Optimal Rules of Thumb for Consumption and Portfolio Choice

David A. Love*

November 21, 2011

Abstract

Conventional rules of thumb represent simple, but potentially inefficient, alternatives to dynamic programming solutions. This paper seeks an intermediate ground by developing a framework for selecting optimal rules of thumb. Defining rules of thumb as simple functions of state variables, I solve for the optimal parameters of specific rules of thumb for portfolio choice and consumption. In the case of portfolio choice, I find that optimal linear age rules lead to modest welfare losses relative to the dynamic programming solution and that a linear rule based on the ratio of financial wealth to total lifetime resources performs even better. Consumption rules generate larger welfare losses—from 1–8% of annual consumption—but an effective rule is to consume 70–80% of annuitized lifetime wealth.

JEL classification: G11; G22; D91; E21

Keywords: Portfolio choice; precautionary saving; rules of thumb.

*Dept. of Economics, Williams College, Williamstown, MA 01267, david.love@williams.edu.

But really there are no rules but rules of thumb.

Clive James

1. Introduction

Rules of thumb are common in personal finance. A quick internet search returns a "save 10% of pre-tax income" rule for saving, a "4% withdrawal" rule for dissaving in retirement, and a "place 100 minus your age in stocks" rule for asset allocation.[1] Rules like these may be easy to understand, but they differ sharply from the optimal decisions that emerge from a standard life-cycle model. In particular, conventional rules fail to account for differences in household characteristics, and they do not respond to changing realizations of risk, resources, and spending needs over the life cycle. One critic, Laurence Kotlikoff, has even gone so far as to characterize the recommendation of simple rules of thumb as a form of "financial malpractice."[2] At the very least, the prevalence and arbitrary nature of these rules invite a series of questions. How large are the welfare losses associated with applying a particular rule of thumb? How robust are rules to differences in investor characteristics? Can simple rules ever approach the efficiency of the optimal model? If so, what would such rules look like?

I address these questions by developing a framework for selecting "optimal rules of thumb" in consumption and portfolio choice. The framework starts from the observation that the solution to any dynamic decision problem takes the form of policy rules mapping states into actions (Lettau and Uhlig, 1999). While dynamic programming provides a method for finding *optimal* rules, the non-linear nature of such rules defies easy distillation into conventional forms of financial advice. Rules of thumb, in contrast, have the advantage of being easy to

[1] The unacknowledged source of many of these rules seems to be Burton Malkiel's (2011) "A Random Walk Down Wall Street," which provides considerably more subtle advice than the abridged versions offered on many personal-finance websites and news articles. For instance, while Malkiel recommends a life-cycle allocation of stocks that approximates a linear age rule, he emphasizes that investors need to take into account housing costs, health, risk tolerance, and income uncertainty. As we will see, the solution to the optimal life-cycle allocation problem is not inconsistent with the more nuanced advice in Malkiel's book.

[2] The quote comes from an article posted on Bankrate.com by Jay MacDonald titled "Figuring Retirement Savings Spend-Down Rate," updated Oct. 1 2007, available at `http://www.bankrate.com/brm/news/retirementguide2007/20070501_spend-down_rate_a1.asp?caret=4b`. Kotlikoff's remarks referred to advice such as the 4% spend-down rule, which he argued can lead to inefficiently high amounts of saving.

calculate, objectively communicable, and independent of individual judgement (Baumol and Quandt, 1964). But they may also be inefficient. This paper seeks a middle ground between the efficiency of the optimal solution and the simplicity of conventional advice. Defining rules of thumb to be simple functions of state variables, I solve for the set of parameters that maximizes welfare given a specific function type. That set of parameters characterizes an optimal rule of thumb.

Previous work has examined the performance of specific rules of thumb for saving (Winter et al., 2011) and asset allocation (Cocco et al., 2005; Gomes et al., 2008). The primary contribution of this paper is to move beyond analyzing the welfare properties of specific rules to measuring the performance of the best rules within a given class. There are at least two advantages to the approach. First, because the rules have been optimized, the framework helps answer the question of whether *any* version of a rule holds promise as an alternative to the more complex dynamic programming solution. If the welfare losses associated with an optimized rule are large, it suggests that we may want to search for a different class of rules altogether. Second, the optimal rules may be of interest in their own right. I find, for example, that while common personal finance rules tend to be inefficient, some new rules, and even new parameterizations of existing rules, perform surprisingly well.

I explore two different types of life-cycle rules of thumb: portfolio allocation rules, assuming optimized consumption; and consumption rules, assuming that households can only invest in the risk-free asset. The portfolio rules take the form of linear functions of either age or the ratio of financial wealth to a measure of total lifetime resources. The first of these has the advantage of simplicity, while the second responds to financial variables that theory suggests should influence portfolio choice (Bodie et al., 1992). The consumption rules fall into two categories as well. According to the first rule, households consume a fraction of permanent income during the working life and withdraw from savings at a constant rate in retirement. This type of rule is meant to capture the spirit of the conventional advice that households save 10% of income while working and draw down 4% of assets in retirement. The second rule, which is more in line with economic theory, sets consumption equal to a fraction of annuitized present value resources.

How well do the optimal rules of thumb perform? In terms of portfolio choice, the optimal linear rules are only moderately inefficient compared to the dynamic programming solution. If individuals adhere to an optimal linear age rule for the remainder of life, for example, welfare losses generally amount to less than 0.5% of annual consumption.[3] I find, however, that while each of the rules leads to only modest welfare losses from the perspective of younger workers, the age-based rule becomes increasingly inefficient over time, while the wealth-based rule remains effective all the way through the retirement period. Not surprisingly, the welfare losses for both types of rules improve measurably if individuals are allowed to update the rules at different ages. In fact, the optimal wealth-based rule is capable of getting very close to the welfare achieved using dynamic programming. Allowing for updating at ages 40 and 65, the welfare losses for a 20-year-old college graduate fall below 0.06 percent of annual consumption, or about $32 a year.

Compared with the portfolio rules, the optimal consumption rules generate much larger welfare losses. The first rule, which sets consumption equal to a fraction of income during working life and assumes a constant withdrawal rate in retirement, leads to welfare losses of 4–8% of consumption per year. These losses fall if households are allowed to update at later ages, but not by much. Not surprisingly, this rule is outperformed by the one based on the annuitized value of total wealth. Here, the welfare losses generally range between 1% and 4% of annual consumption, with some losses falling below 1% in the presence of updating. It turns out that a relatively efficient rule of thumb is to consume between 70 and 80 percent of annuitized wealth.

I also consider the robustness of the rules to uncertainty about the underlying preference parameters, such as risk aversion or impatience. I solve for robust rules of thumb that optimize welfare subject to the functional form contraint, assuming that the rule maker has knowledge only of the distribution of the relevant parameter and not its precise value. While the optimal

[3] I measure the performance of the rules by computing consumption-equivalent welfare losses relative to the dynamic programming solution. In contrast to previous work examining rules of thumb, however, the current paper does not examine the welfare properties of rules exclusively from the perspective of individuals at the beginning of the working life. Although this may be a natural starting point for welfare analysis, it has the drawback that discounting can make some rules attractive early in life that turn out to be ineffective later on. The approach taken in this paper is to instead measure welfare losses from the perspective of different ages.

rules are not dramatically different in the presence of preference uncertainty, the welfare losses tend to be substantially higher. In the case of consumption rules of thumb, for example, the welfare losses rise by as much as 2% of annual consumption when the model incorporates uncertainty in risk aversion or the discount factor. Thus, while parameter uncertainty does not lead to large changes in the optimal rules themselves, it can make the rules less attractive relative to the dynamic programming solution.

An immediate concern about the approach in this paper is that it seems strikingly inefficient. If we have the optimal decision rules for consumption and asset allocation in hand, what is the benefit of examining the "best" rules of thumb that, by definition, cannot improve upon the decision rules we already have? After all, even if households themselves cannot easily perform the numerical calculations underlying the optimal solution, there are companies like ESPlanner that provide this service for an annual fee.[4] One response to this criticism is that some rules of thumb may generate welfare losses that are actually smaller than the fees charged for more advanced solutions. Another response is that many households lack either the resources or the education to take advantage of more tailored financial advice, and these households would be well served by having access to effective rules of thumb.

A deeper concern about the exercise is that the rules of thumb only get close to optimal behavior if the benchmark model is correctly specified. Attanasio and Weber (2010) point out that there is no single life-cycle model, but rather a general framework organized around the principle that households maximize lifetime utility subject to resource constraints. The literature has developed a rich array of particular life-cycle models that impose varying structures on preferences and constraints, but there is no "right" specification. This paper considers rules to be optimal if they maximize an objective function subject to constraints (including the functional constraint defining the type of rule). If a household's specific preferences or constraints differ from those assumed in the benchmark model, there might be other rules that outperform the optimal ones found here.

The rest of the paper proceeds as follows. Section 2 relates the current study to previous

[4]ESPlanner is a company started and run by Laurence Kotlikoff that offers financial advice based on simulations of individually tailored life-cycle models. Kotlikoff calls this the "economics approach" to financial planning (Kotlikoff, 2007), and he contrasts it with conventional advice recommending simple saving and spending targets.

work on optimal consumer behavior and rules of thumb. Section 3 develops the framework for optimal rules of thumb. Section 4 presents a standard model of saving and asset allocation and discusses the benchmark parameterization. Section 5 examines the welfare costs of applying optimal rules of thumb in different settings. The final section offers concluding thoughts about the limitations of the framework and some ideas for extensions.

2. Literature review

The concept of an optimal rule of thumb is an old one, going back to Herbert Simon's notion of satisficing (Simon, 1978) and the "optimally imperfect" firm-level decisions in Baumol and Quandt (1964). While Baumol and Quandt (1964) did not have the benefit of the modern computing speeds needed to analyze dynamic life-cycle problems, their basic approach in many respects resembles the one pursued here. For instance, they too define rules of thumb in terms of specific function types, and they measure the performance of various rules against an optimal benchmark. The application of optimal rules to the life-cycle problem, however, is closer to Allen and Carroll (2001), who examine whether individuals can learn an approximate solution to a buffer-stock model of consumption through experienced utility. They find that they can, but only after an extremely large number of search periods. The current paper takes a more centralized approach to the problem, and offers an approximation that requires, if not a "supercomputer and a doctorate" (Allen and Carroll, 2001), at least modern processing speeds and an efficient solution algorithm.[5]

A handful of previous studies analyze the welfare properties of specific rules of thumb in lifecycle decisions. Winter et al. (2011) focus on the performance of several consumption and saving rules of thumb, while Cocco et al. (2005) and Gomes et al. (2008) consider rules of thumb for portfolio allocation. For different parameterizations and rules, each of these papers calculates a compensating welfare measure that would make a representative individual

[5] Allen and Carroll (2001) use the phrase to describe the challenge of arriving at the exact solution to the dynamic programming problem, not the approximate version in their learning algorithm. They write, "Despite its heuristic simplicity, the *exact* mathematical specification of optimal behavior is given by a thoroughly nonlinear consumption rule for which there is no analytical formula...[I]t is hard to see how a consumer without a supercomputer and a doctorate could be expected to determine the exact shape of the nonlinear and nonanalytical decision rule."

indifferent between that rule and the solution to the dynamic programming problem. They find that some rules perform reasonably well, while others generate larger welfare losses. But a comparison of the results in the two papers indicates that the welfare consequences of adopting rules of thumb tend to be at least an order of magnitude higher in the case of consumption rules than in portfolio allocation.

