

Rules of thumb in life-cycle savings models*

Ralf Rodepeter and Joachim K. Winter**

First draft: October 1999

This version: August 2000

Abstract: We analyze life-cycle savings decisions when households use simple heuristics, or rules of thumb, rather than solve the underlying intertemporal optimization problem. The decision rules we explore are a simple Keynesian rule where consumption follows income; a forward-looking rule that corresponds to the permanent income hypothesis; a static rule where only a fraction of positive income shocks is saved; and two rules that have been found in experimental studies. Using these rules, we simulate life-cycle savings decisions numerically and compute the utility losses relative to the backwards solution of the intertemporal optimization problem. Our central finding is that the utility losses induced by rule-of-thumb behavior are relatively low. We show that simple decision rules can be successful if they reflect one of the central savings motives in the life-cycle model, long-term consumption smoothing or short-term buffer-stock saving. Our simulations suggest that behaving optimally, in the sense of solving an intertemporal optimization model, is not necessarily much better than using much simpler heuristics which do not require computationally demanding tasks such as backward induction.

Keywords: savings, life-cycle models, bounded rationality, rules of thumb

JEL classification: D91; E21

* We should like to thank Michael Adam, Axel Börsch-Supan, Angelika Eymann, Werner Güth, Silke Januszewski, Ronald Lee, Dan McFadden, Matt Rabin, Paul Ruud, and seminar participants at Berkeley, Mannheim, NYU and the TMR Workshop on Savings, Pensions, and Portfolio Choice, Deidesheim, for helpful discussions and for comments on earlier versions of this paper. Simone Kohnz and Melanie Lührmann provided valuable research assistance. Financial support was received from the Deutsche Forschungsgemeinschaft (Sonderforschungsbereich 504 at the University of Mannheim) and the European Union (TMR project "Savings, pension, and portfolio choice"). Joachim Winter gratefully acknowledges the hospitality of the Department of Economics at the University of California, Berkeley, where parts of this paper have been written.

** Corresponding author.

Address: Sonderforschungsbereich 504, Universität Mannheim, D-68131 Mannheim.

E-mail: winter@uni-mannheim.de; *Phone:* +49-621-181-3438; *Fax:* +49-621-181-3451.

1 Introduction

It is a well-known finding from psychological research on decision-making that individuals use heuristics, or rules of thumb, in making judgments and decisions. In economics, rule-of-thumb behavior is an important aspect of bounded rationality (Simon (1955)).¹ Life-cycle consumption and savings decisions are a case in point. There is a large literature on such models and their solution. In realistic versions which incorporate income uncertainty, the solution of the underlying intertemporal optimization problem is rather complicated, and it requires backward induction because no closed-form solution for current consumption as a function of the relevant state variables exists. It has frequently been argued that individuals are unable to perform the calculations which are required to solve the underlying intertemporal optimization problem by backwards induction; see, *inter alia*, Wärneryd (1989), Pemberton (1993), Thaler (1994), and Hey (1999). The aim of this paper is to analyze life-cycle savings decisions when households use simple heuristics, or rules of thumb, rather than solve the underlying intertemporal optimization problem.

Before we review the empirical evidence on how individuals solve the life-cycle savings problem, we take a step back. It is important to be clear about the fact that we do not know, and indeed might never know exactly, the cognitive process by which individuals make their consumption decisions, i. e., we do not know their preferences. We can, however, assume that if individuals have preferences over all possible states of nature at the the current and any future date, there will be some intertemporal utility function that individuals maximize. Using observed behavior, we then can make inferences about individuals' preferences.

The standard approach in the life-cycle literature is to assume that preferences are additively separable over time and that there is some discounting of future utility. More specifically, it is standard to assume that the rate at which individuals discount future utility is constant and that the within-period utility is of the constant relative risk aversion (CRRA) type.² There exists a well-defined intertemporal optimization problem which corresponds to these intertemporal preferences. This problem is well understood, and it serves as a powerful tool in applied research; Browning and Lusardi (1996) review this vast literature.

The open question, however, is: Do individuals actually solve this problem? It is important to see that this question actually consists of two parts: First, are our assumptions about individuals' preferences correct? And second, given that our assumptions about individuals' preferences are correct, do individuals behave optimally? In this paper, we concentrate on

¹ We do not provide a review of the literature on boundedly rational behavior in economics here. For excellent surveys, see Camerer (1995), Rabin (1998), and McFadden (1999).

² There are many papers which depart from this standard model, and many good reasons, too. An important example is the literature on hyperbolic discounting (see Laibson (1998) for a review), and there are also many models with alternative within-period utility functions, such as Caballero (1990).

the second question by comparing the utility outcomes of using simple heuristics, or rules of thumb, with the benchmark given by the solution to the underlying intertemporal optimization problem. We take preferences as given, and we maintain the assumptions of the standard life-cycle model as outlined above.

There is a large and still growing empirical literature which addresses the research questions raised in the previous paragraph from different perspectives.³ It is our interpretation of this literature that the question whether rational behavior is an empirically valid assumption in life-cycle models is still open. A few examples serve to illustrate this point. In experimental studies of intertemporal decision making, rational behavior is frequently rejected. In the context of life-cycle models, some experimental studies test whether subjects perform backward induction in cognitive tasks that involve some dynamic trade-off (e.g., Hey and Dardanoni (1988), Carbone and Hey (1997)). In their experiments, backward induction, and hence rational behavior, is strongly rejected. In an experimental study of a search model (which is an intertemporal decision task slightly different from life-cycle decision making, but more akin to experimental study), Moon and Martin (1990) find that individuals use quite good heuristics, and with additional simulations, they show that these heuristics can be very close to optimal search rules. In his experimental study of a simple savings task, Anderhub (1998) also finds that individuals use relatively sophisticated heuristics, but do not use backwards induction. It should therefore not be surprising that in econometric studies using field data, it is typically difficult to distinguish between optimal and heuristic behavior, and intertemporal optimization cannot be rejected. For example, in dynamic discrete choice problems such as retirement decisions, the solution to an intertemporal optimization seems to be a valid approximation of actual behavior (e.g., Rust and Phelan (1997)).

Given the state of the empirical literature, we can add another research question to those stated before: If individuals do not solve the underlying optimization problems when making intertemporal decisions, why do we still observe decisions that are so close to model predictions? Should we be surprised if we observe *as if* behavior? One possible explanation which we explore in this paper is that individuals behave boundedly rational in the sense that they follow rules of thumb, and that their decisions lead to observed behavior that is similar to optimal behavior based on the solution of the underlying optimization problem. We will show in this paper that the utility loss associated with boundedly rational behavior can be small even in a complex intertemporal decision problem such as the life-cycle consumption-savings model.

³ We do not attempt a thorough review of the experimental and econometric literatures on choice over time here; the reader is referred to Loewenstein (1992), Rust (1994), Camerer (1995), Rabin (1998), and Hey (1999) for comprehensive reviews from different perspectives.

Rules of thumb have, of course, been analyzed in savings models before. In the macroeconomics literature, simple rules of thumb have been used in tests of the life-cycle/permanent income hypotheses based on aggregate consumption data, starting with the seminal papers by Hall (1978) and Flavin (1981).⁴ In these models, the assumption is that some fraction of the population behaves according to some simple rule of thumb such as “just consume your current income in every period”. In this paper, we explore this rule and four other rules of thumb that have been used in the economics literature on life-cycle savings behavior or which correspond to behavior which has been observed in laboratory experiments. We do not consider how such rules might arise endogenously in a learning context (as in Lettau and Uhlig (1999)) or due to institutional arrangements in the economy, but we return to this issue in the concluding section.

Our methodological approach is to compute life-cycle savings decisions under exogenously specified behavioral rules and to compare the outcomes with the optimal solution, using a compensating variation approach based on life-time utility. In doing this, we ignore any costs of computing optimal decisions, an issue to which we return below.⁵ Our approach is most closely related to earlier work on near-rational behavior in intertemporal consumption and savings problems by Cochrane (1989), Glaeser and Paulson (1997) and Lettau and Uhlig (1999), and to the study of heuristics in a search problem by Moon and Martin (1990).

