictive power of the model.

7.3 Heterogeneous Values for \bar{c}

In the baseline version, the parameter \bar{c} has been treated as an exogenous constant. It represents a reference point with which individuals compare their own consumption level and captures the desire to “keep up with the Joneses.” This assumption may be relaxed in two directions. First, it may be that different income groups have different consumption reference points and these increase with income. Second, in the presence of economic growth, this reference point may increase over time. In this subsection, I consider the first extension. The second one is addressed in Appendix B.

In order to get a feeling for how the predictions of the goal model change if \bar{c} is heterogeneous across income groups, I proceed as follows. I take the best-fit specification as the starting point. For this specification, \bar{c} amounts to 29. I now assume that this value applies to the third income quintile and that, for a quintile i , \bar{c} is determined according to

$$\bar{c}_i = 29 + 0.2(w_i - w_3) \quad (9)$$

for $i = 1, \dots, 5$, where w_i denotes the median income level for quintile i . Thus, the deviation of \bar{c}_i from 29 is proportional to $w_i - w_3$.

The results are shown in the bottom section of Table 2 (second and third row). The prediction error amounts to 0.06, whereas the prediction error for the best-fit specification is only 0.02. Thus, while the general shape of the empirical curves is matched rather well, the best-fit specification with a constant \bar{c} explains the data better.

8 Comparison to Existing Preference Models

In this section, I first compare the numerical predictions for the goal model to the corresponding predictions for a subclass of hyperbolic absolute risk aversion (HARA) prefer-

ences. The latter are a natural candidate for explaining an increasing pattern of saving rates and equity shares within the range of existing preference model. Second, it is also of interest to compare the goal model to an existing model of habit formation.

Starting with HARA preferences, I consider the specification

$$U(c_1, \tilde{c}_2) = \frac{1}{1-\eta} (c_1 - \bar{c}_1)^{1-\eta} + \frac{\beta}{1-\eta} E[(\tilde{c}_2 - \bar{c}_2)^{1-\eta}], \quad (10)$$

where β represents a discount factor, \bar{c}_1 , \bar{c}_2 are age-specific exogenous subsistence consumption levels, and η corresponds to relative risk aversion in case that $\bar{c}_1 = \bar{c}_2 = 0$. E denotes the mathematical expectation operator. It is important to note that non-zero values for \bar{c}_1 , \bar{c}_2 give rise to non-flat profiles of saving rates and equity shares. In contrast, both profiles are flat for $\bar{c}_1 = \bar{c}_2 = 0$.²⁸

Figures 6 and 7, as well as the second section in Table 3, show calibrations for a parameter constellation that has been chosen as being plausible ex ante. In particular, \bar{c}_1 and \bar{c}_2 have been set to life-cycle-specific poverty thresholds for the year 2001, which amount to 14.1 and 10.7 thousands of year-2001 U.S. dollars, respectively.²⁹ The parameter η is set to 3 and the annualized value of β is set to .96.

It turns out that this specification fits the data worse than the baseline version of the goal model. In particular, predicted saving rates are very high. Moreover, the saving rate profile is falling instead of increasing. Thus, for a priori plausible parameter values, the model misses the empirical picture even from a qualitative point of view.

The observation that the HARA model does not explain the data well for a set of baseline parameters does not yet establish the claim that the goal model is more satisfactory overall, however. It may be possible that the latter just enjoys a higher flexibility to explain any pattern in the data. To explore this issue more formally, I run a calibration for the HARA model where I endow the latter with a high number of degrees of freedom.

²⁸Equity share profiles are non-flat only if $\bar{c}_2 \neq 0$.

²⁹Source is the U.S. Census Bureau. These figures represent poverty thresholds for a weighted average of three-person households with age less than 65 for working life, and a weighted average of two-person households with age greater or equal to 65 for retirement, respectively.

I do so in a way that it should be beyond doubt that they exceed the degrees of freedom associated with the goal model.

Specifically, I identify a parameter constellation for HARA preferences that minimizes the root mean squared error function defined in (7). However, in contrast to the procedure chosen for the goal model in the last section, I allow for many more degrees of freedom. I do so by allowing the parameters \bar{c}_1 , \bar{c}_2 and η to take on different values *for each of the five income quintiles*. The parameters \bar{c}_1 , \bar{c}_2 may take on any integer value between zero and 14 for each income quintile.³⁰ The parameter η may take on any value between one and ten for each quintile. In contrast, β is held fixed across income levels and set to .96.