While the welfare approach taken in this paper is similar, there are several key departures. First, instead of focusing on specific rules of thumb, I consider different classes of rules and solve for the optimal parameterization. The differences in the welfare losses associated with an arbitrary rule of thumb advocated in popular finance and an optimized one often turn out to be substantial. Second, the welfare losses in the papers above are computed from the perspective of an individual at the beginning of the working life, and the parameterization of rules is held fixed for the duration of the life cycle. In contrast, I evaluate welfare from the perspective of different ages to show how the rules perform over the life cycle, and I allow for the possibility that individuals can update their rules in select periods. Finally, the welfare losses in Cocco et al. (2005), Gomes et al. (2008), and Winter et al. (2011) naturally depend on assumptions about individual preferences. This paper incorporates parameter uncertainty directly into the welfare calculations and solves for "optimal" rules of thumb that are robust to preference heterogeneity.

Do households follow rules of thumb in consumption and portfolio allocation? The evidence is mixed.[6] Campbell and Mankiw (1989, 1990) show that the responsiveness of aggregate consumption to changes in current income is consistent with about half of all households consuming their current income. These results stimulated a large literature examining the fraction of "rule of thumb" consumers, with estimates ranging from 15–85% of households (Weber, 2000). One needs to be cautious, however, in interpreting these fractions as representing rule-of-thumb behavior. The excess sensitivity of consumption to income may instead be due to standard life-cycle factors such as liquidity constraints or precautionary motives (Attanasio, 1999), or it may simply reflect bias in the econometric estimates of Euler equations (Weber, 2000).

Other evidence for rules of thumb comes from examining the relationship between household

[6] Attanasio and Weber (2010) provide an extensive review of the life-cycle consumption literature.

wealth and consumption patterns. Bernheim et al. (2001) find that, first, in contrast to the predictions of the life-cycle model, there does not appear to be any correlation between the level of wealth at retirement the growth rate of consumption, and second, that households with lower savings experience larger drops in consumption at retirement. They conclude that the observed patterns of consumption may be more consistent with rule-of-thumb behavior or mental accounting than with the standard consumption model. In contrast, Scholz et al. (2006) compare the actual savings of households in the Health and Retirement Study with the household-specific predicted levels from a life-cycle model and find a remarkably close match between the model's predictions and observed behavior. Further, they examine alternative heuristic rules and find that the life-cycle model does a significantly better job of matching the household-level data than "naïve" rules proposed in previous studies.

In terms of portfolio choice, the literature has focused more on reconciling the predictions of the model with the aggregate evidence on household allocations than on estimating the amount of rule of thumb behavior per se. Some of the most important contributions in this area have augmented the model in Cocco et al. (2005) to include participation costs and Epstein-Zin preferences (Gomes and Michaelides, 2005), housing (Cocco, 2005), luxury goods (Wachter and Yogo, 2010), and household debt (Becker and Shabani, 2010). This research does not, however, make predictions on a household-by-household basis in the spirit of Scholz et al. (2006), and it is difficult to reject the possibility that some households are relying on rules of thumb.

3. Optimal rules of thumb

Rules of thumb may emerge in a variety of economic settings, from firms' decisions about production, inventory, or hiring to household decisions about consumption and portfolio allocation. This section develops a framework for analyzing rules of thumb that is general enough to handle a range of interesting possibilities. Following Lettau and Uhlig (1999), I consider a general dynamic decision problem that can be solved optimally using dynamic programming or approximated using rules of thumb.[7]

An individual makes decisions over the discrete time horizon $\{t_0, \ldots, T\}$. In each period t, the individual observes a state vector, $s_t \in \mathcal{S}_t$, which summarizes the current set of relevant information. The individual then takes a vector of actions, $a_t \in \mathcal{A}_t$, which influences the evolution of the state vector in the next period.

Given n state variables and m possible actions, a decision rule is a vector-valued function mapping states into actions: $d_t : \mathcal{S}_t \subseteq \mathbb{R}^n \to \mathcal{A}_t \subseteq \mathbb{R}^m$. A policy $h_t = \{d_t(s_t), ..., d_T(s_T)\}$ is a sequence of decision rules specifying actions over the remaining planning horizon. The value of adopting a policy h_t at time t and state s_t is given by the function $v_t(h_t, s_t)$:

$$v_t(h_t, s_t) = \mathbf{E}_t \sum_{\tau=t}^{T} g_\tau(d_\tau, s_\tau), \tag{1}$$

where $g_t(.)$ is a period value function, and $\mathbf{E}_t$ is an expectations operator.[8] Given state vector s_t and decision $d_t(s_t)$, the state vector next period evolves according to the Markov transition matrix $\pi(s_{t+1}|s_t, d_t)$. Letting $\mathcal{H}_t$ denote the set of all feasible policies at time t, the maximum obtainable value of $v_t(h_t, s_t)$ is given by:

$$v_t^*(s_t) = \sup_{h_t \in \mathcal{H}_t} v_t(h_t, s_t), \ \forall s_t \in \mathcal{S}_t. \tag{2}$$

Define a rule of thumb to be a policy $\tilde{h}_t(\theta)$ in which the decision rules are constrained to take a specific functional form, $d_t = f(s_t; \theta)$, where θ is a vector of parameters. For example, the 100-minus-age rule of thumb for portfolio choice is one parameterization of the linear function, $f(age_t; \theta) = \theta_0 - \theta_1 age_t$, with $\theta_0 = 100$ and $\theta_1 = 1$. The value of adopting the rule $\tilde{h}_t(\theta)$ is given by:

$$v_t(\tilde{h}_t(\theta), s_t) = \mathbf{E}_t \sum_{\tau=t}^{T} g_\tau(f(s_\tau, \theta), s_\tau). \tag{3}$$

[7] Lettau and Uhlig (1999) approximate the optimal solution using a "classifier system" from the field of artificial intelligence and apply the method to a stylized consumption problem. The model makes an important contribution to dynamic learning, but it does lend itself to the more complex versions of the life-cycle model considered here.

[8] Note that the period value function, $g_t(.)$, depends on the time period. In a life-cycle model with utility function $u(.)$ and bequest function $B(.)$, for example, $g_\tau = \beta^{\tau-t}u(C_\tau)$ if the individual is alive in period τ, and $g_\tau = \beta^{\tau-t}B(X_\tau)$ otherwise, where β is the discount factor, C_τ is consumption, and X_τ is cash on hand.

Define an *optimal rule of thumb* to be a parameterization θ^* that maximizes the value of adopting a particular rule of thumb, $\tilde{h}_t(\theta)$:

$$\theta^* = \arg\max_{\theta} v_t(\tilde{h}_t(\theta), s_t). \tag{4}$$

Returning to the linear rule of thumb in portfolio choice, the optimal rule of thumb is characterized by the parameter pair, $\{\theta_0^*, \theta_1^*\}$, that maximizes the expected discounted present value of lifetime utility. In some cases the optimal rule of thumb will be the same regardless of the time period, but this will not true in general. The rules of thumb that perform the best in early periods may differ from those that perform well in later ones.

3.1. Updating the rule

The analysis above assumes that individuals commit to using a particular rule for the remainder of life. A more realistic assumption would grant individuals the option to update their rules after a certain amount of time. Begin by considering the simplest case of updating, where individuals can update to a more efficient rule at some future time period $t = t_1$. Let θ_{t_1} denote the parameter vector in the updating period and θ_0 denote the parameter vector in the first period. The value associated with choosing θ_0 and θ_{t_1} when the first period's state is s_0 is given by:

$$v_0^u(s_0, t_1) = v_0(\tilde{h}_0(\theta_0), s_0) + \beta^{t_1}\mathbf{E}_0\left[v_{t_1}(\tilde{h}_{t_1}(\theta_{t_1}), s_{t_1}) - v_{t_1}(\tilde{h}_{t_1}(\theta_0), s_{t_1})\right], \tag{5}$$

where β is a discount factor, and $v_0(.)$ and $v_{t_1}(.)$ are defined in equation (3). The first term in equation (5) is the value of sticking with the rule chosen in period 0, while the second term is the discounted expected value of the option to switch to θ_{t_1}. More generally, suppose that the individual updates the rules in periods $t_1, t_2, ..., t_m$. Letting $t_0 = t$, the value of updating is given by

$$v_t^u(s_t, t_1, ..., t_m) = v(\tilde{h}_t(\theta_t), s_t) + \sum_{i=1}^{m}\beta^{t_i}\mathbf{E}_t\left[v_{t_i}(\tilde{h}_{t_i}(\theta_{t_i}), s_{t_i}) - v_{t_i}(\tilde{h}_{t_i}(\theta_{t_{i-1}}), s_{t_i})\right]. \tag{6}$$

An *optimal rule of thumb with updating* is given by the parameterization $\theta^*_{t_0}, ..., \theta^*_{t_m}$ that maximizes equation (6):

$$\{\theta^*_{t_i}\}_{i=0}^m = \arg\max_{\{\theta_{t_i}\}_{i=0}^m} v_t^u(s_t, t_1, ..., t_m). \tag{7}$$

As with all options, the value of the option to update must be non-negative. There are at least two ways to compute the value functions with updating. The first method uses information about the distribution of the state vector in the updating periods, $t_1, t_2, ..., t_m$, to compute the expected option values of updating. Gomes and Michaelides (2005) show how this can be done in the context of optimal saving and portfolio choice using the decision rules and the discretized distribution of returns and permanent income. The second method for computing the value function in equation (6) takes advantage of the recursive structure of dynamic programming. Suppose that we are considering only a single updating period, t_1. For each feasible state vector, s_{t_1}, it is possible to compute an optimal rule of thumb that satisfies the definition given in equation (4), with the corresponding value function $v_{t_1}(\tilde{h}_{t_1}(\theta^*_{t_1}), s_{t_1})$. Moving back one period to $t_1 - 1$, the method substitutes $v_{t_1}(.)$ for next period's value function, regardless of the policy rule implemented in period $t_1 - 1$.

3.2. Robust rules of thumb

In their definition of a rule of thumb, Baumol and Quandt (1964, pg. 24) require that "the variables which are employed in the decision criteria are objectively measurable" and that "decision criteria are objectively communicable, and decisions do not depend on the judgment of individual decision makers." While these requirements still permit a wide range of functional choices for the decisions rules, they arguably rule out functions that depend on subjective beliefs or preference parameters.

But even if the rule itself does not depend explicitly on individual preferences, a "good" rule of thumb should be robust to observed variation in the parameters of the utility function. Suppose, for example, that a financial planner would like to offer the "best" rule of thumb for an individual investor of a given age, education, wealth level, and so on. The planner does not know the individual's risk aversion with precision, but instead has an idea of the parameter's distribution. In this situation, the planner may want to select a rule of thumb that minimizes welfare losses taking into account any uncertainty about the preference parameters. An *optimal robust rule of thumb* minimizes the expected loss associated with adopting a rule of thumb, where the expectations take into account both state and preference uncertainty.

Let ξ be a preference parameter with cumulative distribution function $G(\xi)$. An optimal robust rule of thumb solves:

$$\theta_R^* = \arg\max_{\theta} \int v_t(\tilde{h}_t(\theta), s_t, \xi) dG(\xi). \tag{8}$$

The optimal robust rule explicitly accounts for variation in the investor's parameter vector. In practice, however, there are limitations to the amount of uncertainty the model can handle due to the curse of dimensionality.[9]

4. Saving and portfolio choice

With minor variations, I adopt the model of consumption and portfolio choice in Cocco et al. (2005). Time is discrete. The individual lives for a maximum of T periods, retires at date T_R, and lives from one period to the next with probability $\psi(t)$. In each period t, the individual consumes C_t and allocates ς_t percent of wealth in the risky asset, which offers a gross rate of return R_t^s, and allocates the remainder in the risk-free asset, which offers a gross return R^f. Saving and consumption must be financed out of cash on hand of X_t, which consists of saving from the previous period plus current income, Y_t:

$$X_t = R_t(X_{t-1} - C_{t-1}) + Y_t, \tag{9}$$

where $R_t = \varsigma_t R_t^s + (1 - \varsigma_t)R^f$ is the gross portfolio rate of return. In the versions of the model that focus on consumption rules of thumb, I assume that individuals can only invest in the risk-free asset, which is equivalent to requiring $\varsigma = 0$.

