



LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

VOLKSWIRTSCHAFTLICHE FAKULTÄT



Winter, Joachim; Schlafmann, Kathrin und Rodepeter,
Ralf:

Rules of Thumb in Life-Cycle Saving Decisions

Munich Discussion Paper No. 2011-16

Department of Economics
University of Munich

Volkswirtschaftliche Fakultät
Ludwig-Maximilians-Universität München

Online at <https://doi.org/10.5282/ubm/epub.12334>

Rules of Thumb in Life-Cycle Saving Decisions*

Joachim K. Winter

University of Munich

Munich Center of the Economics of Aging

ifo Institute for Economic Research

Kathrin Schlafmann

University of Munich

Ralf Rodepeter

BMW Museum, Munich

29 September 2011

Abstract: We analyse life-cycle saving decisions when households use simple heuristics, or rules of thumb, rather than solve the underlying intertemporal optimization problem. We simulate life-cycle saving decisions using three simple rules and compute utility losses relative to the solution of the optimization problem. Our simulations suggest that utility losses induced by following simple decision rules are relatively low. Moreover, the two main saving motives reflected by the canonical life-cycle model – long-run consumption smoothing and short-run insurance against income shocks – can be addressed quite well by saving rules that do not require computationally demanding tasks such as backward induction.

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

JEL classification: D91; E21

* Corresponding author: Joachim Winter, Department of Economics, University of Munich, Ludwigstr. 33, D-80539 Munich, Germany. Email: winter@lmu.de.

This paper is a substantially revised and extended version of unpublished research by Rodepeter and Winter (1999). We would like to thank Michael Adam, Axel Börsch-Supan, Thomas Crossley, Angelika Eymann, Werner Güth, Silke Januszewski, Ronald Lee, Alexander Ludwig, Annamaria Lusardi, Daniel McFadden, Matthew Rabin, Paul Ruud, Daniel Schunk, the special issue editor (Rachel Griffith) and referees, and participants at numerous seminars and conferences for helpful discussions and comments. We gratefully acknowledge financial support by the Deutsche Forschungsgemeinschaft (DFG) through SFB 504 at the University of Mannheim (Rodepeter and Winter) and through GRK 801 at the University of Munich (Schlafmann). Kathrin Schlafmann also gratefully acknowledges support by LMU Mentoring.

1 Introduction

Much recent research on households' saving and investment behaviour has focused on financial literacy, that is, on individuals' knowledge of such fundamental financial concepts as compound interest, and their ability to apply such knowledge when making financial decisions; see Lusardi and Mitchell (2007) and van Rooij *et al.* (2011), *inter alia*. In this paper, we take a different but related perspective on households' savings decisions. We start from the – at least to us – natural assumption that even financially sophisticated individuals do not solve intertemporal optimization models when they make their saving and investment decisions. Rather, they often use simple or sophisticated decision rules, which we call rules of thumb in this paper. The question we ask is: How large is the utility loss that households incur when they use such rules of thumb rather than solve an intertemporal optimization problem?

It is a well-known finding from psychological research on decision-making that individuals use heuristics, or rules of thumb, in making judgements and decisions. In economics, rule-of-thumb behaviour has been recognised as an important aspect of bounded rationality since the seminal work by Simon (1955).¹ Life-cycle consumption and saving 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. It requires backward induction, and no closed-form solution for current consumption as a function of the relevant state variables exists. Many authors argue 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 (2005).

Standard economic theory is based on the notion that if individuals have preferences over all possible states of nature at the current and any future date, and if their behaviour is time consistent, there exists some intertemporal utility function that individuals maximise. The standard approach in the literature on households' life-cycle behaviour 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 that corresponds to these intertemporal preferences.

¹ For surveys of this literature, see Camerer (1995), Conlisk (1996) and Rabin (1998). Gigerenzer and Todd (1999) provide a psychological perspective on bounded rationality and heuristics.

² There are many papers which depart from this standard model. An important example is the literature on hyperbolic discounting which departs from additive separability and exponential discounting; see Laibson (1998).

This problem is well understood, and the standard model of life-cycle saving serves as a powerful tool in applied research and policy analysis.³

Many important questions remain open, however. Do individuals behave according to the solution of an intertemporal optimization problem? If not, how do individuals make intertemporal choices? Further, if the assumptions of the standard life-cycle model are not empirically warranted, can it nevertheless deliver predictions that are valid in practical applications? There is a large and still growing empirical literature that addresses these issues from different perspectives.⁴ It is our reading of this literature that the question whether rational behaviour 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 behaviour is frequently rejected. In the context of life-cycle models, a series of experimental studies test whether individuals perform backward induction in cognitive tasks that involve some dynamic trade-off, see Johnson *et al.* (2001), Hey and Dardanoni (1988), and Carbone and Hey (1999). In their experiments, backward induction, and hence rational behaviour, is strongly rejected. In experimental studies of search models (which is an intertemporal decision task slightly different from life-cycle decision making, but more akin to experimental study), Moon and Martin (1990) and Houser and Winter (2004) find that individuals use heuristics that are quite close to optimal, but still different from optimal search rules (which have to be computed by backward induction). Additional evidence on the prevalent use of heuristics in intertemporal or dynamic decision problems comes from a series of experiments by Hey and Lotito (2009), Hey and Panaccione (2011), and Hey and Knoll (2010). More closely related to the topic of this paper, Anderhub *et al.* (2000) and Müller (2001) document that individuals use relatively sophisticated heuristics, but do not use backwards induction, in experimental studies of a simple saving task. Finally, Binswanger and Carman (2011) present survey data on the use of rules of thumb in actual household decisions. They show that three types of households – planners, rule-of-thumb users, and “unsystematic” savers – each make up about a third of their sample. Interestingly, the wealth outcomes of rule-of-thumb users are similar to those of planners.