As shown in the third panel of Table 2, the profile of best-fit parameter values for the five income quintiles are as follows. For \bar{c}_1 the best-fit profile is given by (14, 14, 14, 14, 10), for \bar{c}_2 by (1, 0, 2, 10, 5). Finally, for η the profile is (10, 10, 3, 2, 2). The prediction error amounts to 0.04 (compared to 0.02 for the goal model). In light of the vast degrees of freedom that this specification of the HARA model has been endowed with, it is rather surprising that this fit is worse than in case of the best-fit parameter constellation for the goal model. Remember that for the goal model the best-fit parameter values have been derived under the natural restriction that parameters must be identical across income values.

While a variation of the value of \bar{c}_1 across income levels does not contribute to a better fit of the data, a large flexibility concerning \bar{c}_2 and η is clearly important. In particular, a good fit is achieved only at the price of large jumps in the values of \bar{c}_2 and η across income levels. Furthermore, the profile of \bar{c}_2 is not even monotonic. This suggests that the HARA model may not be particularly useful for understanding why and how saving rates and equity shares increase with lifetime/permanent income.

Taking an extreme view, one may see the goal model as a collection of three different utility functions (although indeed very crude ones that have little in common with what we usually consider as utility functions). Allowing parameters of HARA preferences to

³⁰I do not allow \bar{c}_1 and \bar{c}_2 to take on values higher than 14 since optimal choices may then not be defined for the first quintile.

take on different values for each income quintile means literally assuming that preferences are represented by a collection of five different utility functions. Furthermore, these preferences may differ in multiple dimensions since there are several parameters that may take on different values across income levels. Given that the HARA model explains the data less well than the goal model even in this case, it seems fair to conclude that the explanatory power of the goal model cannot merely be attributed to a high degree of flexibility.

I finally turn now to a comparison of the goal model with an existing model of habit formation. In particular, I consider a habit formation model where intertemporal preferences are given by

$$U(c_1, \tilde{c}_2) = \frac{1}{1-\varepsilon} (c_1 - \delta H_1)^{1-\varepsilon} + \frac{\rho}{1-\varepsilon} E[(\tilde{c}_2 - \delta H_2)^{1-\varepsilon}] \quad (11)$$

(see Section 5 in Gomes and Michaelides, 2003). H_1 and H_2 denote the habit levels in the first and second period, respectively. The parameter δ , taking on values between zero and one, determines the impact of the habit level on (marginal) utility. The parameter ε is equal to the coefficient of relative risk aversion in case that $\delta = 0$. The first-period habit level H_1 is taken as exogenous.³¹ H_2 is determined according to

$$H_2 = (1 - \lambda) H_1 + \lambda c_1. \quad (12)$$

H_2 is thus endogenous, unlike \bar{c}_2 in the case of HARA preferences. The parameter λ , taking on values between zero and one, determines the relative influence of first-period consumption on H_2 . The habit model (11) represents a so-called additive model of habit formation, in contrast to a multiplicative version. The additive specification has been chosen for its closer resemblance to the goal model.³²

³¹Compare this to the parameter \bar{c}_1 in the case of HARA preferences.

³²For the multiplicative version, the instantaneous utility function is given by $u(c) = \frac{1}{1-\varepsilon} \left(\frac{c}{H^*}\right)^{1-\varepsilon}$ (see Gomes and Michaelides, 2003). The calibration results for the multiplicative model are very similar to the ones for the additive model reported below.

As in the HARA case, I first simulate optimal savings and asset allocation decisions under the habit formation model (11), (12) for a set of parameter values that are chosen ex ante. Specifically, I set ε equal to 3, the annualized value of ρ equal to .96 and H_1 to 10. The latter value has been chosen to be identical to the baseline value for \bar{c}_1 in the case of HARA preferences. Referring to Gomes and Michaelides (2003), I set δ equal to .8 and λ equal to .5. The results are shown in the second-to-last section of Table 3. They confirm a finding that has been observed in other studies (Gomes and Michaelides, 2003; Polkovnichenko, 2007): existing habit models tend to predict very high savings rates.