[9] This paper uses a simple grid search method to solve for the optimal rule-of-thumb parameters. If there are M discrete parameter values and P points in the discretized distribution of preferences, the state space expands by a factor of $M \times P$ relative to that in a conventional dynamic programming model.

Following Carroll (1997), the income process consists of a deterministic function of age, a transitory shock, and a random-walk persistent shock. Permanent income, P_t, evolves according to $P_t = P_{t-1}G_tN_t$, where G_t captures the age earnings profile, and N_t is a log-normally distributed shock. Current income then equals the realized value of permanent income times a log-normally distributed transitory shock, Θ_t: $Y_t = P_{t-1}G_tN_t\Theta_t$. In some specifications of the model, current income in retirement is interpreted as income net of medical costs (see Love, 2010), so that uncertainty reflects changes in medical expenditures rather than labor income.

Individuals value consumption according to the isoelastic function $u(C_t) = C_t^{1-\rho}/(1-\rho)$, and they value bequests according to the function $B(X_t) = b\left(X_t/b\right)^{1-\rho}/(1-\rho)$. Regardless of the policy governing consumption and portfolio choice, utility in period t is given by:

$$U_t = \mathbf{E}_t \sum_{i=0}^{T-t} \beta^i \left[\Psi_{t+i,t}u(C_{t+i}) + (1-\Psi_{t+i,t})B(X_{t+1})\right], \tag{10}$$

where β is the time-invariant discount factor and $\Psi_{t+i,t}$ is the probability of surviving to period $t+i$ conditional on being alive in period t. The value function for the consumer's problem is then given by:

$$V_t^*(X_t, P_t) = \max_{C_t,\varsigma_t} \left\{u(C_t) + \beta\psi_t\mathbf{E}_tV_{t+1}^*(X_{t+1}, P_{t+1}) + \beta(1-\psi_t)\mathbf{E}_tB_{t+1}(R_{t+1}(X_t - C_t))\right\}, \tag{11}$$

subject equation (9), where ψ_t is the conditional probability of surviving to period $t+1$ given that the individual is alive in period t. I normalize the problem by permanent income, and then solve the model using Carroll's (2001) method of endogenous grid points. A detailed account of the solution can be found in the appendix.

4.1. Welfare

Following Cocco et al. (2005), I compute welfare costs using a measure of equivalent consumption. In particular, I solve for the constant stream of lifetime consumption that would deliver an equivalent amount of lifetime utility as would be obtained by applying a particular policy, whether it be optimal or a rule of thumb, for the remainder of life. The welfare cost associated with following a given policy is then the percentage increase in consumption required to make the individual indifferent between that policy and the optimal one. Let V_t be the value associated with a rule of thumb, and let V_t^* be the value associated with the optimal policy. The consumption equivalent streams for V_t and V_t^* are implicitly given by:

$$V_t = \sum_{i=t}^{T} \beta^{i-t} \Psi_{t+i,t} U(C), \tag{12}$$

and

$$V_t^* = \sum_{i=t}^{T} \beta^{i-t} \Psi_{t+i,t} U(C^*). \tag{13}$$

Letting $\kappa = 1/\sum_{i=t}^{T} \beta^{i-t}\Psi_{t+i,t}$, $C = U^{-1}(\kappa V_t)$ and $C^* = U^{-1}(\kappa V_t^*)$. Homogeneity of $U(.)$ implies that

$$\frac{C^* - C}{C} = \frac{U^{-1}(V_t^*) - U^{-1}(V_t)}{U^{-1}(V_t)}. \tag{14}$$

Although this formulation of welfare is standard, one may reasonably object that the kinds of households who find rules of thumb appealing are unlikely to arrive at a close approximation of their own expected discounted lifetime utility, even with rules of thumb substituting for more advanced rules derived from a dynamic programming problem. The rules may be simple, but forecasting the welfare effects of applying those rules requires a sophisticated understanding of actuarial and financial risk, as well as a stable set of intertemporal preferences. But in a sense, this is in keeping with the normative spirit of the exercise as long as the supplier of the optimal rules is a sophisticated financial planner rather than an individual decision-maker.

4.2. Income process

I estimate income profiles and the covariance structure of earnings using panel data from the 1970–2007 waves of the Panel Study of Income Dynamics (PSID). As is standard in the literature, I estimate separate income processes for household heads with less than 12 years of education, 12–15 years of education, and 16 or more years of education. The sample is restricted to households with male heads who are not part of the SEO oversample of low-income households. Income is a post-government concept that sums total family labor income, public and private transfers, and public and private pensions and subtracts total taxes, all deflated into 2010 dollars using the CPI-U. Because portfolio returns are endogenously determined in the model, I exclude all sources of asset income from the measure of income. This definition of family income is appropriate for thinking about the impact of background risk coming from an uncertain stream of income over the life cycle.

For each education group, I estimate fixed-effects regressions of the natural logarithm of income on a full set of age dummies, marital status, family composition, and year dummies. I run the regressions for respondents aged 20–65 for high school graduates and dropouts, and for respondents aged 22–65 for college graduates, excluding students and retirees. Because the estimates of permanent and transitory variance are sensitive to extreme outliers, I also eliminate the top 0.5% of one-year and two-year changes in income.[10] This eliminates 157 observations for the college sample, 311 for high school graduates, and 93 for dropouts.

Following Cocco et al. (2005), I construct income profiles by fitting a third-degree polynomial to the full set of age dummies for each education group and adding the regression constant and the coefficient on married. The profiles therefore represent the average income trajectories for married households without additional household members living at home. As a result, they suppress the effects of potentially important changes in family composition over the life cycle due to children and transitions in marital status (Love, 2010). To compute the replacement rate of income in retirement, I first calculate the average post-government income for households aged 65–85, whose respondents report working less than 300 hours in the year. The replacement rate is then set to the ratio of retirement income to the profile income in the period just before retirement.

Carroll and Samwick (1997) develop an efficient way to estimate the variance structure of income. After stripping away the trend component of income growth (predicted family income based on the full set of covariates), they show that the d-year difference in log incomes, denoted r_d, is given by the d-year difference in permanent and transitory income: $r_d = p_{t+d} - p_t + \epsilon_{t+d} - \epsilon_t$, where $p_t = \ln(P_t)$, and ϵ_t is the transitory shock. Letting η_t denote the permanent

[10] I treat outliers for both one-year and two-year changes because the PSID switches from an annual to a biannual survey after 1997.

shock in period t, the income difference can be written as the cumulative sum of permanent shocks and the d-year difference in the transitory shock:

$$r_d = \sum_{i=1}^{d} \eta_{t+i} + \epsilon_{t+d} - \epsilon_t.$$

The variance is therefore given by:

$$Var(r_d) = d\sigma_\eta^2 + 2\sigma_\epsilon^2.$$

For each household, I compute r_d for all values of $d > 2$. (Carroll and Samwick show that the procedure is robust to serial correlation in the transitory error up to $MA(q)$ so long as $d > q$.) I then estimate σ_η^2 and σ_ϵ^2 by running an OLS regression of the variance of r_d on a vector of d's and a constant vector of 2's.

Some parameterizations of the model assume that there is no uncertainty in retirement income, which is a reasonable assumption if one focuses on the annuity payments derived from public and private pensions. But as several researchers have emphasized (see, e.g., Palumbo (1999) and French and Jones (2011)), out-of-pocket medical costs can lead to large variations in the amount of retirement resources net of medical expenses. Medical costs can be viewed as responding to sudden increases in the marginal utility of expenditures due to a deterioration in health status or as an exogenous change in necessary expenditures; either way, uncertain medical costs constitute a source of background risk that may affect the demand for risky assets. Love (2010) estimates the variance process by education group for income net of medical costs using panel data from the 1992–2006 waves of the Health and Retirement Study. I adopt those estimates for the model specifications that allow for retirement risk.

Table 1 reports the polynomial coefficients on age, the replacement rate, and the variance decomposition for the working and retirement period. Figure 1 shows the resulting average income profiles for each of the education groups. The age pattern of income follows the familiar hump shape, with peaks between ages 50 and 60. While incomes for dropouts and high school graduates closely track one another (apart from a level difference), college graduates have a much steeper profile earlier in life. In terms of the variance structure, the estimates for the transitory variances are significantly higher than those found in Carroll and Samwick (1997) and slightly higher than those in Cocco et al. (2005), while the estimated permanent variances are somewhat smaller.[11]

4.3. Asset returns

Campbell and Viceira (2002) estimate the standard deviation and mean of the excess returns of stocks over the risk-free rate using annual data on the S&P 500 for the period 1880–1995.[13] They estimate a mean excess return of 6.24% and a standard deviation of 18.11%. (The postwar data series shows a higher excess returns and lower volatility. Using quarterly data for 1952–1997, they find a mean return of 7.12% and a standard deviation of 6.10%.) I set the standard deviation of the stock return to 18%, the risk-free rate to 2%, and the excess return equal to 4%, which is about two percentage points below the long-run average. As Mehra and Prescott (1985) famously demonstrated, it is difficult to reconcile the high historical risk premium with conventional assumptions about risk aversion. Some combination of a lower premium and higher risk aversion is necessary to push households away from a leverage-constrained corner solution of 100% stocks.

The model also allows asset returns to be correlated with permanent income. As reported in Gomes and Michaelides (2005), Campbell et al. (2001) estimate a correlation coefficient of about 15% between permanent income and excess returns. Cocco et al. (2005), in contrast, also estimate the correlation coefficient and cannot reject the possibility that it is zero. I solve some versions of the model assuming a 15% correlation during the working life and a 0% correlation in retirement (since annuitized income should not relate in any significant way to the performance in the stock market). In other versions of the model, I set the correlation to

[11] Carroll and Samwick only use the 1981–1987 waves of the PSID and exclude households whose income fell below 20% of their average over the sample period. Thus, part of the difference between my estimates and Carroll and Samwick's arises from the difference in sample periods and the criteria for removing outliers. The difference between my estimates and those in Cocco, Gomes, and Maenhout, in contrast, is largely explained by the fact that I use only $d > 2$ year income differences in estimating the variance structure in order to account for the possibility of MA(2) serial correlation in the transitory shock, whereas they use all combinations of r_d.[12] If I estimate the variances using all values of d, the estimates are close to those in Cocco, Gomes, and Maenhout (2005).

[13] The data were originally compiled byGrossman and Shiller (1981) and later updated by Campbell (1999).

zero for the entire life-cycle.

4.4. Preferences and other parameters

Table 2 lists the set of parameters used to solve the baseline model. The discount factor is set to 0.98, which is in in the range of structural estimates.[14] Survival probabilities come from the 2007 Social Security Administration Period Life Tables. I do not model spousal mortality and use only male survival probabilities. This is a common practice in the life-cycle literature, but it has the drawback of shortening the effective decision horizon relative to a more realistic description of a two-person family unit.

The baseline parameterization for the portfolio choice model sets the coefficient of relative risk aversion to 5, which is at the upper range of most structural estimates.[15] As mentioned above, lower values of risk aversion, in conjunction with large observed risk premia on stocks, tend to generate corner solutions in portfolio choice. In the models without portfolio choice, I set the baseline value of risk aversion to 3, which is more in line with structural estimates.