In our analysis, we take a normative perspective, maintaining the assumption that individuals have standard intertemporal preferences. Our point of departure is the observation that individuals have limited computational capabilities. We characterise individuals who

³ Browning and Lusardi (1996), Browning and Crossley (2001), and Attanasio and Weber (2010) review this vast literature. Low *et al.* (2010) present a recent, sophisticated version of this approach that encompasses wage and employment risk. Blundell *et al.* (2008) study empirically whether households insure against income shocks; they document partial insurance for permanent shocks and almost complete insurance of transitory shocks among U.S. households.

⁴ For reviews of the literature on choice over time, see Loewenstein (1992), Camerer (1995), Rabin (1998), and Hey (2005).

follow simple heuristics, or rules of thumb, rather than use the decision rule given by the solution to the dynamic optimization model that corresponds to maximizing their preferences. More specifically, we compute life-cycle saving decisions under three different exogenously specified saving rules and compare the outcomes with the optimal solution. The criterion used for this comparison is a consumption equivalent: We express losses in life-time utility associated with rule-of-thumb behaviour in terms of additional consumption required to give individuals the same utility they would achieve if they behaved optimally. The life-cycle model we use is characterised by both income and life-time uncertainty. Both the model and the approach used to solve the dynamic optimization problem numerically are standard in the literature, following Deaton (1991) and Carroll (1992, 1997). Our approach is related to earlier work on near-rational behaviour in intertemporal consumption and saving problems by Cochrane (1989) and Lettau and Uhlig (1999).

Rules of thumb have been analysed in life-cycle saving models before, in particular in tests of the life-cycle/permanent income hypothesis in the macroeconomics literature. Starting with the seminal paper by Hall (1978), a series of studies assume that a fraction of the population behaves according to some simple rule of thumb such as “just consume your current income in every period” while the rest of the population behaves optimally.⁵ Estimates of the fraction of rule-of-thumb consumers in the population range between zero and well above 50 percent and depend heavily on assumptions about households’ preferences and econometric estimation approaches. In this paper, we explore such a simple consumption-equals-income rule and two other saving rules that have been used in the economics literature on life-cycle saving behaviour. More recently, Scholz *et al.* (2006) studied saving behaviour using data from a sample of older American households. They show that observed saving decisions are closer to the solution of an intertemporal optimization model than they are to two simple rules of thumb (one assumes a constant saving rate while the second is based on age and income specific average saving rates). Also related to our research question is a recent study by Calvet *et al.* (2007) who quantify the welfare losses that result from sub-optimal portfolio choices (which are perhaps driven by rules of thumb as well) using Swedish data.

We should point out that we use the term rule of thumb for *any* decision rule that is (i) not the solution to an underlying utility maximization problem and (ii) easy to derive and apply for individuals with limited computational capabilities. This terminology is consistent with the use in behavioural economics and psychology, as discussed by Goodie *et al.* (1999). In most of the economics literature, the term “rule of thumb” is typically used in a much more narrow sense, namely, for households that spend a fixed fraction or all of their income in every period; see Deaton (1992b) and Browning and Crossley

⁵ The macroeconomics literature on rules of thumb includes papers by Flavin (1981), Hall and Mishkin (1982), Cochrane (1989), Campbell and Mankiw (1990), and Weber (2000, 2002).

(2001).⁶ This is only one of the rules of thumb we consider. As will become clear below, the other rules of thumb that we analyse are forward-looking and therefore much closer to the life-cycle framework than the “fixed consumption” rule of thumb usually postulated in economics.

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

2 The benchmark life-cycle saving model

We assume that an individual’s or household’s optimal life-cycle consumption and saving behaviour can be derived from a well-defined intertemporal optimization problem, given additively separable preferences with constant exponential discounting. We use a version of the standard life-cycle model with borrowing constraints and both life-time and income uncertainty which 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 usually refer to individual decisions.

Households are assumed to maximise, 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 of being alive in period t conditional on being alive in period τ . To simplify notation, we also use a binary random variable that indicates whether an individual is alive in period t :

$$S_t = \begin{cases} 1 & \text{if the individual is alive in period } t \\ 0 & \text{if the individual is not alive in period } t \end{cases}$$

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

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

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

$$A_t \geq 0 \quad \forall t = \tau \dots T \quad (3)$$

⁶ An exception is Deaton (1992a) who specifies a more complicated rule of thumb; see the discussion of Rule 3 in Section 3 below.

Maximization of expected discounted utility given by (1) is subject to a standard law of motion for assets A_t (2) and a borrowing constraint (3). Note that while the household will optimally hold zero assets at the end of terminal period T , the individual might die before T with non-zero assets, i. e., there are accidental bequests in our model.

Income Y_t is determined by an exogenous process which is widely employed in the life-cycle literature (see, e. g., Carroll, 1992, 1997; Cocco *et al.* , 2005). This process features a deterministic hump-shaped pattern as well as both permanent and transitory shocks. In particular, income Y_t can be written as

$$Y_t = S_t P_t V_t \tag{4}$$

where P_t is the permanent income component and V_t is the transitory shock. Recall that the random variable S_t reflects life-time uncertainty and takes the value 1 as long as the individual is alive while it is set to zero thereafter. The permanent income component P_t itself follows a random walk with a drift

$$P_t = G_t P_{t-1} N_t \tag{5}$$

where G_t incorporates the deterministic and exogenously fixed hump-shaped life-cycle pattern of income and N_t is the permanent shock. Both transitory shocks V_t and permanent shocks N_t are i.i.d. log-normally distributed random variables with unit expectation and variance σ_v^2 and σ_n^2 , respectively.

Finally, we assume that the within-period utility function is of the Constant Relative Risk Aversion (CRRA) type,

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{6}$$

where γ is the coefficient of relative risk aversion (and the inverse of the intertemporal elasticity of substitution).

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 saving 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 three such rules.