I turn now to a specification where H_1 , δ , and λ are chosen such as to minimize the prediction error. All three parameters may take on different values for each income quintile. In particular, for each income quintile, H_1 may take on any value from the set $\{0, 2, 4, \dots, 30\}$, and δ and λ may take on any value from the set $\{0, 0.1, 0.2, \dots, 1\}$. As before, ε is set to 3, and the annualized value of ρ to .96. The predictions are shown in the bottom section of Table 3. The prediction error amounts to .11. This is to be compared with 0.04 for the best-fit HARA specification (with heterogeneous parameters), and with 0.02 for the best-fit specification of the goal model (with homogeneous parameters). The prediction error of the habit model is, thus, very high in light of the many degrees of freedom under which the optimization of fit has been carried out.

The profile for the best-fit values for δ over the five income quintiles is (0.5, 0.6, 0.1, 0, 0). Thus, the best fit is obtained if, for the three upper quintiles, habit formation has a weak influence on marginal utility or is even absent. The best-fit profile for λ is (1, 0, 0.3, 1, 1). For H_1 it is (30, 30, 16, 0, 0). Both profiles are characterized by a rather erratic structure.

At first sight, the poor performance of existing habit formation models may come as a surprise. After all, the goal model also incorporates habit formation, if in a different way. To understand why existing models of habit formation models perform purely, one needs to note that the desire to achieve the habit level in the future leads to rather high savings rates, unless δ is close to zero. Thus, a specification with a relatively high level of δ performs poorly because of the prediction of high saving rates. On the other hand, if δ is low, the habit model converges to a standard constant relative risk aversion (CRRA)

model. Thus, there is no room any more for saving rates and equity shares to increase with income. As a result of this, the habit model's performance is poor again, since it does not match the increasing profile of empirical savings rates and equity shares. In sum, there is no combination of parameter values that would allow the habit model to match the data well. In contrast to existing habit models, the goal model – parsimonious as it is – captures habit formation in a way that leads, simultaneously, to comparatively low savings rates and to an increasing pattern of saving rates and equity shares.

9 Conclusion

In Binswanger (2008) I have derived a new bounded-rationality life cycle model within a fully-fledged dynamic environment and I have shown that this model allows for a better understanding of how saving and asset allocation behavior vary over the life cycle. In this paper it is shown that a very simple two-period version of the same model sheds light on why saving and asset allocation decisions vary substantially across different income groups.

The model has potentially important implications for retirement saving (Binswanger, 2007). If the insurance goal is indeed the primary determinant of savings decisions for lower income levels, then zero or very low equity shares among lower-income earners may not be seen as a mistake. However, the dominance of the insurance goal may also be seen as reflecting an excessive degree of loss aversion. Whether or not the goal model should be accepted from a normative point of view remains certainly an open question. However, there may be individuals for whom the prospect of not achieving a certain “insurance goal” evokes real discomfort. For such individuals, the goal model may provide a reasonable and simple guideline for their savings and asset allocation choices even from a normative point of view. More generally, given the complexity of life cycle saving, simply trying to achieve a list of basic goals may prove quite useful for boundedly rational agents, after all.

Future research should address the comparison between the quantitative predictions of

existing preference models and of the goal model with respect to heterogeneity in savings and asset allocation behavior within a fully-fledged multiperiod setup. The results of this paper may prove helpful in understanding and interpreting the basic mechanisms behind such a more involved analysis.

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

$$\mathcal{L} = \gamma \ln(w - b - s - \bar{c}) + \ln(bE\tilde{x} + sE\tilde{y}) + \lambda[b\underline{x} + s\underline{y} - \alpha(w - b - s)] + \mu b + \nu s, \quad (13)$$

where λ , μ and ν are multipliers. It can easily be checked that the objective function is strictly concave in b and s . The first-order Kuhn-Tucker conditions for this program are

$$\frac{E\tilde{x}}{bE\tilde{x} + sE\tilde{y}} - \frac{\gamma}{w - b - s - \bar{c}} + (\alpha + \underline{x})\lambda + \mu = 0, \quad (14)$$

$$\frac{E\tilde{y}}{bE\tilde{x} + sE\tilde{y}} - \frac{\gamma}{w - b - s - \bar{c}} + (\alpha + \underline{y})\lambda + \nu = 0, \quad (15)$$

$$\lambda \geq 0, \quad b\underline{x} + s\underline{y} - \alpha(w - b - s) \geq 0, \quad \lambda [b\underline{x} + s\underline{y} - \alpha(w - b - s)] = 0, \quad (16)$$

$$\mu \geq 0, \quad b \geq 0, \quad \mu b = 0, \quad (17)$$

$$\nu \geq 0, \quad s \geq 0, \quad \nu s = 0. \quad (18)$$

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 (14)-(18) 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 (14) and (15) 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

$$(\underline{x} - \alpha\gamma)w \leq (\alpha + \underline{x})\bar{c}. \quad (19)$$