Finally, I solve the model both with and without an active bequest motive. When the motive is operational, I follow Gomes and Michaelides (2005) and set the bequest parameter $b = 2.5$.

4.5. Optimal solution

Figure 2 shows the average of 20,000 simulated paths of consumption, wealth, income, and the portfolio share for a college graduate without a bequest motive. (The profiles for high school graduates and dropouts look similar but have different levels of wealth.) The top panel of the figure indicates that wealth reaches a peak around $800,000 at retirement and then declines gradually during retirement. Consumption continues to grow throughout most of retirement

[14]Cagetti (2003), for example, estimates the discount rate and risk aversion pair that minimizes the distance between mean and median wealth levels in the Survey of Consumer Finances (SCF) and the Panel Study of Income Dynamics. Depending on education groups, the estimates based on median wealth in the SCF range between 0.923 for high school graduates and 0.988 for college graduates. French (2005) estimate a structural life-cycle model incorporating labor supply and health, and they find values of β in the range of 0.981–1.04, depending on the specification.

[15]For example, Cagetti's (2003) estimates of risk aversion (see footnote above) range between 2.57 and 4.05, while French (2005) reports values of risk aversion between 2.2 and 5.1.

and only comes back down due to the increasing rate of mortality discounting. The high growth rate of consumption is due to two influences: the high portfolio rate of return relative to the discount rate, and the need to maintain a precautionary buffer of savings against medical cost shocks. Turning to the bottom panel of the figure, the portfolio share in stocks remains at a corner solution of 100% stocks until around age 30, after which point it declines to around 40% at retirement. During retirement, the share flattens and then rises slightly near the tail end of life.

The path of the optimal stock share reflects the changing importance of financial wealth in financing lifetime consumption (Bodie, Merton, and Samuelson, 1992). In early years, financial wealth is low, and consumption needs are mostly financed out of human capital. At this stage, even large fluctuations in the value of stocks would have only a minor impact on consumption risk. During the working years, financial wealth plays an increasingly important role in financing lifetime consumption (financial wealth rises while human capital declines), and the optimal allocation becomes more conservative. Finally, in retirement, financial wealth and human capital decline in tandem, and consumption can depend either more or less on financial resources. If the ratio of financial wealth to total resources remains stable, the optimal portfolio share should remain approximately stable as well.

Looking at the declining average shares in Figure 2, one can see why a rule of thumb that decreases the stock share with age might hold promise. It is important to keep in mind, however, that the path in the figure represents the average of 20,000 portfolio decisions at each age. Individual decisions about saving and portfolio allocation can differ substantially from those depicted in the average profiles. Figure 3 shows how the 10th, 50th, and 90th percentiles of the allocation distribution compare to the mean allocation displayed in the figure above. For most of the working life, the shares at the 10th and 90th percentiles are separated by more than 20 percentage points, and the difference widens even more in retirement. This spread in the optimal allocations at each age highlights a potential drawback of using "one-size-fits-all" rules of thumb that do not respond to changing financial circumstances.

5. Rules of thumb for savings and portfolio choice

This section explores the welfare consequences of adopting optimal rules of thumb for saving and allocation decisions under different assumptions about updating and preference uncertainty. The optimization of the rules is performed using a simple grid search over a subset of the parameter space.[16] Despite the inefficiencies of a grid search, the method has the advantage of locating a close approximation of the global optimum even in the likely presence of multiple local optima.

5.1. Portfolio choice rules

I consider two functional forms for portfolio rules. The first constrains portfolio weights to be a linear function of age, and the second constrains them to be a linear function of the ratio of financial wealth to total wealth, including human capital. The motivation for the first rule comes from the common financial advice that investors place a percentage equal to 100 minus their age in equities. As Cocco, Gomes, and Maenhout (2005) have shown, this rule is consistent with the general reduction in the optimal exposure to equities with age during the middle portion of the life cycle, but it does a poor job of matching optimal allocations both early and late in the life cycle.[17]

The second rule is motivated by the observation that age matters in portfolio choice primarily because it varies systematically with the ratio of financial to total wealth over the life cycle. The more that the ability to finance consumption depends on financial wealth, the more sensitive investors should be to the consequences of financial market risk. The second rule sets the portfolio share in stocks equal to a linear function of the ratio, $\mu(.)$, of end-of-period savings to total wealth, comprising savings and human capital:

$$\mu(A_t, H_t) = \frac{A_t}{A_t + H_t},$$

[16]The spacing of the grid points, as well as the lower and upper bounds of the search space, differ across the various models. The number of parameter combinations (and therefore separate models that need to be solved) range from about 250 to 20,000, depending on the range and number of parameters.

[17]Early in the life cycle, their model predicts that most investors are actually leverage constrained to hold 100% in equities, and later in life, the optimal shares either flatten out or even rise as background risk diminishes and financial wealth becomes a less important source of consumption finance relative to future pension income.

where $A_t = X_t - C_t$ is end-of-period savings, and H_t approximates the expected present discounted value of future income. In constructing a measure of future income, the goal was to come up with a methodology that would be accessible to households with limited access to actuarial and financial information. For the purposes of the second rule of thumb, I define human capital as the sum of permanent income, discounted using the risk-free rate, over a horizon equal to the expected remaining years of life. That is:

$$H_t = \mathbf{E}_t \sum_{i=1}^{n} (1 + r^f)^{-i} P_{t+i}, \tag{15}$$

where n is the number of additional expected life years (rounded up), not including the current period. To compute this ratio, households need to know the risk-free rate, their expected longevity, and an estimate of their remaining income stream. They also need to know how to compute present values. Thus, relative to an age-based rule, which takes seconds to calculate, a rule based on the ratio of financial to total wealth requires more information and a more sophisticated calculation.[18] An interesting question is whether the welfare gains associated with using a more sophisticated wealth-based rule are "enough" to justify the calculation costs.

5.1.1. Performance of portfolio rules

I begin with simple allocation rules of the form: $\varsigma_t = 1 - \theta_1 S_t$, where S_t is a state variable equal to either age or the ratio of financial wealth to total wealth (μ), θ is a slope parameter, and households are prohibited from taking either a short or leveraged position in stocks. Figure 4 plots the optimal values of θ over the course of an average life cycle for both the age-based rule (left panel) and the wealth-based rule (right panel), assuming that households maintain the rules from each age going forward.[19] In both cases, the optimal value of the parameter for a

[18] Present discounted value calculations may not seem forbidding to economists, but Lusardi and Mitchell (2011) report that only about half of Americans aged 50 and older can answer two (very) simple questions about compound interest and inflation. The interest rate question, for example, asks: "Suppose you had \$100 in a savings account and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow? More than \$102, Exactly \$102, Less than \$102, Do not know?"

[19] The jaggedness in the right panel of the figure is due to the fact that the wealth rule approximates human capital using the rounded value of expected remaining life years, leading to discrete changes in human capital as

20-year-old college graduate is around 0.95. (The figures for the other education groups display similar patterns but with slightly lower parameter values.) Interestingly, the optimal age rule for a young worker corresponds quite closely to the conventional advice to place 100 minus one's age in stocks. As the individual approaches retirement, however, the slope coefficient declines substantially, falling to around 0.7 in retirement. The decline in the value of θ reflects the desire of households to maintain a somewhat higher concentration of wealth in stocks during the tail end of the life cycle (see Figure 2).

The optimal value θ declines in the case of the simple linear wealth rule as well. But here, the decline is slighter and more gradual, averaging around 0.85 in retirement. Thus, compared with the linear age rule, individuals would seem to have less incentive to update the rule at later ages. Figure 5 compares the welfare losses (measured in dollars of additional annual consumption) associated with maintaining the simple age- and wealth-based rules that are optimal for a 20-year-old for the duration of life. While the welfare losses do not favor either rule unambiguously during the working period, the wealth rule performs dramatically better in the retirement period. The age rule is effective in earlier periods because it mimics the tendency of the optimal allocation to fall as an increasing share of present and future consumption is financed out of financial assets. In retirement, however, assets decline at a similar rate as the reduction in human capital, and the importance of assets in financing consumption remains relatively stable. The wealth rule responds to this change, while the age rule does not.

The welfare losses associated with using the simple linear rules above are relatively modest—averaging between 0.5 and 1 percent of annual consumption. One can do even better, however, by optimizing over both the intercept and the slope. The next rules I consider take the form: $\varsigma_t = \theta_0 - \theta_1 S_t$, where S_t again denotes either age or the ratio of financial wealth to total wealth, and I again rule out both short and leveraged positions in the stock market. Table 3 reports the welfare losses and the optimal values of the parameters for the age- and wealth-based rules, assuming that the rules are maintained from each age going forward. The first thing to notice about the results in the table is that the optimal parameter pair changes substantially

expected remaining years fall from one value to another. These plots are based on the average wealth trajectory over the life-cycle, so these discrete movements are not smoothed over different realizations of wealth.

as individuals age, particularly in the case of the age rule. The optimal age rule, in percentage points, for a 20-year-old college graduate is approximately 120 minus 1.3 times age, which is not too far off the modified "Malkiel rule," which suggests placing 120 minus one's age in stocks. At age 40, however, the optimal rule changes to 80 minus 0.7 times age, and at 60 the rule is 42.5 minus 0.13 times age. Thus, from age 20 to 60, the slope coefficient falls by an order of magnitude. The optimal parameters for the wealth rule change with age as well, but not nearly as much. The results in the table suggest that the rule of thumb that looks best from the perspective of a young worker may look less attractive later in life.

What are the welfare consequences of adhering to the optimal rule for a 20-year-old for the remainder of life? Figure 6 displays the welfare losses, expressed in dollars of additional consumption, associated with maintaining either the optimal age or the optimal wealth rule. In this case, the wealth-based rule strictly dominates the age-based rule, and the welfare losses diverge dramatically in the retirement period. The figure implies that a typical college graduate would incur welfare losses of less than $200 annually by using the same wealth-based rule in each and every year of life. The optimal wealth rule—approximately 75% times 100 minus the percentage ratio of financial wealth to total wealth—is less catchy than the popular 100-minus-age rule, but the improvement in performance may be worth it.

The welfare losses in Figure 6 suggest that there may be substantial gains associated with updating the rules at later ages, particularly in the case of the linear age rule. Table 4 reports the optimal parameter values and welfare losses assuming that the individual can update to the optimal portfolio rules at ages 20, 40, and 65, respectively. Not surprisingly, the optimal rules at age 20 reflect a more aggressive allocation toward stock. Younger households build up little financial wealth relative to human capital and therefore should not be as sensitive to fluctuations in asset markets. Further, because the households can update to a less stock-heavy allocation at middle age and retirement, they can afford to invest aggressively early in life. With the possibility of updating, the welfare losses associated with adopting optimal rules of thumb fall to 0.12 percent of annual consumption in the case of a linear age rule for college graduates, and 0.06 percent of annual consumption in the case of the linear wealth rule. These are extremely small welfare losses relative to the dynamic programming solution—on the order of about \$30 to \$60 per year.

[Under construction] Further, the results above are robust to uncertainty in the assumed values of risk aversion. I solve for optimal rules of thumb assuming that risk aversion is normally distributed with a mean of 5 and a standard deviation of 1. The optimal rule, in this case, minimizes the expected welfare losses, where the expectation is taken over the possible values of the preference parameters (see Section 3.2).[20] In general, the robustness exercise leads to relatively small changes in the optimal parameter values, but it has a measurable effect on the welfare losses. For example, consider the human capital rule for college graduates. With uncertainty in risk aversion, the optimal values of each of the parameters fall by about 0.025, and the welfare loss rises from 0.405% to 0.588% of annual consumption at age 20. Thus, introducing uncertainty about the value of risk aversion reduces the attractiveness of rules of thumb, but it does not lead to large changes in the rules themselves.