Before we analyse 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 characterised by the following first-order condition:

$$u'(C_t) = \frac{1+r}{1+\rho} s_{t+1}^t E_t(u'(C_{t+1})). \quad (7)$$

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 being alive in period $t+1$. While the intuition of Euler equations such as (7) – 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 backwards induction. As noted by many authors before, this procedure is computationally demanding, and we can safely assume that individuals do not actually solve this problem when making their consumption and saving decisions; see, e.g., Hey and Dardanoni (1988). Moreover, Pemberton (1993, p. 5) points out that the intuition behind the Euler equation does not help to find simpler behavioural rules that would generate *as if* behaviour.

To solve the intertemporal optimization problem 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,

$$X_{t+1} = (1+r)(X_t - C_t) + Y_{t+1}. \quad (8)$$

In order to reduce the number of state variables we follow the approach in Carroll (1992) and standardise all variables by the permanent income component. The solution to the optimization problem is then computed by backwards induction over value functions starting in the last period in life T where households consume all remaining wealth. Taking this into account the optimal behaviour in period $T-1$ can be computed, etc.

From a more technical point of view we employ equally spaced grids for both standardised cash on hand and standardised savings. Piecewise Cubic Hermite Interpolation is used for points not on the grid. Shocks to the income process are approximated by Gauss-Hermite quadrature.

3 Three rules of thumb for life-cycle saving decisions

In this section, we present three decision rules that allow individuals to make their life-cycle saving 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-

saving decisions. All these decision rules have been proposed in previous literature. Table 1 contains an overview.

Insert Table 1 about here.

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 (1992a). As we explain below, Deaton designed this rule with the explicit goal that it should be easy to compute but still match optimal behaviour closely. All the 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 Rule of thumb No. 1 (Keynes)

The first rule of thumb we consider is the simplest rule one can think of – just consume your current income and don’t save at all:

$$C_t = Y_t \tag{9}$$

This rule is related to the famous consumption function by Keynes (1936). We should note that a (macroeconomic) Keynesian consumption rule will typically allow for some constant fraction of current income to be saved in each period. Such a rule does not directly translate to the microeconomic life-cycle saving problem since it lacks an element of dissaving (either in rainy days or in retirement).⁷ Here, we use the extreme version that rules out any saving as a benchmark case. In the formal analysis of life-cycle consumption and saving 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).

3.2 Rule of thumb No. 2 (Permanent Income)

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

⁷ We return to the issue of whether rules of thumb could be modified or combined to account for different savings motives and stochastic shocks in the concluding section.

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

$$\sum_{i=t}^T Y_t^P (1+r)^{t-i} = A_t + H_t, \quad (10)$$

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

$$H_t = Y_t + E \left(\sum_{i=t+1}^T Y_i (1+r)^{t-i} \right). \quad (11)$$

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

$$C_t = Y_t^P = \frac{r}{1+r} \frac{1}{1 - (1+r)^{-(T-t+1)}} (A_t + H_t). \quad (12)$$

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

$$C_t = Y_t^P = \frac{1}{T-t} (A_t + H_t). \quad (13)$$

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 behaviour performs relative to the benchmark solution.

3.3 Rule of thumb No. 3 (Deaton)

Deaton (1992a) 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. As Deaton

states, 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). 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:

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

Below, we follow Deaton in setting this fraction to 30%. Deaton explicitly states that he specified this decision rule, including the choice of $\zeta = 30\%$, entirely *ad hoc*. 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.

3.4 Saving motives reflected by simple rules of thumb

Before we turn to simulating life-cycle saving decisions using these three rules of thumb, it is useful to briefly review the central motives for saving that operate in our benchmark life-cycle model. Table 2 contains an overview of how these saving motives are reflected in the three simple decision rules that we consider. It is important to recognise 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 saving motives are relevant. This implies that there is no universal ranking of these decision rules in terms of their usefulness to individuals.

Insert Table 2 about here.

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). Among our three rules of thumb, Deaton’s rule is the only one which allows for a precautionary saving motive, while the “consumption equals income” rule includes no saving motive at all. The permanent income rule is forward-looking in the 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 saving decisions. Therefore, this rule reflects the consumption smoothing motive of saving.

4 Simulation and evaluation of rule-of-thumb behaviour

In this section, we present simulation results and compute the utility losses associated with using three 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 consumption equivalent measure, i. e., the percentage increase in consumption needed in every period and every state which would provide an individual with the same expected life-time utility under a given behavioural rule as she would obtain had she solved the underlying optimization problem. In a different context, this measure of welfare loss has been used by Krueger and Ludwig (2007).

4.1 Calibration of the life-cycle model and welfare measure

In Table 3, we report the benchmark parameter values used to calibrate the model. We chose these values to be in the range of values typically used in the literature.⁸ The parameter values in the income process have been taken from Cocco *et al.* (2005). The authors report estimates of the deterministic, hump-shaped life-cycle profile as well as the estimated variance of permanent and transitory shocks. Furthermore, they differentiate between three education groups: households whose head does not have a high school degree, whose head has a high school degree, or whose head has a college degree, respectively. We stratify our simulations by these three education groups. This approach allows us to analyse how the performance of a rule of thumb is affected by the characteristics of the income process faced by the household. Moreover, stratification by education addresses some of the heterogeneity in saving and income insurance that has been documented by Blundell *et al.* (2008), *inter alia*.

Insert Table 3 about here.

In order to quantify the welfare loss associated with using rules of thumb instead of the optimal solution to the maximization problem we compute a consumption equivalent measure. To be precise we compute the percentage increase in consumption a household following a given rule of thumb would need in each period and in each state in order to have the same expected life-time utility as if she was using the optimal solution. Since the within-period utility function is of the CRRA form there exists a closed form solution for this measure once the expected life-time utility has been computed under both the rule of thumb and the rational behaviour.