The remainder of this paper is structured as follows. In Section 2, we present a version of the standard life-cycle model of savings decisions which allows for both life-time and income uncertainty. Next, we describe five rules of thumb which can be used to make savings decisions in this framework (Section 3). In Section 4, we simulate and compare savings decisions based on these rules of thumb. Section 5 concludes.

2 The benchmark life-cycle savings model

We assume that an individual’s or household’s optimal life-cycle consumption and savings behavior can be derived from a well-defined intertemporal optimization problem, given additively separable preferences with constant exponential discounting and CRRA within-period utility function. We use a version of the standard life-cycle model with implicit borrowing constraints and both life-time and income uncertainty which has been analyzed by Rodepeter and Winter (1998); it is an extension of the model proposed by Carroll (1992, 1997). In the remainder, it is understood that the decision-making unit is the household even though we

⁴ Other papers which consider the use of rules of thumb in macroeconomic models are Campbell and Mankiw (1990), Ingram (1990) and Krusell and Smith (1996).

⁵ Such costs have been considered in some models of rule-of-thumb behavior such as Shi and Epstein (1993) and Hindy and Zhu (1997).

usually refer to individual decisions; this corresponds to the household-level data we use to calibrate the model for our simulations.

Individuals are assumed to maximize, at each discrete point τ in time, the expected discounted stream of utility from future consumption. The per-period utility function is denoted by $u(C_\tau)$, to be specified below. Future utility is discounted by a factor $(1 + \rho)^{-1}$, where ρ is the time preference rate. The interest rate is denoted by r . The maximum age a person can reach is T , and we define s_t^τ as the probability to survive period t conditional on having survived period τ . To simplify notation, we also use a binary random variable that indicates whether an individual survives period t conditional on having survived period $t - 1$:

$$S_t = \begin{cases} 1 & \text{if the individual survives period } t \\ 0 & \text{if the individual does not survive period } t \end{cases}$$

The individual's intertemporal optimization problem can be stated as follows. In the planning period τ , the maximization problem is given by:

$$\max_{\{C_t\}_{t=\tau}^T} E_\tau \sum_{t=\tau}^T (1 + \rho)^{\tau-t} s_t^\tau u(C_t) \quad \text{s.t.} \quad (1)$$

$$A_t = (1 + r)(A_{t-1} + Y_{t-1} - C_{t-1}) \quad (2)$$

$$A_\tau \geq 0 \quad (3)$$

$$A_T \geq 0 \quad (4)$$

Maximization of expected discounted utility given by (1) is subject to standard restrictions, an asset recursion (2) and non-negativity conditions for initial and terminal assets (3) and (4). Note that while we require assets to be zero in the terminal period T , the individual might die before T with non-zero assets, i. e., there are accidental bequests in our model. We do not include an explicit borrowing constraint; rather, we impose the borrowing constraint implicitly.⁶ As Zeldes (1989) has shown, a borrowing constraint arises endogenously if consumption cannot go to zero in each period (i. e., if the marginal utility of consumption goes to infinity as consumption goes to zero), and if there is a positive probability of income dropping to zero in each period. The former is ensured by an appropriate functional form of the utility function $u(C_t)$, the latter by the specification of the income process.

The income process, Y_t , is formulated in terms of a long-term income component, P_t , as in many standard life-cycle models with income uncertainty (see, e.g., Carroll (1992)). Note that this long-term income component is not exactly the same as permanent income in the

⁶ Deaton (1991), among others, considers *explicit* liquidity constraints.

traditional sense, although the literature usually refers to P_t as permanent income. Specifically, we define current income, Y_t , as

$$Y_t = S_t V_t P_t. \tag{5}$$

Here, the long-term income component, P_t , is weighted with two random variables. First, as an extension to the Carroll model, we take into account life-time uncertainty via the “survival” variable S_t . Recall that this variable reflects life-time uncertainty and takes the value 1 as long as the individual is alive while it is set to zero thereafter. Second, labor income is weighted with V_t , a random variable with unit expectation that allows for periods with zero income. This zero-income variable is specified as

$$V_t = \begin{cases} 0 & \text{w.p. } p \\ 1/(1-p) & \text{w.p. } (1-p) \end{cases} \tag{6}$$

where p is an exogenous small probability. In the life-cycle literature, this zero-income shock is introduced to assure that borrowing constraints arise endogenously. Hence, this assumption should be interpreted as a technical device rather than a feature which makes the model more realistic. One could think of these zero-income periods as periods during which the individual is unemployed while after retirement, zero-income periods might be thought of as periods in which unforeseen circumstances (such as large health expenditures) depress disposable income.⁷ To keep the model simple, the process that governs these zero income realizations is assumed to be serially uncorrelated.

The long-term income component itself is assumed to follow a random walk with drift, an assumption which is standard in the literature. Earnings shocks affect the income process via the equation

$$P_t = G_t P_{t-1} N_t, \tag{7}$$

where G_t is the exogenously fixed and deterministic rate of wage growth, and N_t is a log-normally distributed random variable with unit expectation and variance σ which captures income uncertainty. Note that when income follows a random walk, a shock to current long-term income shifts the entire path of future income.

Finally, we assume that the within-period utility function is of the constant relative risk aversion (CRRA) type,

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{8}$$

where $\gamma \geq 1$ is the coefficient of relative risk aversion (and the inverse of the intertemporal elasticity of substitution).

⁷ We admit that in countries such as Germany where a tight social safety net protects individuals from ending up with zero consumption, such interpretations would seem less realistic than in, say, the United States.

As in any model of intertemporal decision making, the individual's decisions can be described by a time-invariant decision rule, i. e., a mapping from states into actions. In the life-cycle savings model, such a decision rule will be a function $C_t = C_t(A_t, Y_t)$ that maps current assets and current income into saving decisions. As noted before, we take the decision rule given by the dynamic programming solution to the intertemporal optimization problem as a benchmark. All other decision rules (i. e., any function that maps states into actions) are interpreted as rules of thumb or heuristics. In the next section, we present five such rules.

Before we analyze how rules of thumb perform relative to the benchmark solution, we conclude this section by briefly sketching how the solution to the intertemporal optimization model given by (1) – (5) can be computed. While there does not exist a closed-form solution, the optimal allocation of consumption over time is characterized by the following first-order condition:

$$u'(C_t) = \frac{1+r}{1+\rho} s_{t+1}^t E_t(u'(C_{t+1})). \quad (9)$$

This is a modified version of the well-known standard Euler equation in which next period's expected marginal utility is weighted with the conditional probability of surviving period t . From this condition, one can see that including mortality risk, via the survival probabilities $s_{t+1}^t < 1$, increases the individual's impatience. The effect is similar to increasing the rate of time preference, ρ . However, the impatience effect of mortality risk is not constant, but increases over time.

While the intuition of Euler equations such as (9) – balancing marginal utility across periods – is clear, there does not, in general, exist a closed-form solution which would allow individuals to compute their optimal consumption decision in each period. Rather, every consumer has to solve, in each decision period, the entire life-time optimization model by backward induction. As noted by many authors before (see, e. g., Hey and Dardanoni (1988), Pemberton (1993), and Rust (1994)), this procedure is computationally demanding, and we can safely assume that individuals do not actually solve this problem when making their consumption and savings decisions. Moreover, Pemberton (1993, p. 5) points out that the intuition behind the Euler equation does not help to find simpler behavioral rules that would generate *as if* behavior. We argue below that such heuristics might indeed exist.

To solve the intertemporal optimization model for the case with implicit borrowing constraints numerically, we apply the cash-on-hand approach by Deaton (1991) in the version developed by Carroll (1992). Cash on hand, denoted by X_t , is the individual's current gross wealth (total current resources), given by the sum of current income and current assets,

$$X_t = (1+r)(X_{t-1} - C_{t-1}) + Y_t. \quad (10)$$

As Deaton (1991) shows, the solution to the intertemporal optimization problem is a function of cash on hand, so we are looking for a policy function of the form $C_t = C_t(X_t)$. Trivially, the individual consumes all remaining wealth in the last period of life. For the remaining periods,

the model can be solved by backward induction starting from the last period, T . The approach we use to compute the optimal consumption path is standard in the literature in life-cycle savings decisions (e. g., Carroll (1992)). However, since our version of the model allows for both life-time and income uncertainty, our algorithm is more complicated and computationally more demanding because in our extended model, we have to keep track of survival probabilities. Rodepeter and Winter (1998) provide a detailed discussion of this algorithm.