Either $\gamma < \frac{\underline{x}}{\alpha}$ or $\gamma \geq \frac{\underline{x}}{\alpha}$. Consider the first case. Then (19) 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 $\gamma \geq \frac{\underline{x}}{\alpha}$, (19) 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 (16) holds.

It remains to check (17) and (18). 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 (14)-(18) 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 (14) holds for the indicated values of the multipliers (use $w - b - s = c_1$). Inspection of (14) and (15) reveals that the latter condition holds whenever the former holds. To check this, solve (14) for $\frac{\gamma}{w-b-s-\bar{c}}$ and insert the resulting expression as well as the expression for λ into (15). Consider next (16). 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, (16) holds. Check next (17). 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 (18). 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 (14)-(18) 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 (14) and (15) hold (use $w - b - s = c_1$). Consider next (16). Insert for c_1 and s in the expression for λ . Rearranging shows that $\lambda \geq 0$ whenever

$$(\alpha\gamma - \underline{y})w \geq -(\alpha + \underline{y})\bar{c}. \quad (20)$$

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 (16) holds.

Consider next (17). 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 (17) holds. Finally, it is straightforward that (18) 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 (14) and (15) hold. (Insert for c_1 , b and s to check the latter condition). Consider now (16). 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 $\gamma < \bar{\bar{\gamma}}$ and $w > \bar{\bar{w}}$.) This and $\lambda = 0$ imply (16). Consider next (17). 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 (17) holds. Finally check (18), which holds in case that $s \geq 0$. This, again, is true because $w > \bar{c}$.

Appendix B: Time-Varying \bar{c}

In the main part of the analysis, the parameter \bar{c} is treated as an exogenous constant. This makes the model particularly tractable and allows to concentrate on the main points of interest. One problematic implication of this assumption is, however, that (aggregate) saving rates and equity shares increase over time in the presence of economic growth and thus a growing income level. This prediction is counterfactual. If the model is embedded in a framework of economic growth, then, within a steady state, a natural benchmark prediction would be one of no trend. In this appendix, I outline a modification of the model that can achieve this.

Consider an economy with an overlapping generations structure that is characterized by economic growth. Denote A_t the stock of “knowledge” in the economy at time t . Imagine that the evolution of the stock of knowledge is the main driver of economic growth. Assume that, in each period, there is one young and one old generation alive. Denote the labor income of a young individual i in period t by w_{it}^y and denote average income of the young in period t by $\bar{w}_t \equiv \int w_{it}^y dP_i$, where P_i indicates the distribution of incomes over individuals. For simplicity, P is taken as time-invariant. The old generation earns no labor income. Assume that, in a steady state, \bar{w}_t is a linear function of A_t .³³ Thus, we may write

$$\bar{w}_t = \phi A_t \quad (21)$$

for some coefficient $\phi > 0$ that depends on the specific model of economic growth.

In the presence of economic growth, it is natural that \bar{c} also grows over time and a time index should be added. In a steady state, \bar{c}_t must grow at the same rate as \bar{w}_t which, in turn, grows at the same rate as A_t .³⁴ Denote the growth rate of the latter variable by

³³This holds, for instance, if the aggregate production function is of Cobb-Douglas type and there is labor-augmenting technical progress; or it may hold for some variants of the so-called *AK*-Model.

³⁴See Carroll et al. (1997), and Alvarez-Cuadrado et al. (2004) for an analysis of economic growth under existing preferences with a keeping-up-with-the-Joneses motive of consumption.

g_A .

To check whether a steady state exists, we have to work through a fixed-point type of argument. Guess that \bar{c}_t grows at the rate g_A . Denote by c_{it}^y the consumption of a young individual i in period t . Assume, for simplicity, that $\bar{c}_t = \int c_{it}^y dP_i$, i.e. \bar{c}_t is equal to the average consumption of the young generation. The fixed-point argument entails that, if the c_{it}^y behave in such a way that \bar{c}_t grows indeed at the rate g_A – as has been guessed –, then a steady state exists. The crucial feature of this fixed-point argument is that the c_{it}^y are a function of \bar{c}_t .