⁸ The parameter values used in this study also lie in the range of parameter values obtained using structural estimation by Gourinchas and Parker (2002) and Cagetti (2003).

The analytic solution for the consumption equivalent can be derived in the following way. Under CRRA preferences the expected life-time utility of a household following a given rule of thumb EU^{ROT} is given by

$$EU^{ROT} = E \left[\sum_{t=0}^T (1 + \rho)^{-t} s_t^0 \frac{(C_t^{ROT})^{1-\gamma}}{1-\gamma} \right] \quad (14)$$

where s_t^0 is the ex ante probability of being alive in period t . The consumption equivalent CE is then defined as

$$E \left[\sum_{t=0}^T (1 + \rho)^{-t} s_t^0 \frac{(C_t^{ROT} \cdot (1 + CE))^{1-\gamma}}{1-\gamma} \right] = EU^{OPT} \quad (15)$$

$$\Leftrightarrow EU^{ROT} \cdot (1 + CE)^{1-\gamma} = EU^{OPT} \quad (16)$$

where EU^{OPT} is the expected life-time utility of a household who has solved the maximization problem. Rearranging equation (16) yields the closed form solution for the consumption equivalent measure:

$$CE = \left(\frac{EU^{OPT}}{EU^{ROT}} \right)^{1/(1-\gamma)} - 1 \quad (17)$$

In order to obtain the consumption equivalent we thus simulate the discounted life-time utility of 100,000 households both under the optimal solution and under the given rule of thumb. The expected life-time utility under each rule is then computed as the average simulated discounted life-time utility. Equation (17) then directly gives the consumption equivalent measure.⁹

4.2 Simulation results

In the following, we first show, for each of the three education groups, the age profiles of consumption spending and asset holdings that arise for the alternative decision rules and in three different scenarios: (i) income certain, life length certain; (ii) income stochastic with i.i.d. shocks, life length uncertain; and (iii) income stochastic as a random walk, life length uncertain. We then summarize the implied utility losses that arise, expressed as consumption equivalents.

Figure 1 shows consumption spending (left column) and asset holdings (right column) for the three educational groups (from top to bottom: no high school, high school, college)

⁹ As an alternative welfare measure we also computed the percentage change in *income* which would give households the same expected utility under a given rule of thumb as under the optimal behaviour. In this case, however, there is no closed form solution since households adjust their behaviour as income changes. The computation is therefore more involved. But under the rules of thumb considered here, households spend most (or even all) of their income anyway, in particular in earlier periods of life. The two welfare measures (percentage increase in consumption and in income, respectively) thus turn out to be almost identical.

when there is no uncertainty – the income path is deterministic, and the length of life is known. In this case, realized income is equal to expected income in every period, and so Deaton’s Rule 3 collapses to Rule 1, the Keynesian rule. Both rules imply that households consume their deterministic income in every period. Thus, the dashed line that depicts consumption of these households in the graphs in the left column also shows the deterministic income path.

In the absence of impatience (i.e., interest rate r equal to discount rate ρ), behaviour according to Rule 2 (Permanent Income) would be identical to the optimal behaviour of the rational consumer. However, in the benchmark parameterisation, households are impatient, so optimal behaviour (solid line) shifts more consumption to earlier periods in life than permanent income rule (dash-dotted line) would suggest. Note also that since households are borrowing constrained, Rule 2 households consume less than their permanent income early in life. Finally, due to impatience, rational households would like to have a decreasing consumption profile over their life cycle; however due to the borrowing constraint, these households cannot consume more than their income early in life.

The differences between education groups are also interesting. With increasing education, the income profile becomes steeper. The implications are straightforward. Marginal utility of consumption is high for low levels of consumption at the beginning of life and lower for higher consumption levels later in life. For given education and income profiles, consumption levels are identical for all rules as long as households are still borrowing constrained. Between households that follow different rules, differences in consumption arise only later in life, and these differences are more pronounced the higher education and the steeper the income profile.

The right column of graphs in Figure 1 shows asset holdings for the three education groups and alternative decision rules. As noted above, Rule 1 and 3 (Keynesian and Deaton) households consume their current income in every period so they do not build up any assets. Asset holdings are larger for Rule 2 (permanent income) households than for rational households since the latter prefer earlier consumption because of their impatience.

Insert Figure 1 about here.

In Figure 2, we introduce uncertainty: Income follows a stochastic process with a deterministic component and i.i.d. shocks, and the length of life is uncertain. In this environment, rules of thumb can now show whether they succeed in helping households insure against income shocks and life-time uncertainty in the absence of insurance markets. As noted above, we simulated realizations of the stochastic income process for 100,000 households; the graphs in Figure 2 show averages across these simulations, again by education group and decision rule. The graphs in the left column again show consumption spending, those in the right column show asset holdings.

Note first that differences in consumption, and thus also in asset holdings, arise already early in life when households are borrowing constrained, in contrast to the deterministic income case in Figure 1. This is because some households realize positive shocks which are, depending on the decision rule used, not consumed fully. Also for this reason, consumption decisions implied by Rule 1 (Keynesian) and Rule 3 (Deaton) are now different. Another difference to the earlier results is that because of uncertain lifetime, households effectively become more impatient, and impatience increases with age. Thus, consumption is shifted towards younger ages. In particular, rational consumers save less, so their behaviour moves closer to that of Rule 1 and 3 households; conversely, it differs more from that of Rule 2 (permanent income) households. The asset holdings reflect these effects of uncertainty.

A peculiar feature of Rule 3 is that it does not allow for asset decumulation once households are retired. This drawback of this rule is due to its simplicity. One could also imagine that households who follow this rule switch to a different one after retirement. However, we refrain from modeling rule switching in this paper.

Insert Figure 2 about here.