3 Five rules of thumb for life-cycle savings decisions

In this section, we present five decision rules that allow individuals to make their life-cycle savings decisions. Such rules of thumb might be used by individuals which are either unwilling or unable to compute optimal decision rules such as those derived in the previous section. In Section 4, we use these decision rules to simulate intertemporal consumption-savings decisions. The first three rules are standard decision rules used in the savings literature, while the remaining decision rules are derived from experimental studies of saving behavior by Anderhub (1998).

Table 1: Non-optimal decision rules in a life-cycle savings model

Decision rule	Description	Based on
Benchmark	Solution to the underlying intertemporal optimization problem	Carroll (1992), Rodepeter and Winter (1998)
Rule No. 1	Consumption equals current income	Keynes (1936)
Rule No. 2	Consumption equals permanent income	Friedman (1957)
Rule No. 3	Consumption equals cash on hand up to mean income, plus 30% of excess income	Deaton (1992)
Rule No. 4	Naive intertemporal allocation based on survival probabilities	experiments by Anderhub (1998)
Rule No. 5	Naive intertemporal allocation based on the expected length of life	experiments by Anderhub (1998)

Table 1 contains an overview of all five decision rules considered in this paper. The first decision rule is the standard “consume your current income” rule by Keynes (1936). The second rule corresponds to Friedman’s (1957) “permanent income” decision rule. The third rule is taken from Deaton (1992). As we explain below, Deaton designed this rule with the explicit goal that it should be easy to compute but still match optimal behavior closely. The remaining rules are based on results from laboratory experiments conducted by Anderhub (1998). All five rules of thumb we consider are relatively easy to compute, although some might seem to be quite involved. Most importantly, however, these rules do not require using

backward induction. Each rule provides a closed-form solution for current consumption given expectations about future income (i. e., given survival probabilities and the expected path of future income).

3.1 Decision rules developed in the savings literature

Rule of thumb No. 1

The first rule of thumb we consider is the simplest rule one can think of – just consume your current income:

$$C_t = Y_t \tag{11}$$

This rule is, of course, the core of the famous consumption function by Keynes (1936), although a Keynesian consumption rule will typically allow for a constant fraction of current income to be saved in each period. By ruling out any saving, we concentrate on an extreme case. In the formal analysis of life-cycle consumption and savings decisions, this rule has been used by, *inter alia*, Hall (1978), Flavin (1981), and Campbell and Mankiw (1990). As simple as it is, this decision rule seems to be natural from a psychological perspective, see Wärneryd (1989).

Rule of thumb No. 2

The second decision rule is the permanent income rule proposed by Friedman (1957). This rule is much more complicated than the Keynesian consumption rule, but still much easier to apply than the optimal decision rule. Friedman hypothesized that consumption is a function of permanent income which is defined as that constant flow which yields the same present value as an individual’s expected present value of actual income. In Friedman’s original work, individuals use a weighted average of past income to compute permanent income. In our simulations, we impose rational expectations about future income so that we can compute permanent income based on the realizations of calibrated income processes. Specifically, we start with the identity

$$\sum_{i=t}^T Y_t^P (1+r)^{t-i} = A_t + H_t, \tag{12}$$

where Y_t^P is permanent income as of period t , A_t are current assets, and H_t is the present value of (non-asset) income given by

$$H_t = Y_t + E \left(\sum_{i=t+1}^T Y_i (1+r)^{t-i} \right). \tag{13}$$

Assuming that individuals consume their permanent income in every period, and re-arranging these identities, we obtain the “permanent income” decision rule,

$$C_t = Y_t^P = \frac{r}{1+r} \frac{1}{1 - (1+r)^{-(T-t+1)}} (A_t + H_t). \quad (14)$$

Setting the interest rate to zero for the moment, this reduces to

$$C_t = Y_t^P = \frac{1}{T-t} (A_t + H_t). \quad (15)$$

Here, one can see that individuals distribute their (expected) total wealth equally over their remaining life time, smoothing consumption, but not insuring themselves against utility losses from negative income shocks as in the life-cycle model presented in Section 2. However, individuals update their expectations about future realizations of the income process. If the stochastic component shows persistence or follows a random walk, permanent income reflects all past and current shocks.

Note that in the absence of income uncertainty (or in the case of certainty equivalence), there is no need for precautionary saving, and this rule of thumb corresponds to the solution of the underlying optimization problem (if one further ignores time preference). In the life-cycle model with income uncertainty presented in Section 2, the permanent income rule deviates from the benchmark solution. However, as Pemberton (1993) argues, this rule is both forward-looking and easy to compute. Therefore, it might be reasonable to assume such a decision rule for individuals which are “farsighted rather than myopic” and whose “concern is for ‘the future’ rather than with a detailed *plan* for the future” (p. 7, emphasis in the original). Pemberton refers to the underlying concept as “sustainable consumption”. Our simulations allow us to evaluate how such forward-looking behavior performs relative to the benchmark solution.

Rule of thumb No. 3

Deaton (1992) considers a static consumption rule which is relatively easy to compute. It is much simpler than the permanent income rule, but it is not forward-looking. Deaton assumes that individuals consume cash on hand, X_t , as long as cash on hand is less than expected income. If the income realization exceeds expected income, individuals save a constant fraction, ζ , of excess income (and consume the rest right away). Formally, Deaton’s decision rule can be written as:

$$C_t = \begin{cases} X_t & \text{if } Y_t \leq E(Y_t) \text{ and } X_t \leq E(Y_t) \\ E(Y_t) & \text{if } Y_t \leq E(Y_t) \text{ and } X_t > E(Y_t) \\ E(Y_t) + \zeta(Y_t - E(Y_t)) & \text{if } Y_t > E(Y_t) \end{cases}$$

Below, we follow Deaton in setting this fraction to 30%.⁸ Deaton explicitly states that he specified this decision rule, including the choice of $\zeta = 30\%$, entirely *ad hoc*. His goal was to approximate the solution of the underlying optimization problem (a life-cycle model similar to ours) with a rule that “should be simple, simple enough to have plausibly evolved from trial and error” (p. 257). The intriguing feature of this rule is that while being based on just easy-to-compute expected income, it approximates the optimal solution quite well in Deaton’s application. We will show below that this is also true in our slightly more involved life-cycle model. Another interesting property of this decision rule is that the corresponding savings function is always below the optimal savings function, i. e., consumption is always too high.

3.2 Decision rules derived from experiments by Anderhub (1998)

Decision rules No. 4 and No. 5 are derived from experiments on optimal savings behavior conducted by Anderhub (1998). Before we describe these rules in detail, some general remarks on these experiments are in order. The specific aim of Anderhub’s experiments was to analyze the effect of life-time uncertainty on the intertemporal allocation of consumption. Subjects were asked to distribute a given amount of tokens (i. e., wealth) over future periods. The number of these periods, interpreted as the length of life, was stochastic with a minimum of three and a maximum of six periods. The uncertainty about the total number of periods was resolved sequentially. After the last period, the allocation of tokens (i. e., consumption) in all periods was translated into final payoff using an additively separable function.⁹ Anderhub’s experiments can be taken to roughly capture the intertemporal aspects of a life-cycle saving problem with life-time uncertainty. However, we should stress that Anderhub’s goal was not to mimic a realistic life-cycle consumption-savings problem (such as the one presented in Section 3) in a laboratory situation.¹⁰

A few distinct patterns of behavior emerge in the experiments conducted by Anderhub (1998). These behavioral patterns share the property that individuals try to account for survival probabilities in some intuitive and easy-to-compute way. This seems to be a desirable property of any rule of thumb used in a life-cycle savings problem with life-time uncertainty. As noted before, there was no evidence that individuals follow a backward solution strategy in these experiments.

⁸ The simulation results we present below turned out not to be very sensitive to this calibration of ζ .

⁹ In an alternative treatment, Anderhub used a pay-off function in which period consumptions enter multiplicatively rather than additively. His results were in general robust to this variation.