To verify this formally, guess that \bar{c}_t grows at the rate g_A . Hence, we can write

$$\bar{c}_t = \psi A_t, \quad (22)$$

Assume that the income of a young individual i at time t is given by

$$w_{it}^y \equiv \chi_i^y \bar{w}_t, \quad (23)$$

i.e. χ_i^y characterizes the income distribution. The latter is assumed to be time-invariant, as a result of a constant distribution of talents in the economy. Consider now the regime corresponding to Proposition 3. Using Proposition 3, and (21), (22), (23), the share of income spent on consumption of a young individual is then given by

$$\frac{c_{it}^y}{w_{it}^y} = \frac{1}{1 + \gamma} \frac{\psi}{\chi_i^y \phi} + \frac{\gamma}{(1 + \gamma)\Omega} (\underline{x}E\tilde{y} - \underline{y}E\tilde{x}). \quad (24)$$

This implies that $c_{it}^y = \omega_i^y w_{it}^y$, for ω_i^y chosen according to (24). Using (23), it follows that

$$\bar{c}_t = \int \omega_i^y \chi_i^y \bar{w}_t dP_i \quad (25)$$

Thus, it follows from (21) that \bar{c}_t grows indeed at the rate g_A and, hence, a steady state exists. This argument can be extended to the case where some w_{it}^y fall into the regimes of Proposition 1, 2, or 4.

Since aggregate consumption grows at the same rate as aggregate income, there is no trend in the aggregate saving rate. The same can be shown to hold for aggregate equity shares. Crucially, however, the *individual* saving rates and equity shares are still increasing in χ_i^y and thus increasing in individual incomes w_{it}^y *within a given period*.

Thus, the baseline version of the goal model can be extended in such a way that the main cross-section predictions continue to hold: saving rates and equity shares increase with income within a given period. At the same time, there is no time trend for (aggregate) saving rates and equity shares. Overall, this shows that the simplifying assumption of a constant \bar{c} is rather innocuous and that the insights of the model extend directly to the case where this parameter is endogenous.

References

- Alvarez-Cuadrado, F., Monteiro, G., Turnovsky, S. J., 2004. Habit formation, catching up with the Joneses, and economic growth. *Journal of Economic Growth* 9, 47-80.
- Bernheim, D. B., 1998. Financial illiteracy, education and retirement saving. In: Mitchell, O. S., Schieber, S. J. (Ed.). *Living with defined contribution pensions*. University of Pennsylvania Press, Philadelphia, 38-68.
- Binswanger, J., 2007. Risk management of pensions from the perspective of loss aversion. *Journal of Public Economics* 91, 641-667.
- Binswanger, J., 2008. Towards understanding life cycle saving of boundedly rational agents: A model with feasibility goals. *Center Working Paper* 2008-14.
- Bowman, D., Minehart, D., Rabin, M., 1999. Loss aversion in a consumption-savings model. *Journal of Economic Behavior and Organization* 38, 155-178.
- Campbell, J. Y., Cochrane, J. H., 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 205-251.
- Campbell, J. Y., Viceira, L. M., 2002. *Strategic asset allocation: Portfolio choice for long-term investors*. Oxford University Press, Oxford.
- Carroll, C. D., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics*, 112, 1-55.
- Carroll, C. D., Overland, J., Weil, D. N., 1997. Comparison utility in a growth model. *Journal of Economic Growth* 2, 339-367.
- Clark, A. E., Frijters, P., Shields, M. A., 2008. Relative income, happiness and utility: An explanation for the Easterlin paradox and other puzzles. *Journal of Economic Literature* 46, 95-144.