Finally, Figure 3, which is similarly structured, shows the average age profile of consumption expenditure and asset holdings when income follows a random walk so that shocks have much more severe consequences. The main implication is that rational households save more and have larger asset holdings than in the i.i.d. case. For the three rules of thumb, the differences are minor.

Insert Figure 3 about here.

We now turn to the central question of this paper: How does rule-of-thumb behaviour translate into utility losses, relative to using the optimal decision rule? Table 4 shows utility losses, expressed as consumption equivalents, for a variety of scenarios. More specifically, the numbers in the table are the percentage increases in consumption in each period that would compensate an individual for using a non-optimal decision rule rather than solving the corresponding intertemporal optimization problem.

Insert Table 4 about here.

In Panel A, we report results using the benchmark calibration described in Table 3. When there is no income risk, consumption equivalents (CEs) and thus utility losses are generally very low (less than one percent). The CE is lowest for college graduates since their income profile is the steepest. This implies that their marginal utility is very high early in life relative to low marginal utility later in life when their consumption level is very high. But differences between the decision rules only occur once households are no longer credit constrained and their consumption level is already high. Hence, college graduates require only a small increase in consumption in the early periods to compensate for the relatively low utility losses later in life. The CEs are relatively high for households in the middle

education group that follow Rules 1 and 3: These households face a large drop in income at retirement which is not smoothed by these rules. These results provide a first important insight: Utility losses implied by rules of thumb depend on the income profile, both in early life and around retirement, and they are not necessarily ranked by education.

There are also differences in CEs depending on whether the length of life is uncertain. As mentioned above, life time uncertainty effectively makes agents more impatient. Hence rational households shift more consumption to earlier periods in life. This implies that rational behaviour differs more from the consumption profile implied by the permanent income rule, and thus the utility loss of following this rule increases.

Next, consider the case with stochastic income (the third and fourth lines of Panel A). Rule 1 (Keynesian) households do poorly now since not smoothing transitory income shocks is very costly: their CE goes up to between 15 percent (high education) and 25 percent (low education). CEs for the other two rules also increase relative to the certain income case, but they remain below 10 percent for all education groups. Deaton's Rule 3 does quite well – this is exactly the environment for which it was designed. Rule 2 (permanent income) does relatively poorly for high education households whose income profile is steep so that borrowing constraints bind longer. Whether lifetime is uncertain or not does not have a major effect.

The random walk case (the last two lines of Panel A) is perhaps most realistic. Even though income shocks now have permanent effects, CEs are still relatively low (mostly below 10 percent), in particular for Rules 2 and 3. Here, the permanent income rule does well because it takes shocks into account by updating permanent income. However, there is no smoothing of small shocks early in life because of borrowing constraints, so utility losses relative to optimal behaviour still arise. Only the Keynesian Rule 1 does poorly, even though CEs are smaller than in the i.i.d. case.

Before we turn to a more general interpretation of these results, we briefly discuss Panels B and C of Table 4. In these simulations, we use alternative parameter values to investigate the sensitivity of our results with respect to households' risk aversion and the variance of the income process. In Panel B, we increase risk aversion by setting the coefficient γ equal to 9 rather than 3. If risk aversion is higher, the curvature of the utility function is larger, so the effects of steep income profiles are more pronounced. As long as income is certain, high school graduates that follow any of the rules of thumb come close to optimal behaviour in utility terms. For the other education groups, CEs tend to increase when risk aversion increases for most decision rules but they remain low. When income is uncertain, higher risk aversion affects rule-of-thumb households quite negatively unless they have very steep income profiles. For the low and middle education groups, CEs go up considerably in the i.i.d. case. Finally, when income follows a random walk, high risk aversion households fare very poorly under any rule of thumb; Rule 2 (permanent income) does slightly better.

In Panel C we increase the variance of the permanent income shock in the random walk case by 3 standard errors. These numbers should be compared with those in the last two lines of Panel A. It is apparent that increasing income variance makes rules of thumb perform less well; CEs increase across the board. However, they typically remain below 20 percent which is perhaps surprisingly low given that the income process has now a rather large variance. As before, Rule 2 (permanent income) performs best in the random walk case since it updates permanent income after each shock.

We can draw several conclusions from the numbers in Table 4. 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 saving decisions suffer considerable utility losses relative to the optimal decision rule. Not surprisingly, the magnitudes of these utility losses depend on preferences (here we studied risk aversion) 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. Unless risk aversion is very high ($\rho = 9$), utility losses as expressed by consumption equivalents are below 20 percent, and the permanent income rule stays even below 10 percent.

The main conclusion from this comparison is that 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 analysed 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. When risk aversion is very high, using a rule of thumb is a particularly bad idea. This observation is interesting since one might speculate that in real life, households with high risk aversion might also be those who are financially less sophisticated and thus more likely to use rules of thumb. Similarly, high education households which might be more likely to use sophisticated savings strategies would actually perform relatively better using rules of thumb than low-education households since their income profiles are steeper.

What does a rule of thumb make perform well? Based on our 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

In our simulations of life-cycle consumption and saving decisions under three heuristic decision rules, we found that losses in total life-time utility can, in general, be substantial compared with optimal behaviour (i. e., using the solution of the underlying intertemporal optimization problem for given standard life-cycle preferences). The magnitudes of these losses vary with the assumptions about preference parameters and the properties of the income process. An important result is that for most environments we simulated, there exists some simple rule of thumb which results in only modest utility losses (expressed by compensating variations, these losses are equivalent to between 5% and 10% of life-time income).

We conclude this paper with a discussion of possible extensions of our analysis and implications for future research in the life-cycle consumption and saving framework. The saving rules we analysed in this paper reflect two distinct motives for saving (consumption smoothing and precautionary saving). A natural extension of our approach would be to combine two or more of these rules of thumb. For example, individuals could use a simple static rule such as Deaton's rule (Rule 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 2) which focuses on income 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 in utility terms. The concept of mental accounting introduced to the life-cycle saving literature by Shefrin and Thaler (1988) is closely related to the notion of using multiple rules of thumb. For instance, mental accounting implies that individuals make saving decisions with different time horizons and saving goals in mind. Combining rules of thumb with different objectives such as those analysed in this paper with a mental accounting framework appears to be a fruitful direction for future research.