¹⁰ We would generally interpret experiments such as the one described here as devices to explore how individuals form decisions when they face dynamic problems, and we would be very cautious about direct conclusions regarding the empirical validity of the standard life-cycle model.

For the purpose of this paper, we translate two of the stylized behavioral patterns found by Anderhub into decision rules that can be applied in our life-cycle savings framework. Both are forward-looking and therefore reflect “farsighted rather than myopic” behavior in sense of Pemberton (1993) without the need for backward induction. Since we consider a life-cycle savings model with both income and life-time uncertainty, this translation is conceptually straightforward, but two slight modifications of the standard model presented in Section 2 are required. First, Anderhub did not implement any mechanism that would correspond to discounting future utility. We therefore ignore discounting as well, and we set impatience (i. e., the difference between the interest rate and the rate of time preference) to zero in the benchmark calibration of the underlying intertemporal optimization model. This allows us to focus on the way in which individuals allocate their income over future periods when the length of life is uncertain. Second, also due to the design of the experiment, these rules have no explicit role for income uncertainty – they focus on life-time uncertainty. This is in contrast to decision rule No. 3 which was explicitly designed to provide individuals with a simple device for self-insurance in a world with income uncertainty but no life-time uncertainty.

Rule of thumb No. 4

The first distinct behavioral pattern we translate into a decision rule involves distributing the amount of token money which is currently available evenly across periods. This is done for all possible outcomes (i. e., length of life), and these outcomes are weighted with their *ex ante* probabilities. In our setting this translates into computing the expected value of life-time income and distributing this value over the remaining periods, weighted by the respective survival probabilities. We label this decision rule the “naive intertemporal allocation based on survival probabilities”.¹¹

Rule of thumb No. 5

The final decision rule we consider is the “naive intertemporal allocation based on the expected length of life”. It is derived from another pattern that emerged in Anderhub’s experiments. In every period, individuals compute the expected length of life and then distribute the available amount of token money evenly over these periods. In the context of our model, it translates into a permanent-income rule which takes into account survival probabilities.

¹¹ Anderhub refers to this rule as “weighted even distribution”.

3.3 Savings motives reflected by simple rules of thumb

Before we turn to simulating life-cycle savings decisions using these five rules of thumb, it is useful to briefly review the two central motives for saving in our life-cycle model. Table 2 contains an overview of how these savings motives are reflected in these five simple decision rules. It is important to recognize that decision rules differ in the motives for saving they reflect. As we will see in the next section, in some specifications of the individual’s stochastic environment, not all savings motives will be relevant. This implies that there is no universal ranking of these decision rules in terms of their usefulness to individuals.

Table 2: Savings motives captured by alternative decision rules

	Rule No. 1	Rule No. 2	Rule No. 3	Rule No. 4	Rule No. 5
Consumption smoothing	no	yes	no	yes	yes
Precautionary saving	no	no	yes	no	no

In a world with uncertainty about the length of life and stochastic income, the intertemporal optimization problem of Section 2 is designed to capture both risk aversion (i. e., consumption smoothing over time) and precautionary motives (i. e., self-insurance against negative income shocks). Of the non-optimal decision rules, Deaton’s (1992) rule is the only one which allows for a precautionary saving motive, while the Keynesian consumption rule includes no savings motive at all. The permanent income rule and the two rules derived from Anderhub’s (1998) experiments are forward-looking in these sense that individuals use their expectations about future income, and in the case of persistent shocks also information about past and current shocks, in their consumption and savings decisions. Therefore, these rules reflect the consumption smoothing motive of saving.

Finally, we should note that none of these five rules prevents consumption from falling to zero if the individual is hit by a series of bad income draws. This undesirable outcome cannot occur in the solution to the intertemporal optimization model where individuals always hold some positive wealth as an optimal response to the implicit liquidity constraint imposed by the zero income probability. In our simulations, we deal with this possibility by giving individuals a small amount of consumption for free – i. e., we implement a primitive social safety net to prevent unlucky rule-of-thumb individuals from starving to death. This induces only slight distortions in the outcome of simulated rule of thumb behavior, and our conclusions are not affected.

4 Simulation and evaluation of rule-of-thumb behavior

In this section, we present simulation results and compute the utility losses associated with using five alternative rules of thumb relative to the benchmark solution, taking preferences as given. More specifically, in order to compare utility losses across different decision rules, we compute a compensating variation measure, i. e., the additional income which would give an individual the same life-time utility under a given behavioral rule as he would obtain had he solved the underlying optimization problem. This income differential can also be interpreted as the maximum amount an individual would be willing to pay to obtain the solution to the underlying optimization problem (assuming that following the corresponding rule of thumb is costless). In a similar context, this approach to evaluating rules of thumb relative to a given intertemporal optimization problem has been used by Cochrane (1989) and Lettau and Uhlig (1999).

4.1 Simulation approach and calibration of the life-cycle model

Regarding the numerical solution and simulation of the life-cycle model with income and life-time uncertainty, we follow the approach developed by Rodepeter and Winter (1998) where further details can be found. In Table 3, we report the benchmark parameter values used to calibrate the model. In some of the simulations that follow, we also use different values of the rate of time preference, the risk aversion coefficient, the interest rate, and the standard deviations and zero-income probabilities of the income processes to illustrate how the performance of rules of thumb depends on these parameters.

As a general result, the compensating income variation which reflects the utility loss associated with rule-of-thumb behavior depends mainly on the curvature of the utility function, on impatience (the difference between the rate of time preference and the interest rate), and on income growth rates. In our benchmark calibration, the rate of time preference is equal to the interest rate (i. e., there is no impatience). Rule-of-thumb behavior leads to a higher utility loss if the individual is more risk averse (i. e., if the utility is more curved) or more impatient, and if the life-cycle income profile shows more variation in income growth rates over the life cycle.

The life-cycle income profile used in our simulations was obtained from German household-level data for the 1978–1993 period taken from the *Einkommens- und Verbrauchsstichprobe* (EVS), a dataset that is roughly comparable to the U.S. Consumer Expenditure Survey (CEX), but which has also very detailed information on various income components. The income measure used to construct the life-cycle income profiles which enter our simulations is net income. This income measure is defined as the sum of net labor income and the net balance of recurring public and private transfers. Note that this income measure excludes interest on

Table 3: Parameter values used for calibration of the life-cycle model

Parameter		Benchmark value
Relative risk aversion coefficient	γ	3
Rate of time preference	ρ	3%
Interest rate	r	3%
Conditional survival probabilities	S_t	life-table values
Number of simulation periods ^a	$T - \tau$	$115 - 20 = 95$
Standard deviation of N_t (random walk)	σ_{RW}	0.2
Standard deviation of N_t (i.i.d.)	σ_{IID}	0.5
Zero income probability ^b	p	0
Starting net labor income ^c	Y_τ	DM 27,300

^a In the case of no life-time uncertainty, we fix $T = 80$.

^b Positive probabilities for zero income are only used in our simulations of savings rule No. 3.

^c *Source:* EVS 1978–93; own calculations. In 1993 prices.

current assets and non-recurring private transfers because these income components would distort our simulations of life-cycle savings decisions.¹²

It is important to note that the long-term component of the income process includes pensions from the quite generous German pay-as-you-go system, with a replacement rate of about 70%. Thus, the loss in utility due to lower income during retirement will be relatively small even if there is no life-cycle saving at all. Utility losses from following rules of thumb are generally higher if the public pension system is less generous as in many other countries, or if there are no public pensions at all, as in the pure life-cycle model.

In the case of a stochastic income process, we obtain the deterministic component of income growth, G_t , from the empirical profiles and add simulated realizations of the stochastic component. Based on a specific realization of the income process, we then solve the entire optimization problem and compute savings and consumption decisions and the resulting period utilities over the life-cycle. From these period utilities, we obtain total life-time utility

¹² In Germany, contributions to the public pay-as-you-go pension system are mandatory for a large fraction of the population (excluding most of the self-employed, however). We treat these contributions like taxes; hence, they reduce disposable income during active working life. Symmetrically, pensions are generally treated as part of the household's income. Further details on the construction of the income variable can be found in Rodepeter and Winter (1998).

and compute the compensating variation measure. Finally, we compute the mean of the compensating variation measures for R draws of the stochastic income process.¹³

To be more precise, for a given draw we first compute life-time utility under both the optimal and the alternative decision rules; the latter is, of course, generally lower. We then increase life-time income proportionally (i. e., by a fixed percentage every year) until life-time utility under the alternative decision rule is equal to life-time utility under optimal behavior. We repeat this process for all R income draws and then compute the average compensating variation, expressed as a percentage of life-time income.