5.2. Consumption rules

I consider two classes of consumption rules. The first rule mirrors the common recommendation that households save a constant fraction (usually 10–20%) of income during the working years and withdraw savings at a constant rate (usually 4–5%) in retirement. I assume that households must consume a constant fraction, θ_0, of permanent income during the working life if this amount exceeds cash on hand, and cash on hand otherwise. In retirement, households consume the minimum of cash on hand and the sum of permanent income and a constant fraction, θ_1, of savings.

The second rule is more in line with the consumption smoothing property implied by the standard life-cycle framework. According to the second rule, households consume a constant fraction, θ, of annuitized total wealth (savings plus human capital). Human capital is given by equation (15), and the annuitization factor, AF_t, is given by:

$$AF_t = \frac{1}{\sum_{i=1}^{n}(1+r)^{-i}},$$

[20]The distribution of risk aversion is approximated using 8-point Gauss-Hermite quadrature. The consumption equivalent measure of welfare in this case is given by $\sum_i^n w_i(C_i^*/C_i - 1)$, where the w_i's are the weights associated with the respective values of risk aversion (i.e., the quadrature abscissa).

where n is the expected years of remaining life. Consumption is then equal to the minimum of cash on hand and annuitized total wealth:

$$C_t = \min(X_t, \theta AF_t(H_t + X_t)). \tag{16}$$

While the second consumption rule involves a more complex calculation than the first one, it requires only a simple spreadsheet listing expected income, longevity, and asset returns.

Because the focus of this section is on the performance of consumption rules, I shut down the portfolio decision by assuming that households only invest in a risk-free asset offering a 3% real return. I also modify the baseline assumption of risk aversion by setting it to the more conventional value of 3.

5.2.1. Performance of consumption rules

The first rule of thumb has households saving a constant fraction of permanent income, θ_0, while young and withdrawing a constant fraction, θ_1, from savings when old. The left panel of Table 5 displays the optimal fractions and associated welfare losses from the perspective of different ages, assuming that individuals adhere to the rules for the remainder of life. For both high school and college graduate, the optimal amount of saving amounts to about 10% of permanent income, which is right in line with the common financial advice quoted in the first paragraph of the introduction. The optimal saving rate rises by about 4 percentage points in middle age for college graduates and falls by about 2 percentage points for high school graduates. The difference can be explained by the fact that college graduates have much steeper income profiles than high school graduates and consequently have a stronger incentive to defer saving to future periods.

The optimal rate of dissaving in retirement remains steady at 7–8% for both education groups until later in the retirement period, at which point the optimal withdrawal rates increase substantially. The combination of medical expense risk and uncertain longevity induce households to maintain a sizable amount of savings in early retirement, but eventually mortality discounting provides a strong incentive to consume a higher fraction of remaining wealth.

Interestingly, in all cases, the optimal rule of thumb for withdrawal rates is at least 3 percentage points higher than the standard recommendation of 4%, suggesting that the conventional advice may be inducing households to draw down wealth too slowly.

In contrast to the simple rules of thumb for portfolio choice, the welfare consequences of using the first rule of thumb are substantial. They amount to 4–5% of annual consumption from the perspective of a 20-year-old, and they rise to 6–7.5% of annual consumption (depending on education) from the perspective of a 40-year-old. While large, it is worth noting that these losses are actually one-half to one-third the size of the welfare losses found by Rodepeter and Winter (1999) using more sophisticated rules in a simpler life-cycle setting. Nevertheless, the size of the welfare losses suggests that one should be cautious recommending or adopting a simple rule based on fractions of permanent income and retirement savings, even if that rule is the best one of its class.

The right panel of Table Table 5 reports the optimal parameters and associated welfare losses for the more sophisticated rule based on a fraction, θ_0, of annuitized present value resources. The first thing to note about the parameters is that they remain remarkably consistent across the different ages. The fraction for college graduates, for example rises from 64% at age 20 to about 70% at age 40, with only minor changes at ages 65 and 80. The fractions for high school graduates are about 10 percentage points higher, but they are even more compressed across age groups. Second, looking at these fractions, a natural question is why they are not closer to 100%, which is what a simple version of a life-cycle model with no sources of risk in longevity, asset returns, or income would predict. The answer is precautionary saving. When I solve a model shutting down all sources of risk, the optimal fractions remain constant at 1. The final thing to note is that while the welfare losses from the perspective of a 20-year-old are sizable, at around 2–3.6% of annual consumption, they fall substantially in later years. The key is again precautionary saving. The saving rule is well designed for smoothing resources over the life-cycle, but it does not adjust for the changing importance of buffer-stock saving. The rule therefore imposes larger costs early in life, when current income is low relative to future income.

Table 6 shows how the results change in the presence of uncertainty in either risk aversion or the discount factor. The thought experiment here is to imagine that a financial advisor would like to offer the best consumption rule of thumb knowing only the distribution of risk aversion or impatience. The table reports the optimal parameter values and welfare losses at age 20 for the two types of consumption rules assuming: no uncertainty ("Baseline"), uncertainty in risk aversion ("CRRA"), and uncertainty in the discount factor ("Discount").[21] For both sets of rules, preference uncertainty leads to small changes in the optimal parameter values, but it has a larger impact on welfare losses. Depending on the source of uncertainty, the welfare losses increase by about 0.4–2 percent of annual consumption, with the largest changes arising in the case of an uncertain discount rate. The importance of discount-rate uncertainty should not be surprising in light of its crucial role in wealth accumulation. Adhering to an optimal rule for "typical" households may prove costly to households with divergent levels of patience or, to a lesser extent, risk aversion.

As Table 7 indicates, the ability to update the rules at ages 40 and 75 reduces the welfare costs at age 20 by 1–1.5 percentage points for both sets of rules, with somewhat smaller reductions at age 40.[22] Overall, the results in Tables 5–7 indicate that even an optimal rule based on simple fractions of income and wealth generates large enough welfare losses that it would pay individuals to seek out more sophisticated advice. The optimal rule based on annuitized present value resources, however, performs much better, and the welfare losses are as low as 1-2% of annual consumption. One interpretation is that these losses are surprisingly small given the simplicity of the optimal rule of thumb, especially in light of the much larger welfare losses found in previous work. Another interpretation is that welfare losses in the range of \$182–\$1,543 (see Table 7) may justify paying the required fees for more tailored life-cycle advice offered by companies such as ESPlanner. At the very least, it suggests that there may be large returns to investigating the performance of alternative optimal rules of thumb for consumption and saving.

---

[21]The uncertainty in risk aversion case assumes that ρ is normally distributed with a mean of 3 and a standard deviation of 0.5. The discount rate case assumes that β is normally distributed with a mean of 0.98 and a standard deviation of 0.02. Gauss-Hermite quadrature is used in computing the expected welfare losses.

[22]The reason that I focus on updating at age 75 instead of age 65 is that the optimal fixed-percentage withdrawal rule at age 65 with or without updating is exactly the same.

6. Conclusion

This paper develops a framework for selecting optimal rules of thumb. The framework allows individuals to update their rules at different ages, and it accommodates uncertainty about the underlying preference parameters. Applying the framework to the life-cycle problem, I find that optimal rules of thumb based on either age or the ratio of financial wealth to total wealth lead to only moderate welfare losses relative to the dynamic programming benchmark. In the case of the wealth-based rule, the losses are as low as 0.06% of annual consumption, suggesting that rules of thumb may indeed represent a viable alternative to more sophisticated solutions to portfolio choice.

Optimal rules of thumb are less effective when it comes to saving and consumption decisions. I examine simple rules of thumb modeled after the conventional advice that households save a fraction of income during the working years and withdraw wealth at a constant rate in retirement and find that even optimal rules generate welfare losses in the range of 4–7.5% of annual consumption. Rules based on the annuitized present value of total lifetime resources perform substantially better, with welfare losses as low as 1–2% of annual consumption, but even these losses are large enough to justify a search for more effective alternatives.

There are at least two ways to improve on the performance of the rules examined in this paper. The first would be to introduce elements of realism that are missing from the underlying model of the economy but that may play an important role in shaping decisions about saving and asset allocation. Extensions along these lines include housing (Cocco, 2005; Chetty and Szeidl, 2011), flexible labor supply (Gomes et al., 2008), annuitization (Yogo, 2011), stock-market participation costs (Gomes and Michaelides, 2005), taxation (Poterba and Samwick, 2003), family shocks (Love, 2010), and changes in spending needs over the life cycle (Attanasio and Browning, 1995). Another way to improve the performance of the rules would be to search for more flexible functional forms that respond to key state variables. Introducing a variable capturing the precautionary motive, for example, may lead to substantial improvements in the performance of life-cycle rules of thumb.

Appendix: Solving the Life-Cycle Model

As Carroll (2008) has shown, the assumptions that income follows a unit root and preferences are isoelastic imply that the problem can be normalized by permanent income. Letting $x_t = X_t/P_t$, the decision rules can be written as $c_t(x_t) = C_t(X_t, P_t)/P_t$ and $\varsigma_t(x_t) = \varsigma_t(X_t, P_t)/P_t$. Similarly, the normalized value is $v_t(x_t) = P_t^{\rho-1} V_t(X_t, P_t)$. The optimal solution to the problem is given by the value function:

$$v_t^*(x_t) = \max_{c_t, \varsigma_t} \left\{ u(c_t) + \beta\psi_t \mathbf{E}_t \Gamma_{t+1}^{1-\rho} v_{t+1}^*(a_t R_{t+1} + \Theta_{t+1}) + \beta(1-\psi_t)\mathbf{E}_t \Gamma_{t+1}^{1-\rho} B_{t+1}(a_t R_{t+1}) \right\}, \tag{A-1}$$

where $\Gamma_{t+1} = G_{t+1} N_{t+1}$ is the stochastic growth factor, and $a_t = x_t - c_t$ is end-of-period saving.

The first order conditions for the optimal solution are standard. First note that homogeneity of preferences implies that $\Gamma_t^{-\rho} u_t'(c_t) = u_t'(\Gamma_t c_t)$, and $\Gamma_t^{-\rho} B_t'(R_t a_{t-1}/\Gamma_t) = B_t'(R_t a_{t-1})$. The first order condition for consumption is the Euler equation:

$$u_t'(c_t) = \beta\psi_t \mathbf{E}_t R_{t+1} u_{t+1}'(\Gamma_{t+1}\mathbf{c}_{t+1}) + \beta(1-\psi_t)\mathbf{E}_t R_{t+1} B_t'(R_{t+1} a_t), \tag{A-2}$$

where $\mathbf{c}_{t+1} = \mathbf{c}_{t+1}(R_{t+1} a_t + \Theta_{t+1})$ is the decision rule for consumption in $t+1$. The first-order condition for portfolio choice balances the higher returns of equities against the utility consequences of risk:

$$\beta\psi_t \mathbf{E}_t (R_{t+1}^e - R^f) a_t u_{t+1}'(\Gamma_{t+1}\mathbf{c}_{t+1}) + \beta(1-\psi_t)\mathbf{E}_t (R_{t+1}^e - R^f) a_t B_t'(R_{t+1} a_t) = 0. \tag{A-3}$$

Given end-of-period saving a_t and a decision rule for next-period consumption, equation (A-3) can be used to find the optimal portfolio choice in period t for each a_t. With the optimal portfolio choice in hand, equation (A-2) determines the optimal level of consumption in period t for each level of a_t. Following Carroll's (2006) method of endogenous grid points, the corresponding level of normalized cash on hand is given by: $x_t = a_t + c_t(a_t)$. Interpolating points between the c_t and x_t pairs generates a consumption decision rule that can be used to solve for optimal consumption and portfolio choice in period $t-1$.[23] In practice, I use 30 grid points for end-of-period savings, spaced according to a triple exponential. I compute approximate the distributions for the asset returns and transitory and permanent income shocks using 10-point Gauss-Hermite quadrature.