- Cocco, J. F., Gomes, F. J., Maenhout, P. J., 2005. Consumption and portfolio choice over the life cycle. *Review of Financial Studies* 18, 491-533.
- Constantinides, G. M., 1990. Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy* 98, 519-543.
- Duesenberry, J. S., 1949. *Income, saving, and the theory of consumer behavior*. Harvard University Press, Cambridge MA.
- Dynan, K. E., Skinner, J., Zeldes, S. P., 2004. Do the rich save more? *Journal of Political Economy* 112, 397-444.
- Easterlin, R. A., 1974. Does economic growth improve the human lot? Some empirical evidence. In: David, P. A., Reder, M. W. (Ed.). *Nations and households in economic growth: Essays in honor of Moses Abramowitz*. Academic Press, New York, 89-125.
- Easterlin, R. A., 1995. Will raising the incomes of all increase the happiness of all? *Journal of Economic Behavior and Organization* 27, 35-47.
- Engel, E., 1857. Die Productions- und Consumptionsverhältnisse des Königreichs Sachsen. *Zeitschrift des Statistischen Büreaus des Königlich Sächsischen Ministeriums des Inneren* No. 8 und 9.
- Fliebsbach, K., Weber, B., Trautner, P., Dohmen, T., Sunde, U., Elger, C. E., Falk, A., 2007. Social comparison affects reward-related brain activity in the human ventral striatum. *Science* 318, 1305-1308.
- Foellmi, R., Zweimüller, J., 2006. Income distribution and demand-induced innovations. *Review of Economic Studies* 73, 941-960.
- Foellmi, R., Zweimüller, J., 2008. Structural change, Engel's consumption cycles, and Kaldor's facts of economic growth. *Journal of Monetary Economics* 55, 1317-1328.
- Ford, K. J., Schmitt, N., Schechtman, S. L., Hults, B. M., Doherty, M. L., 1989. *Process tracing methods: Contributions, problems, and neglected research*

- questions. *Organizational Behavior and Human Decision Processes* 43, 75-117.
- Gollier, C., 2001. *The economics of risk and time*. MIT Press, Cambridge MA.
- Gomes, F., Michaelides, A., 2003. Portfolio choice with internal habit formation: A life-cycle model with uninsurable labor income risk. *Review of Economic Dynamics* 6, 729-766.
- Hubbard, R. G., Skinner, J., Zeldes, S. P., 1995. Precautionary saving and social insurance. *Journal of Political Economy* 103, 360-399.
- Laibson, D., 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112, 443-477.
- Laibson, D., Repetto, A., Tobacman, J., 1998. Self-control and saving for retirement. *Brookings Papers on Economic Activity* 1998, 91-196.
- Lusardi, A., Mitchell, O. S., 2007a. Financial literacy and retirement preparedness: Evidence and implications for financial education. *Business Economics* 42, 35-44.
- Lusardi, A., Mitchell, O. S., 2007b. Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics* 54, 205-224.
- Markman, A. B., Brendl, M. C., 2000. The influence of goals on value and choice. *The Psychology of Learning and Motivation* 39, 97-128.
- Maslow, A. H., 1943. A theory of human motivation. *Psychological Review* 50, 370-396.
- Munnell, A. H., Soto, M., 2005. What replacement rates do households actually experience in retirement? Center for Retirement Research Working Papers, Boston College.

- Pashler, H. E., 1998. The psychology of attention. MIT Press, Cambridge MA.
- Payne, J. W., Bettman, J. R., Johnson, E. J., 1993. The adaptive decision maker. Cambridge University Press, Cambridge.
- Polkovnichenko, V., 2007. Life-cycle portfolio choice with additive habit formation preferences and uninsurable labor income risk. *Review of Financial Studies* 20, 83-124.
- Shefrin, H. M., Thaler, R. H., 1988. The behavioral life-cycle hypothesis. *Economic Inquiry* 26, 609-643.

Figure 1: Optimal life cycle portfolio allocation: Theory




Figure 2: Saving rates for the goal model: baseline




Figure 3: Equity shares for the goal model: baseline




Figure 4: Saving rates for the goal model: best-fit




Figure 5: Equity shares for the goal model: best-fit




Figure 6: Saving rates for HARA preferences (baseline)




Figure 7: Equity shares for HARA preferences (baseline)




Figure 8: Saving rates for existing habit formation preferences (baseline)




Figure 9: Equity shares for existing habit formation preferences (baseline)