Unfortunately, combining two or more saving rules in our framework would complicate the analysis considerably, both technically and conceptually. The optimization and simulation problem is obviously much more involved, but more importantly, one would have to make assumptions of how individuals allocate funds to different saving rules. This allocation might change over the life cycle. This would imply that a second class of behavioural decision rules would have to be specified for the allocation of saving into different mental accounts.

It would also be interesting to compare the welfare consequences of suboptimal savings decisions (the subject of the present paper) with those of suboptimal portfolio choices (as in Calvet *et al.* , 2007) when both are driven by rules of thumb. To perform such a comparison, both simulation models would have to be specified and calibrated consistently, or perhaps even embedded in a unified model. Such an analysis would allow us

to address the question of whether it is more important to use good rules of thumb for the consumption-saving decision or for the decision of how to invest whatever amount is saved.

Another important research question that we have not addressed in this paper is: How do rules of thumb actually arise? How do individuals decide which behavioural 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. Rule-of-thumb behaviour could be derived from some meta problem if the cost of computing the solution to the underlying life-cycle optimization problem was taken into account.¹⁰ However, such an approach might run into the conceptual problem of an infinite regress, as discussed by Conlisk (1996).

A promising approach is to explore how rules of thumb arise endogenously from learning behaviour; for example, Lettau and Uhlig (1999) investigate a model of learning rules of thumb in intertemporal decision problems. However, in a life-cycle saving 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 in the life-cycle saving problem is likely to be based on social interactions with other people (family, neighbors, or friends). Brown *et al.* (2009) present evidence for the relevance of learning: In an experimental study with artificial life-cycle decisions, subjects saved too little initially but learned to save optimally as life cycles were repeated. In the domain of entrepreneurial financial decisions, Drexler *et al.* (2010) conducted a field experiment with micro-entrepreneurs in the Dominican Republic and found that simple rule-of-thumb training improved business practices and outcomes. An important question for future research thus seems to be how individuals acquire the rules they use for making financial decisions, and whether policy interventions can help individuals to improve the rules they use. This research agenda is obviously related to the ongoing discussion of how households’ financial literacy can be improved.

Certain institutions might also provide individuals with saving rules so that they don’t need to figure out optimal or heuristic saving rules themselves. An important example is social security which replaces the need for discretionary long-term saving to some extent. Another mechanism that provides rules for long-term saving is housing expenditure. In countries such as the U.S. and the U.K., many households buy their first family homes relatively early in their life cycle, and this decision determines a large fraction of their consumption and saving pattern over future years. In Germany, due to its favorable tax

¹⁰ Computation costs have been considered in models with rule-of-thumb behaviour by Shi and Epstein (1993) and Hindy and Zhu (1997).

treatment, the acquisition of life-insurance policies with substantial saving components during the early stage of the active working life has been quite common in the past, see Sauter and Winter (2010). The acquisition of a family home or a life-insurance policy is a one-time decision that fixes a substantial part of life-cycle saving and it reduces the scope for discretionary saving over remaining years substantially. There is also some recent evidence that individuals are quite willing to follow saving rules provided by institutional arrangements; e. g., Thaler and Benartzi (2004) and Choi *et al.* (2005). The behavioural patterns mentioned in this paragraph (based on direct imitation, social traditions, or institutional arrangements) can be interpreted as following a heuristic saving rule, and they might result in decisions that are quite close to optimal life-cycle behaviour.

Any empirical analysis of rule-of-thumb saving behaviour 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 behaviour 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, some for long-term consumption smoothing, while others might rely on saving rules provided by institutional arrangements. As noted above, combinations of rules might arise as well. To our knowledge, there exist no econometric studies that try to identify individual decisions rules in a life-cycle saving context, but this is clearly an important area for future research that could build on results from laboratory experiments on decision rules in dynamic problems by Müller (2001), Houser and Winter (2004), and more recently Hey and Lotito (2009), Hey and Panaccione (2011), and Hey and Knoll (2010).

Based on our results on the performance of simple saving rules, we would argue that an important direction for research on life-cycle consumption and saving behaviour is how social learning and institutional factors generate simple decision rules, and to test whether individuals actually use such rules with micro data on household consumption and saving.

References

- Anderhub, V., W. Güth, W. Müller, and M. Strobel (2000):** An experimental analysis of intertemporal allocation behavior. *Experimental Economics*, 3, 137–152.
- Attanasio, O. P. and G. Weber (2010):** Consumption and saving: Models of intertemporal allocation and their implications for public policy. *Journal of Economic Literature*, 48, 693–751.
- Binswanger, J. and K. G. Carman (2011):** How real people make long-term decisions: The case of retirement preparation. Unpublished manuscript, Tilburg University.
- Blundell, R., L. Pistaferri, and I. Preston (2008):** Consumption inequality and partial insurance. *American Economic Review*, 98, 1887–1921.