[23] The solution method follows many of the tricks and suggestions in Carroll (2008), which result in dramatic improvements in solution time and accuracy. The improvements are crucial for identifying optimal rules of thumb since the complete life-cycle model will have to be solved numerous time in the search for an optimal parameterization of rules.

References

Allen, T. W., and C. D. Carroll. 2001. Individual Learning About Consumption. Macroeconomic Dynamics 5(02): 255–271.

Attanasio, O. P. 1999. Consumption. In *Handbook of Macroeconomics*, ed. J. B. Taylor and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 11, 741–812. Elsevier.

Attanasio, O. P., and M. Browning. 1995. Consumption Over the Life Cycle and Over the Business Cycle. The American Economic Review 85(5): 1118–1137.

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(3): 693–751.

Baumol, W. J., and R. E. Quandt. 1964. Rules of Thumb and Optimally Imperfect Decisions. The American Economic Review 54(2): pp. 23–46.

Becker, T. A., and R. Shabani. 2010. Outstanding Debt and the Household Portfolio. Review of Financial Studies 23(7): 2900–2934.

Bernheim, B. D., J. Skinner, and S. Weinberg. 2001. What Accounts for the Variation in Retirement Wealth among U.S. Households? The American Economic Review 91(4): pp. 832–857.

Bodie, Z., R. C. Merton, and W. F. Samuelson. 1992. Labor Flexibility and Porfolio Choice in a Life Cycle Model. Journal of Economic Dynamics and Control 16(3–4): 427–449.

Cagetti, M. 2003. Wealth Accumulation Over the Life Cycle and Precautionary Saving. Journal of Business and Economic Statistics 21(3): 339–353.

Campbell, J. Y. 1999. Asset prices, consumption, and the business cycle. In *Handbook of Macroeconomics*, ed. J. B. Taylor and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 19, 1231–1303. Elsevier.

Campbell, J. Y., J. F. Cocco, F. J. Gomes, and P. J. Maenhout. 2001. Investing Retirement Wealth: A Life Cycle Model. In *Risk Aspects of Social Security Reform*. University of Chicago Press, Chicago.

Campbell, J. Y., and N. G. Mankiw. 1989. Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence. NBER Macroeconomics Annual 4: pp. 185–216.

———. 1990. Permanent Income, Current Income, and Consumption. Journal of Business & Economic Statistics 8(3): pp. 265–279.

Campbell, J. Y., and L. M. Viceira. 2002. *Strategic Asset Allocation*. Oxford University Press.

Carroll, C. D. 1997. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. The Quarterly Journal of Economics 112(1): 1–55.

———. 2006. The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems. Economics Letters 91(3): 312–320.

———. 2008. Lecture Notes on Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems. Working Paper: Johns Hopkins University.

Carroll, C. D., and A. A. Samwick. 1997. The Nature of Precautionary Wealth. The Journal of Monetary Economics 40(1): 41–71.

Chetty, R., and A. Szeidl. 2011. The Effect of Housing on Portfolio Choice. NBER Working Paper 15998.

Cocco, J. F. 2005. Portfolio Choice in the Presence of Housing. Review of Financial Studies 18(2): 535–567.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout. 2005. Consumption and Portfolio Choice over the Life Cycle. The Review of Financial Studies 18(2): 491–533.

French, E. 2005. The Effects of Health, Wealth, and Wages on Labor Supply and Retirement. Review of Economic Studies 72(2): 395–427.

French, E., and J. B. Jones. 2011. The Effects of Health Insurance and Self Insurance on Retirement Behavior. Econometrica 79(3): 693–732.

Gomes, F. J., L. J. Kotlikoff, and L. M. Viceira. 2008. Optimal Life-Cycle Investing with Flexible Labor Supply: A Welfare Analysis of Life-Cycle Funds. The American Economic Review 98(2): pp. 297–303.

Gomes, F. J., and A. Michaelides. 2005. Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence. The Journal of Finance 60(2): 869–904.

Grossman, S. J., and R. J. Shiller. 1981. The Determinants of the Variability of Stock Market Prices. American Economic Review 71(2): 222–27.

Kotlikoff, L. J. 2007. Economics' Approach to Financial Planning. Working Paper.

Lettau, M., and H. Uhlig. 1999. Rules of Thumb versus Dynamic Programming. The American Economic Review 89(1): pp. 148–174.

Love, D. A. 2010. The Effects of Marital Status and Children on Savings and Portfolio Choice. Review of Financial Studies 23(1): 385–432.

Malkiel, B. G. 2011. *A Random Walk Down Wall Street: The Time-Tested Strategy for Successful Investing*. W. W. Norton & Company, New York, NY.

Mehra, R., and E. C. Prescott. 1985. The equity premium: A puzzle. Journal of Monetary Economics 15(2): 145 – 161.

Palumbo, M. G. 1999. Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle. Review of Economic Studies 66(2): 395–421.

Poterba, J. M., and A. A. Samwick. 2003. Taxation and Household Portfolio Composition: U.S. Evidence from the 1980's and 1990's. Journal of Public Economics 87: 5–38.

Scholz, J. K., A. Seshadri, and S. Khitatrakun. 2006. Are Americans Saving "Optimally" For Retirement? Journal of Political Economy 114(4): 607–643.

Simon, H. A. 1978. Rationality as Process and as Product of Thought. The American Economic Review 68(2): pp. 1–16.

Wachter, J. A., and M. Yogo. 2010. Why Do Household Portfolio Shares Rise in Wealth? Review of Financial Studies 23(11): 3929–3965.

Weber, C. E. 2000. "Rule-of-Thumb" Consumption, Intertemporal Substitution, and Risk Aversion. Journal of Business & Economic Statistics 18(4): pp. 497–502.

Winter, J. K., K. Schlafmann, and R. Rodepeter. 2011. Rules of Thumb in Life-Cycle Saving Decisions. Munich Discussion Paper 2011-16.

Yogo, M. 2011. Portfolio Choice in Retirement: Health Risk and the Demand for Annuities, Housing, and Risky Assets. AFA 2009 San Francisco Meetings Paper.

Table 1
Income process

	Dropout	High School	College
Fitted age polynomials			
Constant	-0.3474	-0.8002	-2.2218
Age	0.0194	0.0530	0.1395
Age^2/100	0.0611	-0.0090	-0.1504
Age^3/10000	-0.1114	-0.0618	0.0192
Replacement rate	81.76%	76.94%	75.67%
Coefficient estimates			
Married	0.2972	0.3962	0.4609
	(0.0151)	(0.0069)	(0.0115)
HH Size	0.0516	0.0329	0.0314
	(0.0042)	(0.0023)	(0.0037)
Constant	9.1732	9.7380	9.7602
	(1.3049)	(0.0283)	(0.0498)
N	1,721	4,614	2,494
R^2-within	0.149	0.225	0.330
F statistic	28.294	159.617	138.845
Variance decomposition: working life			
Permanent (σ_η^2)	0.0097	0.0087	0.0120
	(0.0006)	(0.0002)	(0.0003)
Transitory (σ_ϵ^2)	0.1203	0.0896	0.0851
	(0.0037)	(0.0013)	(0.0018)
Variance decomposition: retirement			
Permanent (σ_η^2)	0.0018	0.0125	0.0281
	(0.0015)	(0.0026)	(0.0027)
Transitory (σ_ϵ^2)	0.0911	0.0784	0.0767
	(0.0072)	(0.0070)	(0.0075)

This table presents the fitted age polynomials, coefficient estimates, and variance decomposition from separate fixed-effects regressions of the natural logarithm of income on a full set of age dummies, year dummies, marital status, and family composition (number of people living the household, excluding the respondent and spouse, if any). The data are taken from the 1970–2007 waves of the PSID. Income is the sum of household labor income and public and private transfers less income and payroll taxes. The fixed effects regressions are restricted to respondents aged 20–65 for high school graduates and dropouts and to respondents aged 22–65 for college graduates. The replacement rate is defined as the ratio of the average retirement income for individuals aged 65–85 to the average income in the last working year. The estimation procedure for the error structure follows Carroll and Samwick (1997). The variance decomposition for the retirement period applies to income net of medical expenses, and the estimates are taken from Love (2010).

Table 2
Baseline parameters

Parameter	**Value**
Risk aversion (ρ)	5
Bequest parameter (b)	2.5
Discount factor (β)	0.98
Risk-free return (R^f)	2%
Equity premium ($R^e_t - R^f$)	4%
Retirement age (T^R)	65
Maximum age (T)	100
Uncertain income in retirement?	yes
Bequest motive?	no

This table presents the parameter values used for the baseline version of the model.

Table 3
Comparison of Portfolio Choice Rules of Thumb

	College							
	Linear Age Rule				Linear Wealth Rule			
	Age = 20	40	60	80	Age = 20	40	60	80
θ_0	1.200	0.800	0.425	0.475	0.750	0.600	0.475	0.575
θ_1	1.275	0.675	0.125	0.125	0.725	0.475	0.250	0.400
Welfare loss (pct cons)	0.422	0.373	0.131	0.072	0.405	0.217	0.035	0.010
Welfare loss (dollars)	246.09	253.04	98.16	55.60	236.59	147.34	25.94	8.10
	High School							
θ_0	1.250	0.825	0.575	0.675	0.925	0.675	0.625	0.800
θ_1	1.275	0.625	0.250	0.250	1.050	0.550	0.450	0.800
Welfare loss (pct cons)	0.481	0.423	0.245	0.126	0.340	0.177	0.080	0.020
Welfare loss (dollars)	182.76	180.93	112.33	56.94	129.54	76.01	36.86	9.05

This table displays the optimal parameterizations and welfare implications for linear age and linear wealth rules for portfolio choice. The linear age rule sets the portfolio share of equities equal to $\theta_0 - \theta_1 Age$. The linear wealth rule sets the share equal to $\theta_0 - \theta_1$(financial wealth/(financial wealth + human capital). The results assume that the investor applies the portfolio rules that are optimal for the given ages for the remainder of the life cycle. The welfare losses are expressed as both a percent of *annual* consumption and as an *annual* dollar amount. See text for details.

Table 4
Optimal Portfolio Choice Rules with Updating

	College					
	Linear Age Rule			Linear Wealth Rule		
	Age = 20	40	65	Age = 20	40	65
θ_0	1.500	1.350	0.525	1.200	0.675	0.525
θ_1	1.575	1.700	0.250	3.000	0.750	0.325
Welfare loss (pct cons)	0.117	0.216	0.120	0.055	0.142	0.015
Welfare loss (dollars)	68.63	146.54	91.83	32.14	96.24	11.50
	High School					
θ_0	1.625	1.250	0.600	1.250	0.775	0.675
θ_1	1.975	1.450	0.250	2.925	0.900	0.525
Welfare loss (pct cons)	0.176	0.281	0.232	0.055	0.113	0.045
Welfare loss (dollars)	66.92	120.34	107.51	21.13	48.67	20.72

This table displays the optimal parameterizations and welfare implications for linear age and linear wealth rules for portfolio choice. The linear age rule sets the portfolio share of equities equal to $\theta_0 - \theta_1 Age$. The linear wealth rule sets the share equal to $\theta_0 - \theta_1$(financial wealth/(financial wealth + human capital). The results assume that the investor updates to the optimal rules at ages 20, 40, and 65. The welfare losses are expressed as both a percent of *annual* consumption and as an *annual* dollar amount. See text for details.