Before looking at the results in detail, we should point out that in a given version of the life-cycle model, not all of the five decision rules presented in Section 3 make sense, and we do not simulate all decision rules in all stochastic environments. For example, rule No. 3 was designed by Deaton for the case of an i.i.d. income process and fixed length of life. If Deaton's rule were used in a situation without income uncertainty, it would collapse to a permanent income rule.

In all figures that follow, we plot the percentage of additional life-time income (in short, the *compensating income variation*) required to compensate the individual for the utility loss associated with making consumption and savings decisions according to some rule of thumb rather than solving the intertemporal optimization problem. To build some intuition about the mechanics of rule-of-thumb behavior in the life-cycle model, we first present simulation results for models without income uncertainty and then turn to stochastic income processes.

4.2 Models without income uncertainty

The most straightforward consumption rule is to just consume current income in every period; this is the Keynesian consumption rule (decision rule No. 1). In the most simple case with no uncertainty about the length of life or future income, the loss in total life-time utility resulting from rule-of-thumb behavior depends only on the growth rate of income (relative to the time preference rate) and on the curvature of the utility function. Figure 1 shows the compensating income variation as a function of the coefficient of relative risk aversion, γ , when the individual consumes all income in every period and does not save at all. In our benchmark calibration, this relationship is monotonically increasing. This is of course the standard textbook result: A higher risk aversion coefficient is equivalent to a lower elasticity of intertemporal substitution, and thus the utility loss from not smoothing consumption over time will be higher. The amount of additional life-time income required to compensate for not saving at all is substantial. For the benchmark risk aversion coefficient of 3, more than a

¹³ The number of income draws used in our simulations is $R=10,000$. Results are not sensitive to increasing the number of draws further, and for most decision rules, much fewer repetitions proved sufficient to obtain stable results.

10% increase in life-time income is required to make up the utility loss due to rule-of-thumb behavior.

Note that in the baseline case, both the interest rate and the rate of time preference are set to 3%, i. e., there is no impatience. As can be seen from Figure 1, if the individual is impatient in the sense that the time preference rate is higher than the interest rate, the compensation variation of life-time income is generally higher. This result might be surprising at first sight because the utility loss from lower consumption during retirement is more heavily discounted if the individual is impatient, and so the life-time utility loss should be lower. However, due to fact that life-cycle income is increasing both in age and over time during the active working life, there are also strong intertemporal effects which work in the opposite direction. The shape of the income profile also explains the fact that with impatience, the life-time utility loss is initially decreasing as a function of the risk aversion coefficient. These effects will also be at work in other calibrations and for other decision rules, and we will not comment on these any more.

When we introduce uncertainty about the length of life, the utility loss from following a simple Keynesian consumption rule increases substantially even in the absence of any uncertainty about income itself; see Figure 2. This is of course due to the fact that individuals now face the risk of living longer, hence there is the risk of having to cope longer with lower pension income and no savings which leads to lower consumption levels and to lower life-time utility. The effects of variations in the risk aversion rate, the interest and time preference rates are similar as before.

Next, we consider a simple decision rule which allows for saving. Rule No. 2 says that an individual should only consume his permanent income, hence during working life, when income is above permanent income, individuals will save. Following such a rule leads to a smooth consumption path, and when the rate of time preference is equal to the interest rate, the outcome will be identical to the solution of the intertemporal optimization problem. However, as this rule has no role for time preference *per se*, the individual suffers a utility loss relative to the maintained optimization model if there is a wedge between time preference and the interest rate. The corresponding income variation is an increasing function of the difference between interest rate and time preference rate, see Figure 3. For small differences, say up to three percentage points as in many simulation models in the literature, the utility loss is small. For this rule, we do not show results for the version with uncertain length of life; the effects are similar as in the case of the Keynesian consumption rule.

Finally, we check how individuals who follow one of the decision rules that have been found in the experiments by Anderhub (1998) do in terms of their life-time utility relative to optimizers. For these rules, it makes only sense to consider the case of uncertainty about length of life. We find that these rules lead to different savings profiles over the life cycle, but they both imply

some consumption smoothing that takes account of life-time uncertainty. Again, these rules do not allow for time preference, hence the utility loss relative to the underlying optimization model is a function of the difference between time preference and interest rates. In our simulations (see Figure 4), we found that if this difference is about two percentage points, these rules come quite close to optimal behavior. Even for more substantial differences, the utility loss from following these rules of thumb is relatively small when compared to rules No. 1 and No. 2.

At this point, we should once again stress that our translation of Anderhub's experimental findings into decision rules which can be used in our setting was to some extent guided by our knowledge of the structure of the underlying model. We do not claim that individuals use these rules in real-life savings decisions literally, but the strategies they follow might be similar. Our central finding is that while these rules of thumb are relatively simple in the sense that they do not require backwards induction, they result in small losses of life-time utility relative to the solution of the intertemporal optimization problem. These losses are smaller than those associated with a Keynesian consumption rule or a version of the permanent income rule, in particular if we allow for impatience.

4.3 Models with income uncertainty

In this section, we investigate how rules of thumb perform in a world with income uncertainty. Once we introduce income uncertainty in the life-cycle model, precautionary saving arises as a second saving motive (in addition to consumption smoothing). We distinguish two polar cases: one in which income shocks are i.i.d. and thus purely transitory and one in which income follows a random walk so that shocks have permanent effects. After we have analyzed the effects of variations in the preference parameters (risk aversion, rate of time preference) in the previous section, we fix these now at the values reported in Table 3 and concentrate on characterizing the compensating income variation as a function of the variance of the income process. This will help to understand to what extent following simple rules of thumb exposes individuals to income risk.

We should stress that the empirical evidence on the stochastic properties of income processes is mixed; in particular, it is not clear whether income should be modeled as either an i.i.d. or a random walk process. We choose these extremes for illustrative purposes, and our main interest is the effect of increasing income uncertainty. Other things equal, realistic i.i.d. or random walk income processes will differ in the standard deviation of the underlying income shock. Therefore, the compensating variation income measures reported in this section should not be compared across specifications for a given value of the standard deviation.

As a final remark before we present our results, note that the standard version of the life-cycle model presented in Section 2 allows for zero income. This allows to account for extremely

negative income shocks such as an unemployment spell, and it is also a technically convenient way to introduce liquidity constraints. With the exception of rule No. 3, all decision rules we consider do not allow for precautionary saving, and even rule No. 3 does prevent consumption falling to zero. While this shows that in a world with income uncertainty, following a rule that does not allow for precautionary saving is surely sub-optimal, we want to exclude this case in the analysis for the moment. We therefore set the zero-income probability to zero in the benchmark calibration and analyze this possibility only when we simulate rule No. 3.

We do not present results for the simple Keynesian decision rule in which consumption equals income (rule No. 1) because it does not allow for saving and thus the individual can neither smooth consumption nor self-insure against adverse income shocks. When income is uncertain, the utility loss of following such a rule is even larger than shown in Figure 2, and the amount of additional life-time income required to yield the same total utility as optimal behavior would be of course increasing in the variance of the income shock.

In the case of the permanent income rule (rule No. 2), the intuition is similar. Results are shown in Figure 5. Again, this rule does not directly allow for precautionary saving. However, in the case of persistent income shocks, each shock alters the expected time path of future income and does therefore change permanent income and consumption. Therefore, the permanent income rule does somewhat better than a Keynesian rule. Just as in a world with certain income, it does also better because it is forward-looking. Under uncertainty, this means that such a rule takes into account the deterministic part of the life-cycle income profile. For both i.i.d. and random walk income processes, the life-time utility loss increases in the variance of the underlying shock, although as mentioned before, utility losses should not be compared directly for a specific value of the variance.