- Brown, A. L., Z. E. Chua, and C. F. Camerer (2009):** Learning and visceral temptation in dynamic saving experiments. *Quarterly Journal of Economics*, 124, 197–231.
- Browning, M. and T. F. Crossley (2001):** The life-cycle model of consumption and saving. *Journal of Economic Perspectives*, 15(3), 3–22.
- Browning, M. and A. Lusardi (1996):** Household saving: Micro theories and micro facts. *Journal of Economic Literature*, 34, 1797–1855.
- Cagett, M. (2003):** Wealth accumulation over the life cycle and precautionary savings. *Journal of Business and Economic Statistics*, 21, 339–353.
- Calvet, L. E., J. Y. Campbell, and P. Sodini (2007):** Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy*, 115, 707–747.
- 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 Economic Statistics*, 8, 265–279.
- Carbone, E. and J. D. Hey (1999):** A test of the principle of optimality. *Theory and Decision*, 50, 263–281.
- 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, 1–55.
- Choi, J. J., D. Laibson, B. C. Madrian, and A. Metrick (2005):** Passive decisions and potent defaults. In D. A. Wise (Ed.), *Analyses in the Economics of Aging*, 59–78. Chicago and London: Cambridge University Press.
- Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005):** Consumption and portfolio choice over the life cycle. *Review of Financial Studies*, 18, 491–533.
- Cochrane, J. H. (1989):** The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives. *American Economic Review*, 79, 319–337.
- Conlisk, J. (1996):** Why bounded rationality? *Journal of Economic Literature*, 34, 669–700.
- Deaton, A. (1991):** Saving and liquidity constraints. *Econometrica*, 59(5), 1221–1248.
- Deaton, A. (1992a):** Household saving in LDCs: Credit markets, insurance and welfare. *Scandinavian Journal of Economics*, 94, 253–273.
- Deaton, A. (1992b):** *Understanding Consumption*. Oxford, UK: Clarendon Press.
- Drexler, A., G. Fischer, and A. Schoar (2010):** Keeping it simple: Financial literacy and rules of thumb. Discussion Paper No. 7994, CEPR, London.
- Flavin, M. A. (1981):** The adjustment of consumption to changing expectations about future income. *Journal of Political Economy*, 89, 974–1009.

- Friedman, M. (1957):** *A Theory of the Consumption Function*. Princeton, NJ: Princeton University Press.
- Gigerenzer, G. and P. M. Todd (1999):** Fast and frugal heuristics. In G. Gigerenzer, P. M. Todd, and the ABC Research Group (Eds.), *Simple Heuristics That Make Us Smart*, 3–34. New York, NY and Oxford, UK: Oxford University Press.
- Goodie, A. S., A. Ortmann, J. N. Davis, S. Bullock, and G. M. Werner (1999):** Demons versus heuristics in artificial intelligence, behavioral ecology, and economics. In G. Gigerenzer, P. M. Todd, and the ABC Research Group (Eds.), *Simple Heuristics That Make Us Smart*, 327–355. New York, NY and Oxford, UK: Oxford University Press.
- Gourinchas, P.-O. and J. A. Parker (2002):** Consumption over the life cycle. *Econometrica*, 70, 47–89.
- Hall, R. E. (1978):** Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86, 971–987.
- Hall, R. E. and F. Mishkin (1982):** The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica*, 50, 461–481.
- Hey, J. D. (2005):** Do people (want to) plan? *Scottish Journal of Political Economy*, 52, 122–138.
- Hey, J. D. and V. Dardanoni (1988):** Optimal consumption under uncertainty: An experimental investigation. *Economic Journal*, 98, S105–S116.
- Hey, J. D. and J. A. Knoll (2010):** Strategies in dynamic decision making: An experimental investigation of the rationality of decision behaviour. Unpublished manuscript, University of York.
- Hey, J. D. and G. Lotito (2009):** Naive, resolute or sophisticated? A study of dynamic decision making. *Journal of Risk and Uncertainty*, 38, 1–25.
- Hey, J. D. and L. Panaccione (2011):** Dynamic decision making: What do people do? *Journal of Risk and Uncertainty*, 42, 85–123.
- 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.
- Houser, D. and J. Winter (2004):** How do behavioral assumptions affect structural inference? Evidence from a laboratory experiment. *Journal of Business and Economic Statistics*, 22, 64–79.
- Johnson, S., L. J. Kotlikoff, and W. Samuelson (2001):** Can people compute: An experimental test of the life-cycle consumption model. In L. C. Kotlikoff (Ed.), *Essays on Saving, Bequests, Altruism, and Life-Cycle Planning*, 335–385. Cambridge, MA, and London, UK: MIT Press (first circulated in March 1987 as NBER Working Paper No. 2183).
- Keynes, J. M. (1936):** *The General Theory of Employment, Interest and Money*. London: Macmillan.

- Krueger, D. and A. Ludwig (2007):** On the consequences of demographic change for rates of returns to capital, and the distribution of wealth and welfare. *Journal of Monetary Economics*, 54, 49–87.
- 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, 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.
- Low, H., C. Meghir, and L. Pistaferri (2010):** Wage risk and employment risk over the life cycle. *American Economic Review*, 100, 1432–1467.
- Lusardi, A. and O. S. Mitchell (2007):** Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics*, 54, 205–224.
- Moon, P. and A. Martin (1990):** Better heuristics for economic search: Experimental and simulation evidence. *Journal of Behavioral Decision Making*, 3, 177–193.
- Müller, W. (2001):** Strategies, heuristics, and the relevance of risk-aversion in a dynamic decision problem. *Journal of Economic Psychology*, 22, 493–522.
- 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, 11–46.
- Rodepeter, R. and J. Winter (1998):** Saving decisions under life-time and earnings uncertainty: Empirical evidence from West German household data. Discussion Paper No. 98-58, Sonderforschungsbereich 504, Universität Mannheim.
- Rodepeter, R. and J. Winter (1999):** Rules of thumb in life-cycle saving decisions. Discussion Paper No. 99-81, Sonderforschungsbereich 504, Universität Mannheim.
- Sauter, N. and J. Winter (2010):** Do investors respond to tax reform? Evidence from a natural experiment in Germany. *Economics Letters*, 108, 193–196.
- Scholz, J. K., A. Seshadri, and S. Khitatrakun (2006):** Are Americans saving “optimally” for retirement? *Journal of Political Economy*, 114, 607–643.
- 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, 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.
- Thaler, R. H. and S. Benartzi (2004):** Save more tomorrow: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112, S164–S187.
- van Rooij, M. C. J., A. Lusardi, and R. Alessie (2011):** Financial literacy and stock market participation. *Journal of Financial Economics*, 101, 449–472.
- Wärneryd, K.-E. (1989):** On the psychology of saving: An essay on economic behavior. *Journal of Economic Psychology*, 10, 515–541.
- Weber, C. E. (2000):** Rule-of-thumb consumption, intertemporal substitution, and risk aversion. *Journal of Business and Economic Statistics*, 18(4), 497–502.
- Weber, C. E. (2002):** Intertemporal non-separability and “rule-of-thumb” consumption. *Journal of Monetary Economics*, 49, 293–308.