Table 5
Comparison of Consumption Rules of Thumb

	College							
	Fraction of Income/Assets Rule				Annuitized Wealth Rule			
	Age = 20	40	65	80	Age = 20	40	65	80
θ_0	0.900	0.860	–	–	0.638	0.710	0.688	0.693
θ_1	0.070	0.070	0.070	0.135	–	–	–	–
Welfare loss (pct cons)	4.945	7.551	7.094	2.574	3.625	1.964	2.825	1.601
Welfare loss if adhere to age-20 rule	4.945	8.259	7.094	9.902	3.625	3.491	3.504	2.447
	High School							
θ_0	0.910	0.930	–	–	0.740	0.792	0.787	0.762
θ_1	0.080	0.080	0.080	0.160	–	–	–	–
Welfare loss (pct cons)	4.059	5.937	3.887	1.060	2.056	0.952	1.157	1.161
Welfare loss if adhere to age-20 rule	4.059	6.044	3.887	6.625	2.056	1.658	1.630	1.277

This table displays the optimal parameterizations and welfare implications for two types of consumption rules. The first (left columns) sets consumption during the working life to the minimum of a fraction of permanent income and cash on hand, $\min(\theta_0 P_t, X_t)$, and it sets consumption in retirement equal to the minimum of cash on hand and permanent income plus a fraction of assets, $\min(P_t + \theta_1 X_t, X_t)$. The second rule (right columns) sets consumption equal to a fraction, θ_0, of an approximate annuitized value of total wealth (see equation (16)). Except for the last rows for college and high school graduates, the results assume that the household applies the rules that are optimal for the given ages for the remainder of the life cycle. The last rows assume that the optimal rule at age 20 is followed for the remainder of life. The welfare losses are expressed as both a percent of *annual* consumption. See text for details.

Table 6
Robust Rules of Thumb for Consumption

College

	Fraction of Income/Assets Rule			Annuitized Wealth Rule		
	Baseline	CRRA	Discount	Baseline	CRRA	Discount
θ_0	0.900	0.890	0.890	0.638	0.638	0.638
θ_1	0.070	0.065	0.065	–	–	–
Welfare loss	4.945	5.401	6.764	3.625	4.079	4.252
High School						
θ_0	0.910	0.910	0.900	0.740	0.735	0.718
θ_1	0.080	0.075	0.075	–	–	–
Welfare loss	4.059	4.467	5.944	2.056	2.422	3.269

This table displays the optimal parameterizations and welfare losses *at age 20* for two types of consumption rules. The rules are the same as those in Table 5, except that here I consider the possibility that risk aversion and the discount factor may be uncertain from the perspective of the rule maker. In the columns, "CRRA" assumes that the level of risk aversion is distributed normally with a mean of 3 and a standard deviation of 0.5. "Discount" assumes that the discount factor is normally distributed with a mean of 0.98 and a standard deviation of 0.02. The welfare losses are expressed as a percent of annual consumption. See text for details.

Table 7
Optimal Consumption Rules with Updating

College

	Fraction of Income/Assets Rule			Annuitized Wealth Rule		
	Age = 20	40	75	Age = 20	40	75
θ_0	0.950	0.870	–	0.595	0.737	0.695
θ_1	0.050	0.050	0.105	–	–	–
Welfare loss (pct cons)	3.426	6.248	4.091	2.382	0.620	1.883
Welfare loss (dollars)	2217.40	4413.01	3061.21	1543.62	457.79	1420.43
High School						
θ_0	0.900	0.930	–	0.685	0.810	0.772
θ_1	0.050	0.060	0.125	–	–	–
Welfare loss (pct cons)	3.019	5.102	2.065	1.098	0.376	1.048
Welfare loss (dollars)	1326.91	2347.32	922.99	493.87	182.07	477.80

This table displays the optimal parameterizations and welfare implications for two types of consumption rules. The first (left columns) sets consumption during the working life to the minimum of a fraction of permanent income and cash on hand, $\min(\theta_0 P_t, X_t)$, and it sets consumption in retirement equal to the minimum of cash on hand and permanent income plus a fraction of assets, $\min(P_t + \theta_1 X_t, X_t)$. The second rule (right columns) sets consumption equal to a fraction, θ_0, of an approximate annuitized value of total wealth (see equation (16)). The results assume that the household updates to the optimal rules at ages 20, 40, and 65. The welfare losses are expressed as both a percent of *annual* consumption and as an *annual* dollar amount. See text for details.

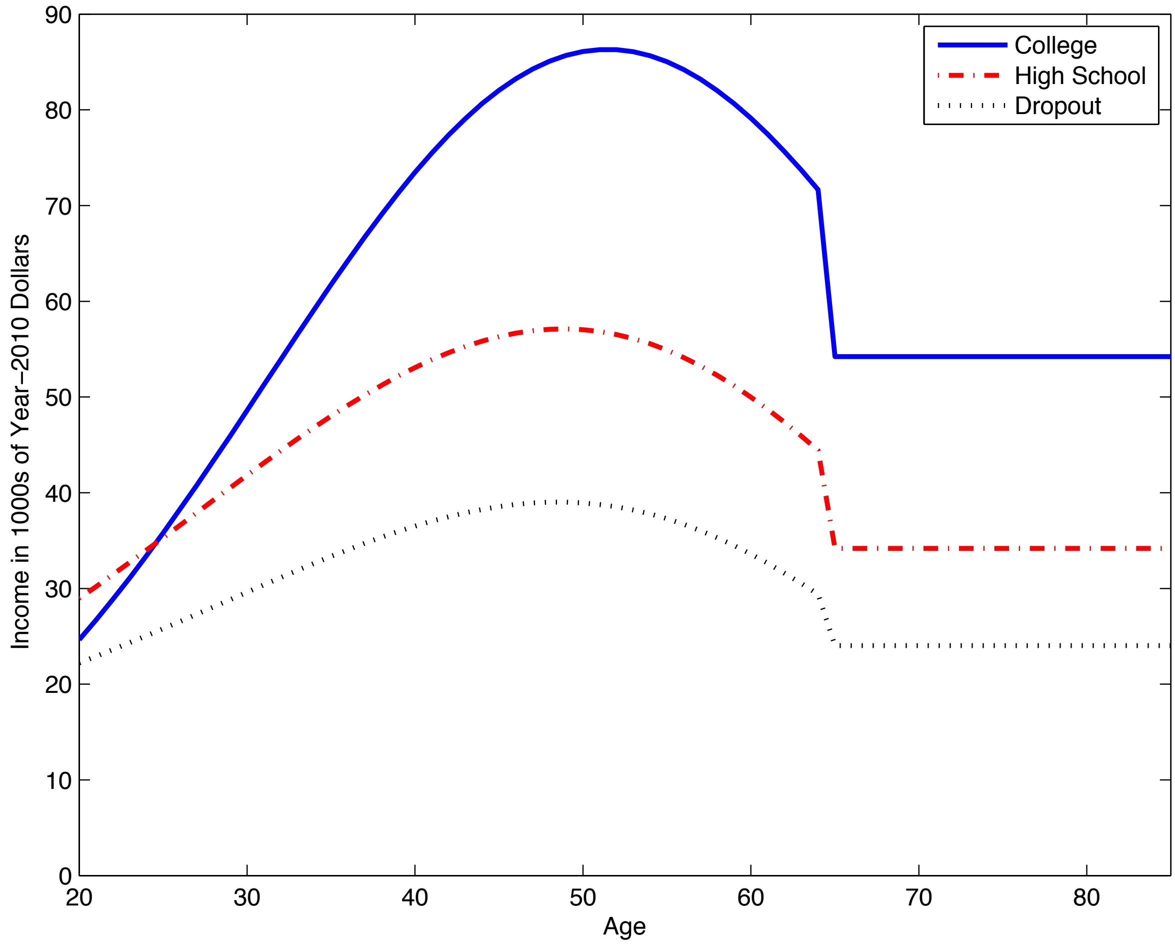

Figure 1
Estimated income profiles: This figure displays the average family income profiles in thousands of year-2010 dollars. The profiles are constructed by fitting a third-degree polynomial through the coefficients on the age dummies in a fixed-effects regression of family income on ages, family composition, and marital status. Retirement income is the average income of retired households between ages 65 and 85. The data come from the 1970–2007 waves of the PSID.

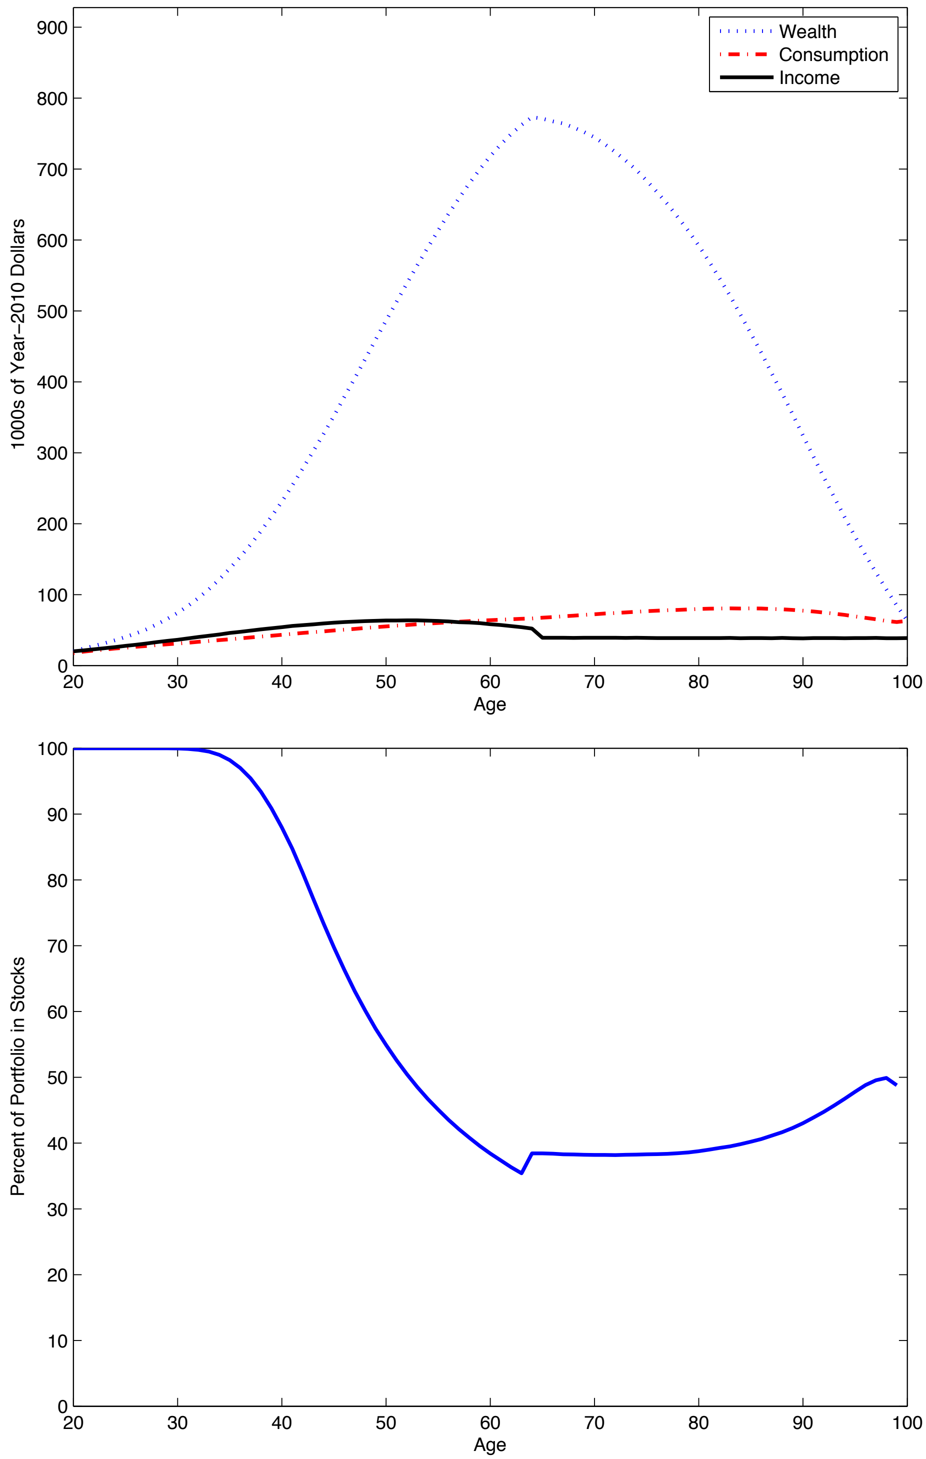

Figure 2
Life-cycle profiles for college graduates: The top panel of the figure shows the average values of wealth, consumption, and income for college graduates for the optimal solution to the baseline model. The bottom panel shows the average allocation in stocks. The baseline model assumes there is not a bequest motive but that there is retirement risk. The income process and model parameters are listed in Tables 1 and 2, respectively.