Next, recall that Deaton's (1992) savings rule (rule No. 3) was specifically designed as a simple rule of thumb that would yield good results in situations with income uncertainty. Because this rule does not require computing any permanent income or some other measure of expected future income, it is relatively simple to follow and might be considered as a true rule of thumb (but as noted before, it is not forward-looking). We restrict our attention to the case of an i.i.d. income process as in Deaton's original work; the effects of allowing for persistent shocks are similar to the other cases considered so far.

It should not come as a surprise that this rule yields relatively poor results when the variance of income is low. In the extreme, for zero variance, it collapses to the Keynesian consumption rule with no saving at all. This can be seen from Figure 6, where once again the compensating variation measure of life-time utility loss is plotted against income variance. However, Deaton's rule allows individuals to self-insure against stochastic income shocks, and therefore it does quite well when the variance of the income process is increased – the compensating income variation is essentially flat as a function of the income variance as long as we ignore the risk of

zero income. When we introduce zero-income risk, the life-time utility loss of rule-of-thumb behavior is substantially higher, as can be seen from Figure 6.

Finally, we consider savings rules No. 4 and No. 5. Recall that these rules have been designed to capture the way in which individuals deal with uncertainty about planning horizons in laboratory experiments; they do not specifically address income uncertainty. It might therefore be expected that the results for these rules are similar to those obtained for the permanent income rule No. 2. As can be seen from Figures 7 and 8, life-time utility losses increase in the variance of the income process. In particular, rule No. 5 implies that individuals follow some sort of permanent income rule, so the shape of the relationship between the compensating income variation and income variance starts is similar to that obtained for rule No. 2 in Figure 5. It begins at a value of 5.3% which is the value obtained in Figure 4 for the case of certain income and a zero difference between interest rate and rate of time preference. In contrast, Rule No. 4 does not have a direct role for permanent income and does much worse if income follows a random walk and shocks are persistent. For very low values of the income variance, this rule does better, starting at 3.5% which again corresponds to the value obtained earlier in Figure 4.

4.4 What makes a rule of thumb perform well?

Based on our simulation results, we can draw a preliminary conclusion. In the case of income certainty, life-time utility losses resulting from following some rule of thumb rather than solving the underlying intertemporal optimization problem are relatively small. When income is uncertain, individuals who follow rules of thumb in their consumption and savings decisions suffer considerable utility losses relative to the optimal decision rule. Not surprisingly, the magnitudes of these utility losses depend on preference parameters and the specific structure of the income process. There are also many cases in which rules of thumb do not imply substantial utility losses, and rules of thumb which are simpler than others (such as Deaton's rule) do not necessarily perform worse.

In Table 4, we provide a systematic comparison of the relative performance of the five decision rules considered in this paper. As discussed before, the structure of the decision rules requires some compromise in the specification of the underlying optimization models. All results in this table are based on the benchmark parameters of Table 3 to allow a fair comparison of the five rules of thumb. Note that while we use the same preference parameters in all specifications, we use different variances for i.i.d. and random walk processes. However, the utility losses reported in this table can be directly compared for a given stochastic process, i. e., along the rows of this table. We should emphasize that these findings are sensitive to the parameter values used in the calibration, but the results in the preceding sections allow to assess how a variation of any of these parameters affects the performance of rule-of-thumb behavior.

Table 4: Life-time utility loss from using non-optimal decision rules

Income process	Length of life	Rule No. 1	Rule No. 2	Rule No. 3	Rule No. 4	Rule No. 5
certain	certain	10.7	0.0			
certain	uncertain	7.2			3.3	5.3
random walk	certain	23.1	7.2			
random walk	uncertain				17.7	12.5
i.i.d.	certain	12.8	1.5	5.0		
i.i.d.	uncertain				4.2	6.8

Note: Life-time utility losses are expressed as the percentage of additional life-time income that would compensate an individual for using non-optimal decision rules rather than solving the corresponding intertemporal optimization problem. Parameter values used for calibrating the model are reported in Table 3.

Source: Own calculations.

The main conclusion from this comparison is that looking both across columns (i. e., comparing the performance of a given decision rule in alternative stochastic environments) and across rows (i. e., comparing alternative stochastic environments themselves), there is considerable variation in the life-time utility loss associated with using rules of thumb. There is no uniformly best rule of thumb, and for most stochastic environments analyzed in this paper, there is some rule of thumb which yields relatively small utility losses (less than 10% of life-time income). In the case of uncertain length of life and a random walk income process, however, utility losses are substantial for all rules of thumb.

Based on these results, we conclude that the key factor that makes a rule of thumb successful is its ability to generate a measure of life-time income that correctly reflects movements in future income. If the life-time income process exhibits a strong deterministic trend and modest shocks with low persistence, this might not be too difficult. We discuss some implications of this finding in the concluding section.

5 Conclusions

Although the predictions of the benchmark life-cycle savings model – such as a hump-shaped savings profile with positive saving during the active employment and dissaving during old age – have been frequently rejected in household level data, the standard model seems still to be widely accepted among applied researchers. This might be due to the fact that many of the well-known failures of the standard model can be attributed to institutional arrangements which are either not properly reflected in simple versions of the model, or lead to measurement

problems in its empirical analysis.¹⁴ These problems are important in practical applications of the life-cycle model, but the more fundamental question is whether the underlying assumption of intertemporal optimization is a valid approximation to real-world decisions.

In this paper we argued that while individuals, most likely, do not solve such highly complicated models, they might actually behave *as if* they were: Using much more tractable heuristics, or rules of thumb, results only in modest utility losses. Simulating life-cycle consumption and savings decisions based on five decision rules, we found that losses in total life-time utility compared with optimal behavior (i. e., using the solution of the underlying intertemporal optimization problem for given preferences) can, in general, be substantial. However, the magnitudes of these losses vary with the assumptions about preference parameters and the properties of the income process. For those specifications we tested, there always exists some simple rule of thumb which results in only modest utility losses. It is, of course, difficult to assess whether utility losses equivalent to between 5% and 10% of life-time income are small or large, but given that individuals seem to be unable to compute the optimal solution anyway, following much simpler rules with utility losses of the magnitudes we found here does not seem too bad.

Given that the savings rules we analyzed reflect two distinct motives for saving (consumption smoothing and precautionary saving), a natural extension of our approach would be to combine rules of thumb. For example, individuals could use a simple static rule such as Deaton's rule (our rule No. 3) to insure themselves against adverse income shocks, and at the same time, they could do some consumption smoothing by using forward-looking rules, i. e., the permanent income rule (rule No. 2) which focuses on income uncertainty and/or either one of rules No. 4 and No. 5 which focus on life-time uncertainty. By combining several rules of thumb, individuals should be able to improve their total life-time utility considerably, and they might actually come quite close to using the solution to the underlying optimization problem. The concept of mental accounting introduced by Shefrin and Thaler (1988) suggests that individuals actually make savings decisions with different time horizons and savings goals in mind. Combining rules of thumb such as those analyzed in this paper with a mental accounting framework would therefore seem a fruitful direction for future research. Unfortunately, such combinations of savings rules would complicate matters considerably, both technically and conceptually. First, the optimization and simulation problem is obviously much more involved, and second, and more importantly, one would have to make assumptions of how individuals

¹⁴ An important example is social security which crowds out private saving at least partially but is hard to incorporate in empirical models. For the case of Germany, see Schnabel (1999) and Rodepeter and Winter (1998).

allocate funds to different savings rules.¹⁵ This would imply that a second class of behavioral decision rules would have to be specified for asset allocation.

An important question which we have not addressed in this paper is: How do rules of thumb actually arise? How do individuals decide which behavioral decision rule they use? In our analysis, we have taken the rules of thumb as exogenously given because our main objective was to evaluate the utility loss associated with using some heuristic rather than computing the optimal solution to the underlying decision problem. We did not model the choice between the optimal strategy and using rules of thumb. There is, of course, a methodological problem here: One can always set up some meta decision problem which addresses this question formally, but this would lead to an infinite regress. A promising approach is to explore how rules of thumb arise endogenously from learning behavior; a recent example is a paper by Lettau and Uhlig (1999) who investigate a model of learning rules of thumb in intertemporal decision problems. However, in a life-cycle savings setting, learning from own mistakes is impossible. Every life-cycle decision is made only once – all decisions are conditional on the planning period (i. e., on age) and cannot be repeated with a different “trial” decisions in the future. Therefore, learning models need to take account of social interactions (see Ellison and Fudenberg (1996) for a formal treatment). Such models would imply that savings decisions are based at least partially on imitation of other individuals’ (family, neighbors, or friends) behavior. To our knowledge, there is no strong empirical evidence on this issue, but this should be an important direction for future empirical research.