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

Decision rule	Description	Source
Benchmark	Solution to the underlying intertemporal optimization problem	Carroll (1992), Rodepeter and Winter (1998)
Rule 1	Consumption equals current income	related to Keynes (1936)
Rule 2	Consumption equals permanent income	Friedman (1957)
Rule 3	Consumption equals cash on hand up to mean income, plus 30% of excess income	Deaton (1992a)

Table 2: Savings motives captured by alternative decision rules

	Benchmark	Rule 1	Rule 2	Rule 3
Consumption smoothing	yes	no	yes	no
Precautionary saving	yes	no	no	yes
Life-time uncertainty	yes	no	no	no

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

Parameter		Benchmark value		
<i>Preference parameters:</i>				
Relative risk aversion coefficient	γ		3	
Rate of time preference	ρ		4%	
Interest rate	r		2%	
<i>Life-time parameters:</i>				
Conditional survival probabilities	s_t^τ		life-table values	
Retirement age			65	
Max. age at death ^a	T		100	
<i>Income process (i.i.d.):</i>				
Variance of transitory shock	$\sigma_{e,IID}^2$	0.25		
Variance of permanent shock	$\sigma_{u,IID}^2$	0		
<i>Income process (random walk):</i>				
		<i>no high school</i>	<i>high school</i>	<i>college</i>
Variance of transitory shock ^b	$\sigma_{e,RW}^2$	0.1056 (0.0080)	0.0738 (0.0034)	0.0584 (0.0045)
Variance of permanent shock ^b	$\sigma_{u,RW}^2$	0.0105 (0.0011)	0.0106 (0.0004)	0.0169 (0.0006)

^a In specifications without life-time uncertainty, the age of death is fixed at 80.


^b Parameters values are from Cocco *et al.* (2005). Standard errors in parentheses.

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


Income	Length of life	No high school			High school			College		
		Rule 1	Rule 2	Rule 3	Rule 1	Rule 2	Rule 3	Rule 1	Rule 2	Rule 3
<i>Panel A: benchmark calibration</i>										
certain	certain	0.2	0.7	0.2	0.9	0.4	0.9	0.1	0.2	0.1
certain	uncertain	0.0	3.0	0.0	0.4	2.7	0.4	0.0	1.1	0.0
i.i.d.	certain	24.7	4.3	6.9	21.8	6.7	6.2	15.9	8.3	3.7
i.i.d.	uncertain	24.2	5.0	6.4	21.1	7.1	5.6	15.6	8.5	3.5
random	certain	19.9	2.8	11.7	14.2	2.7	9.8	14.4	4.3	11.4
random	uncertain	16.9	3.4	8.8	11.0	3.6	6.5	11.1	4.2	8.2
<i>Panel B: high risk aversion ($\gamma = 9$)</i>										
certain	certain	1.5	0.1	1.5	1.7	0.0	1.7	0.0	0.0	0.0
certain	uncertain	0.9	1.1	0.9	1.2	0.4	1.2	0.0	0.0	0.0
i.i.d.	certain	40.5	32.4	32.0	32.9	30.4	29.2	4.4	4.0	2.0
i.i.d.	uncertain	39.4	32.1	32.0	32.8	30.2	29.2	4.4	4.0	2.0
random	certain	179.2	55.5	141.1	140.3	46.8	135.1	170.8	72.0	162.4
random	uncertain	160.2	44.3	125.5	118.5	39.1	113.0	149.3	59.5	140.8
<i>Panel C: high variance of permanent shocks (+ 3 standard errors)</i>										
random	certain	26.4	4.8	17.9	16.3	3.3	11.7	17.3	5.4	14.2
random	uncertain	22.0	4.8	13.4	12.4	4.0	7.9	13.4	5.0	10.4

Note: Life-time utility losses are expressed as the percentage increase in consumption in each period and state that would compensate an individual for using non-optimal decision rules rather than solving the corresponding intertemporal optimization problem. Parameter values for the benchmark model are reported in Table 3.


Figure 1: Behaviour due to different decision rules: income certain, life length certain



Panel A: no high school




Panel B: high school




Panel C: college

Note: The figure plots the behaviour generated by the different decision rules: rational behaviour (solid line), Keynes Rule 1 & Deaton Rule 3 (dashed line), Permanent Income Rule 2 (dash-dotted line)


Figure 2: Mean behaviour due to different decision rules: income i.i.d. life length uncertain



Panel A: no high school




Panel B: high school




Panel C: college

Note: The figure plots the mean behaviour generated by the different decision rules (calculated from 100,000 simulations): rational behaviour (solid line), Keynes Rule 1 (dashed line), Permanent Income Rule 2 (dash-dotted line), Deaton Rule 3 (dotted line)


Figure 3: Mean behaviour due to different decision rules: income random walk, life length uncertain



Panel A: no high school



Panel B: high school



Panel C: college

Note: The figure plots the mean behaviour generated by the different decision rules (calculated from 100,000 simulations): rational behaviour (solid line), Keynes Rule 1 (dashed line), Permanent Income Rule 2 (dash-dotted line), Deaton Rule 3 (dotted line)