Table 1: Numerical predictions for the goal model

	Quintiles				
	First	Second	Third	Fourth	Fifth
<i>Empirical</i>					
Savings rates	0.09	0.13	0.17	0.19	0.24
Equity shares	0	0.09	0.35	0.46	0.59
<i>Baseline: $\alpha = 0.3, \bar{c} = 33, \gamma = 0.96^{-30}$</i>					
Savings rates	0.14	0.15	0.19	0.23	0.25
Equity shares	0	0.05	0.47	0.66	0.78
Prediction error			0.10		
<i>“High α”: $\alpha = 0.35, \bar{c} = 33, \gamma = 0.96^{-30}$</i>					
Savings rates	0.16	0.16	0.21	0.24	0.27
Equity shares	0	0	0.39	0.58	0.71
Prediction error			0.07		
<i>“Low \bar{c}”: $\alpha = 0.3, \bar{c} = 25, \gamma = 0.96^{-30}$</i>					
Savings rates	0.14	0.19	0.22	0.25	0.27
Equity shares	0	0.43	0.65	0.76	0.83
Prediction error			0.19		
<i>“High γ”: $\alpha = 0.3, \bar{c} = 33, \gamma = 0.94^{-30}$</i>					
Savings rates	0.14	0.14	0.16	0.18	0.20
Equity shares	0	0	0.21	0.38	0.49
Prediction error			0.07		
<i>“Best fit”: $\alpha = 0.27, \bar{c} = 29, \gamma = 0.94^{-30}$</i>					
Savings rates	0.13	0.14	0.16	0.18	0.19
Equity shares	0	0.09	0.34	0.48	0.57
Prediction error			0.02		

Note: Prediction error is defined according to (7) in the main text.

Table 2: Numerical predictions for extensions of the goal model

	First	Second	Quintiles Third	Fourth	Fifth
<i>Empirical</i>					
Savings rates	0.09	0.13	0.17	0.19	0.24
Equity shares	0	0.09	0.35	0.46	0.59
<i>Volatility aversion: $\alpha = 0.27$, $\bar{c} = 29$, $\gamma = 0.94^{-30}$, $\phi = 5$</i>					
Savings rates	0.13	0.13	0.16	0.17	0.19
Equity shares	0	0.01	0.29	0.43	0.54
Prediction error			0.04		
<i>Heterogeneous α-values: $\bar{c} = 29$, $\gamma = 0.94^{-30}$</i>					
Replacement rate profile	0.63	0.49	0.44	0.39	0.33
α -profile	0.31	0.24	0.22	0.19	0.16
Savings rates	0.14	0.15	0.19	0.23	0.25
Equity shares	0	0.06	0.48	0.66	0.78
Prediction error			0.10		
<i>Heterogeneous \bar{c}-values: $\alpha = 0.27$, $\gamma = 0.94^{-30}$</i>					
\bar{c} -profile	21	25	29	34	41
Savings rates	0.13	0.15	0.16	0.17	0.18
Equity shares	0	0.21	0.34	0.42	0.48
Prediction error			0.06		

Note: Prediction error is defined according to (7) in the main text.

Table 3: Numerical predictions for alternative preference models

	Quintiles				
	First	Second	Third	Fourth	Fifth
<i>Empirical</i>					
Savings rates	0.09	0.13	0.17	0.19	0.24
Equity shares	0	0.09	0.35	0.46	0.59
<i>HARA baseline: $\bar{c}_1 = 14$, $\bar{c}_2 = 11$, $\eta = 3$, $\beta = 0.96^{30}$</i>					
Savings rates	0.29	0.28	0.28	0.28	0.28
Equity shares	0.03	0.21	0.29	0.32	0.32
Prediction error			0.14		
<i>HARA best fit: $\eta = 3$, $\beta = 0.96^{30}$</i>					
\bar{c}_1 -profile	14	14	14	14	10
\bar{c}_2 -profile	1	0	2	10	5
η -profile	10	10	3	2	2
Savings rates	0.14	0.22	0.22	0.25	0.24
Equity shares	0	0.09	0.36	0.48	0.58
Prediction error			0.04		
<i>Habit baseline: $\delta = 0.8$, $\lambda = 0.5$, $H_1 = 10$, $\varepsilon = 3$, $\rho = 0.96^{30}$</i>					
Savings rates	0.36	0.37	0.37	0.38	0.38
Equity shares	0.14	0.19	0.22	0.24	0.24
Prediction error			0.21		
<i>Habit best fit: $\varepsilon = 3$, $\rho = 0.96^{30}$</i>					
δ -profile	0.5	0.6	0.1	0	0
λ -profile	1	0	0.3	1	1
H_1 -profile	30	30	16	0	0
Savings rates	0.24	0.32	0.28	0.27	0.27
Equity shares	0.04	0.12	0.36	0.41	0.41
Prediction error			0.11		

Note: Prediction error is defined according to (7) in the main text.