At this point, we should also point out that an empirical analysis of rule-of-thumb savings behavior would have to account for the possibility that individuals are heterogeneous with respect to the decision rules they use. Once we give up the fiction of optimal behavior based on the solution to a (unique) underlying optimization problem, the result that all individuals follow the same decision rule does not need to hold any more. Some individuals might care more about short-term precautionary saving, others for long-term consumption smoothing, and, as noted above, combinations might arise as well. To our knowledge, there exist no econometric studies that try to identify decisions rules in a life-cycle savings context, but this is clearly an important area for future research.

Finally, an important part of life-cycle savings decisions is pre-determined by a combination of imitation and institutional arrangements. For example, in countries such as the U.S. and the U.K., many households buy their first family homes in similar phases of their life cycle, and this decision determines a large fraction of their consumption and savings pattern over future years. Similarly, due to its favorable tax treatment, the acquisition of life-insurance policies with substantial savings components during the early stage of the active working life is quite

¹⁵ Meier (1997) analyzes mental accounting in a very stylized simulation model of savings behavior; the allocation of funds across accounts in his model is, however, static.

common in Germany (see Walliser and Winter (1999)). The acquisition of a family home or a life-insurance policy is a one-time decision which might well be influenced by social learning, and it fixes a substantial part of life-cycle saving, reducing the scope for discretionary saving over remaining years substantially. Such behavior, based on direct imitation, social traditions or institutional arrangements, can be interpreted as following a rule of thumb, and it might result in observed behavior that can be quite close to optimal life-cycle savings behavior.

Based on our results on the performance of rules of thumb, we would argue that the most important direction for understanding life-cycle savings behavior is how social learning and institutional factors lead to decision rules which result in outcomes that are close to the solution of the underlying intertemporal optimization problem. As we have shown in this paper, there might be a whole variety of quite different savings rules that meet this criterion.


References

- Anderhub, V. (1998):** Savings decisions when uncertainty is reduced: An experimental study. Discussion Paper No. 98-73, Sonderforschungsbereich 373, Humboldt-Universität zu Berlin.
- Browning, M. and A. Lusardi (1996):** Household saving: Micro theories and micro facts. *Journal of Economic Literature*, 34(4), 1797–1855.
- Caballero, R. J. (1990):** Consumption puzzles and precautionary savings. *Journal of Monetary Economics*, 25, 113–136.
- Camerer, C. (1995):** Individual decision making. In J. H. Kagel and A. E. Roth (Eds.), *Handbook of Experimental Economics*, 587–703. Princeton, NJ: Princeton University Press.
- Campbell, J. Y. and G. N. Mankiw (1990):** Permanent income, current income, and consumption. *Journal of Business and Economics Statistics*, 8(3), 265–279.
- Carbone, E. and J. D. Hey (1997):** How do people tackle dynamic decision problems? Unpublished manuscript, University of York.
- Carroll, C. D. (1992):** The buffer stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity*, No. 2, 61–135.
- Carroll, C. D. (1997):** Buffer-stock saving and the life cycle-permanent income hypothesis. *Quarterly Journal of Economics*, 112(1), 1–55.
- Cochrane, J. H. (1989):** The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives. *American Economic Review*, 79(3), 319–337.
- Deaton, A. (1991):** Saving and liquidity constraints. *Econometrica*, 59(5), 1221–1248.
- Deaton, A. (1992):** Household saving in LDCs: Credit markets, insurance and welfare. *Scandinavian Journal of Economics*, 94(2), 253–273.
- Ellison, G. and D. Fudenberg (1996):** Rules of thumb for social learning. *Journal of Political Economy*, 104, 612–643.
- Flavin, M. A. (1981):** The adjustment of consumption to changing expectations about future income. *Journal of Political Economy*, 89(5), 974–1009.
- Friedman, M. (1957):** *A Theory of the Consumption Function*. Princeton, NJ: Princeton University Press.
- Glaeser, E. L. and A. L. Paulson (1997):** A reconsideration of the (in)sensitivity of tests of the intertemporal allocation of consumption to near rational alternatives. *Journal of Economic Dynamics and Control*, 21, 1173–1180.
- Hall, R. E. (1978):** Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86, 971–987.

- Hey, J. D. (1999):** Do people (want to) plan? Unpublished manuscript, Università di Bari and University of York.
- Hey, J. D. and V. Dardanoni (1988):** Optimal consumption under uncertainty: An experimental investigation. *Economic Journal*, 98, S105–S116.
- Hindy, A., C. H. and S. H. Zhu (1997):** Optimal consumption and portfolio rules with durability and habit formation. *Journal of Economic Dynamics and Control*, 21, 525–550.
- Ingram, B. F. (1990):** Equilibrium modeling asset prices: Rationality versus rules of thumb. *Journal of Business and Economics Statistics*, 8(1), 115–125.
- Keynes, J. M. (1936):** *The General Theory of Employment, Interest and Money*. London: Macmillan.
- Krusell, P. and A. A. Smith, Jr. (1996):** Rules of thumb in macroeconomic equilibrium: A quantitative analysis. *Journal of Economic Dynamics and Control*, 20, 527–558.
- Laibson, D. (1998):** Life-cycle consumption and hyperbolic discount functions. *European Economic Review*, 42(3-5), 861–871.
- Lettau, M. and H. Uhlig (1999):** Rules of thumb versus dynamic programming. *American Economic Review*, 89(1), 148–174.
- Loewenstein, G. (1992):** The fall and rise of psychological explanation in the economics of intertemporal choice. In G. Loewenstein and J. Elster (Eds.), *Choice Over Time*. New York: Russell Sage Foundation.
- McFadden, D. F. (1999):** Rationality for economists. *Journal of Risk and Uncertainty*, 19(1/3), 73–105.
- Meier, M. (1997):** *Das Sparverhalten der privaten Haushalte und der demographische Wandel: Makroökonomische Auswirkungen*. Frankfurt: Peter Lang.
- Moon, P. and A. Martin (1990):** Better heuristics for economic search: Experimental and simulation evidence. *Journal of Behavioral Decision Making*, 3, 177–193.
- Pemberton, J. (1993):** Attainable non-optimality or unattainable optimality: A new approach to stochastic life cycle problems. *Economic Journal*, 103, 1–20.
- Rabin, M. (1998):** Psychology and economics. *Journal of Economic Literature*, 36(1), 11–46.
- Rodepeter, R. and J. K. Winter (1998):** Savings decisions under life-time and earnings uncertainty: Empirical evidence from West German household data. Discussion Paper No. 98-58, Sonderforschungsbereich 504, Universität Mannheim.
- Rust, J. (1994):** Do people behave according to Bellman’s principle of optimality? Unpublished manuscript, University of Wisconsin–Madison.
- Rust, J. and C. Phelan (1997):** How Social Security and Medicare affect retirement behavior in a world of incomplete markets. *Econometrica*, 65(4), 781–831.

- Schnabel, R. (1999):** The golden years of social security: Life-cycle income, pensions and savings in Germany. Discussion Paper No. 99-40, Sonderforschungsbereich 504, Universität Mannheim.
- Shefrin, H. M. and R. H. Thaler (1988):** The behavioral life-cycle hypothesis. *Economic Inquiry*, 26, 609–643.
- Shi, S. and L. G. Epstein (1993):** Habits and time preference. *International Economic Review*, 34(1), 61–84.
- Simon, H. A. (1955):** A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Thaler, R. H. (1994):** Psychology and savings policies. *American Economic Review, Papers & Proceedings*, 84(2), 186–192.
- Walliser, J. and J. K. Winter (1999):** Tax incentives, bequest motives and the demand for life insurance: Evidence from Germany. Discussion Paper No. 99-28, Sonderforschungsbereich 504, Universität Mannheim.
- Wärneryd, K.-E. (1989):** On the psychology of saving: An essay on economic behavior. *Journal of Economic Psychology*, 10, 515–541.
- Zeldes, S. P. (1989):** Optimal consumption with stochastic income: Deviations from certainty equivalence. *Quarterly Journal of Economics*, 104, 275–298.


Figure 1: Compensating income variation: income certain, life-length certain, rule No. 1



Note: Additional life-time income required to compensate total life-time utility loss from following rule-of-thumb No. 1 („Consumption equals current income“) plotted as a function of the coefficient of relative risk aversion, γ .

Source: Own calculations. (D 32)


Figure 2: Compensating income variation: income certain, life-length uncertain, rule No. 1



Note: Additional life-time income required to compensate total life-time utility loss from following rule-of-thumb No. 1 („Consumption equals current income“) plotted as a function of the coefficient of relative risk aversion, γ .

Source: own calculations. (D 33)


Figure 3: Compensating income variation: income certain, life-length certain, rule No. 2



Note: Additional life-time income required to compensate total life-time utility loss from following rule-of-thumb No. 2 („Consumption equals permanent income“) plotted as a function of the difference between the interest rate and the time preference rate, $r - \rho$

Source: own calculations. (D 34)

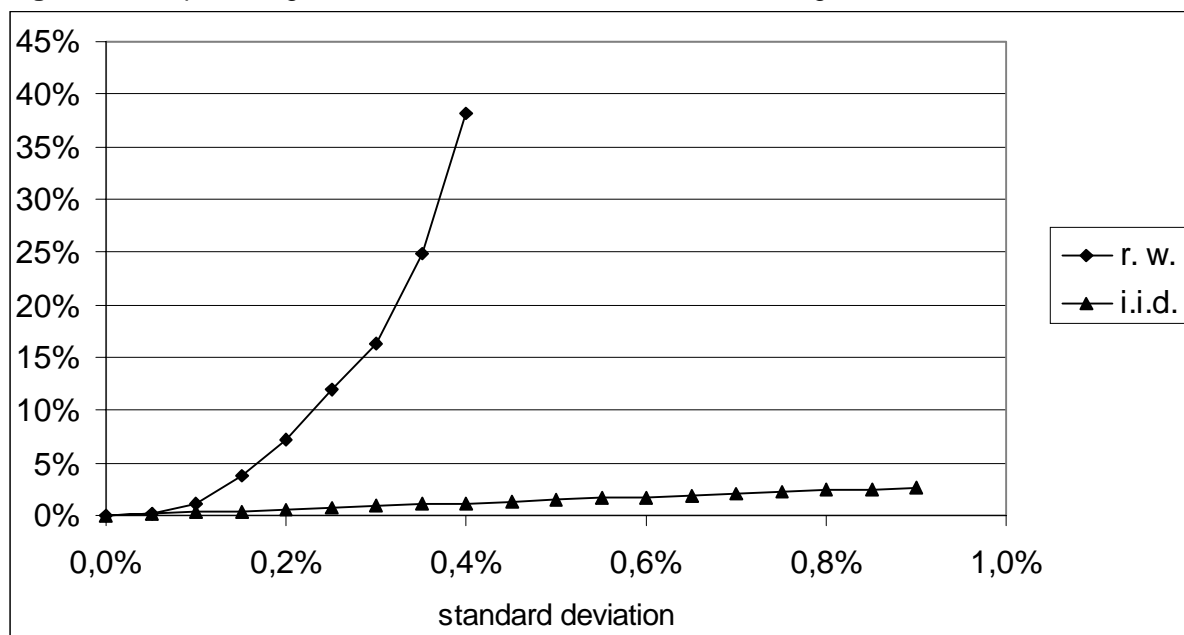
Figure 4: Compensating income variation: income certain, life-length uncertain, rule Nos. 4 & 5



Note: Additional life-time income required to compensate total life-time utility loss from following rule-of-thumb No. 4 and 5 („Probability weighted intertemporal allocation rule“ and „Expected length of life rule“) plotted as a function of the difference between the interest rate and the time preference rate, $r - \rho$

Source: own calculations. (D 35)

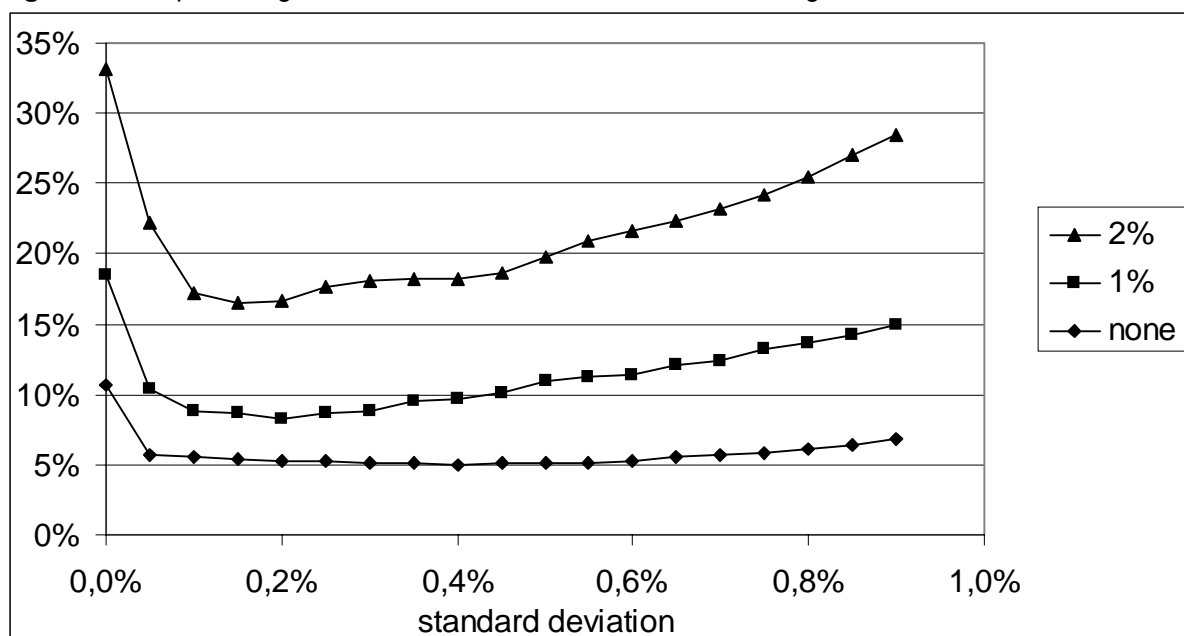
Figure 5: Compensating income variation: income uncertain, life-length certain, rule No. 2



Note: Additional life-time income required to compensate total life-time utility loss from following rule-of-thumb No. 2 („Consumption equals permanent income“) plotted as a function of the standard deviation of the income shock

Source: own calculations, $\gamma = 3$, $r = \rho = 3\%$. (D 38)


Figure 6: Compensating income variation: income uncertain, life-length certain, rule No. 3



Note: Additional life-time income required to compensate total life-time utility loss from following rule-of-thumb No. 3 („Consumption equals cash on hand up to mean income, plus 30% of any excess income“) plotted as a function of the standard deviation of the income shock

Source: own calculations, $\gamma = 3$, $r = \rho = 3\%$. (D 37)


Figure 7: Compensating income variation: income uncertain, life-length uncertain, rule No. 4



Note: Additional life-time income required to compensate total life-time utility loss from following rule-of-thumb No. 4 („Probability weighted intertemporal allocation rule“) plotted as a function of the standard deviation of the income shock

Source: own calculations, $\gamma = 3$, $r = \rho = 3\%$. (D 40)

Figure 8: Compensating income variation: income uncertain, life-length uncertain, rule No. 5



Note: Additional life-time income required to compensate total life-time utility loss from following rule-of-thumb No. 5 („Expected length of life rule“) plotted as a function of the standard deviation of the income shock

Source: own calculations, $\gamma = 3$, $r = \rho = 3\%$. (D 39)