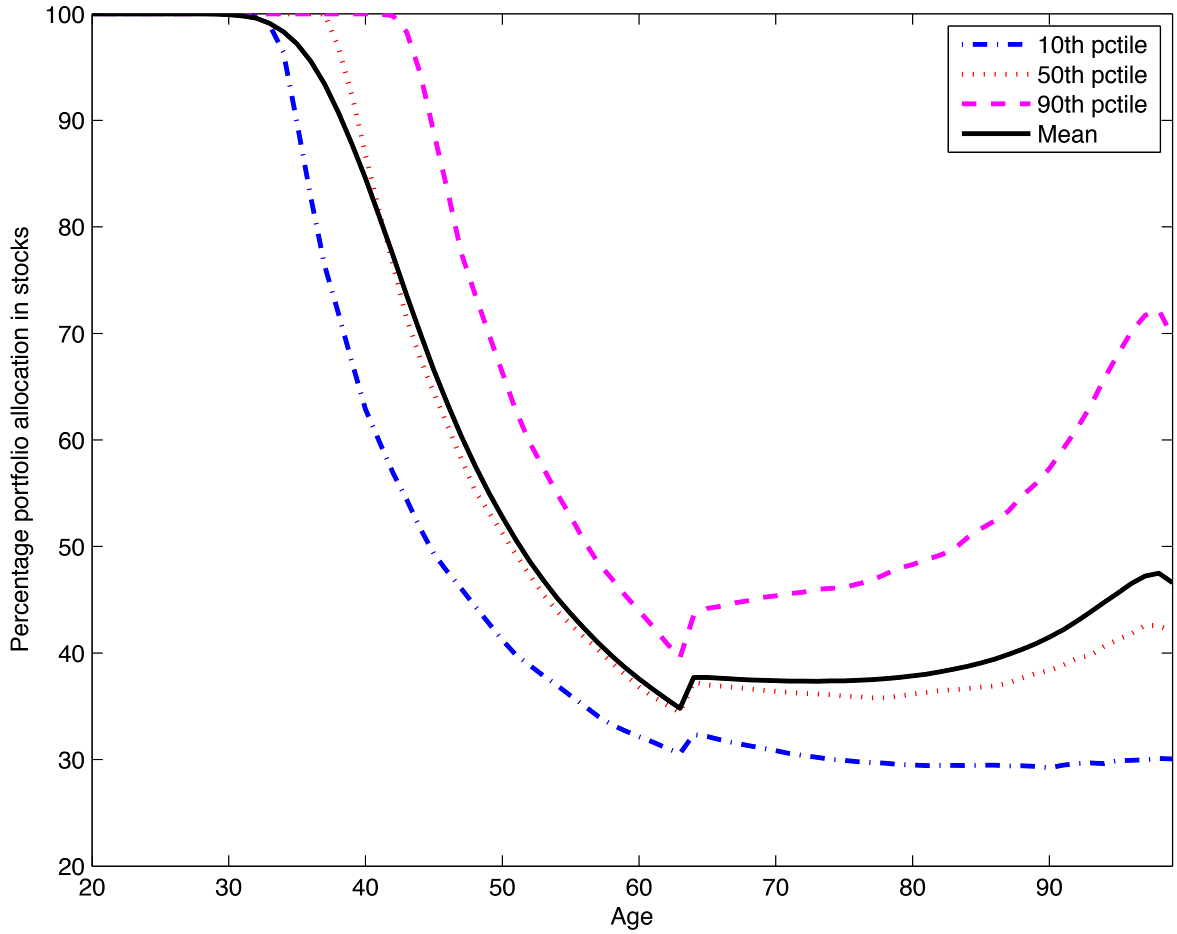

Figure 3
Distribution of optimal share: This figure shows the optimal stock shares by means and 10th, 50th, 90th percentiles for the baseline specification for college graduates.

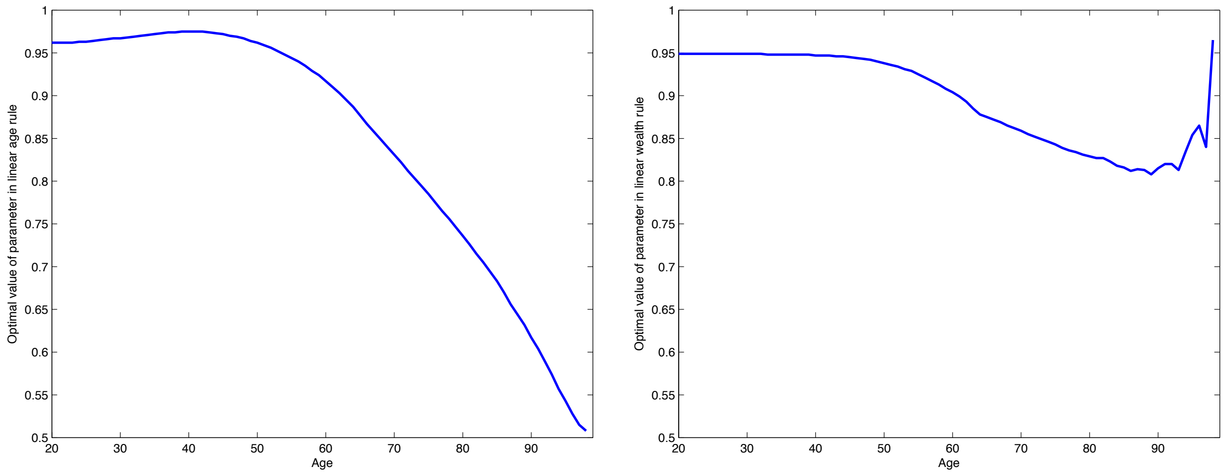

Figure 4
Parameter values for simple age vs. wealth rule: The left panel of the figure shows the optimal value of the slope parameter in the $100 - \theta age$ rule for college graduates in the baseline model; the right panel shows the optimal value of the parameter in the $100 - \theta\mu$ rule. The optimal values assume that the rule is maintained from each age forward.

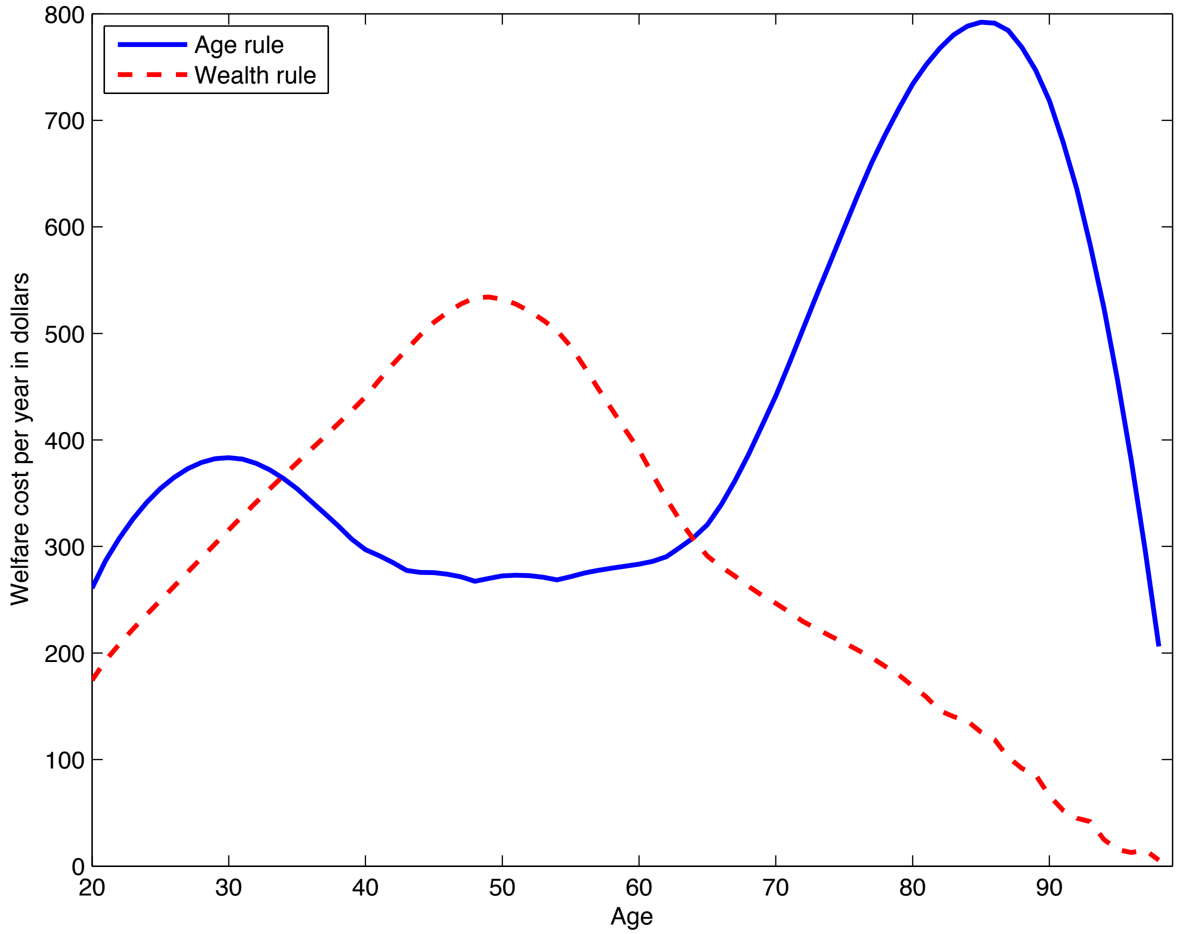

Figure 5
Welfare comparison for simple linear allocation rule: This figure compares the welfare costs of using the optimal linear age rule to the welfare costs of using the optimal linear wealth rule in the baseline specification for college graduates. The optimal rules for a 20-year-old are maintained for the entire life cycle. For each age, the welfare costs represent the *annual* dollar amount of additional consumption that would make the individual indifferent between using the optimal rule and maintaining the rule of thumb for the rest of life.

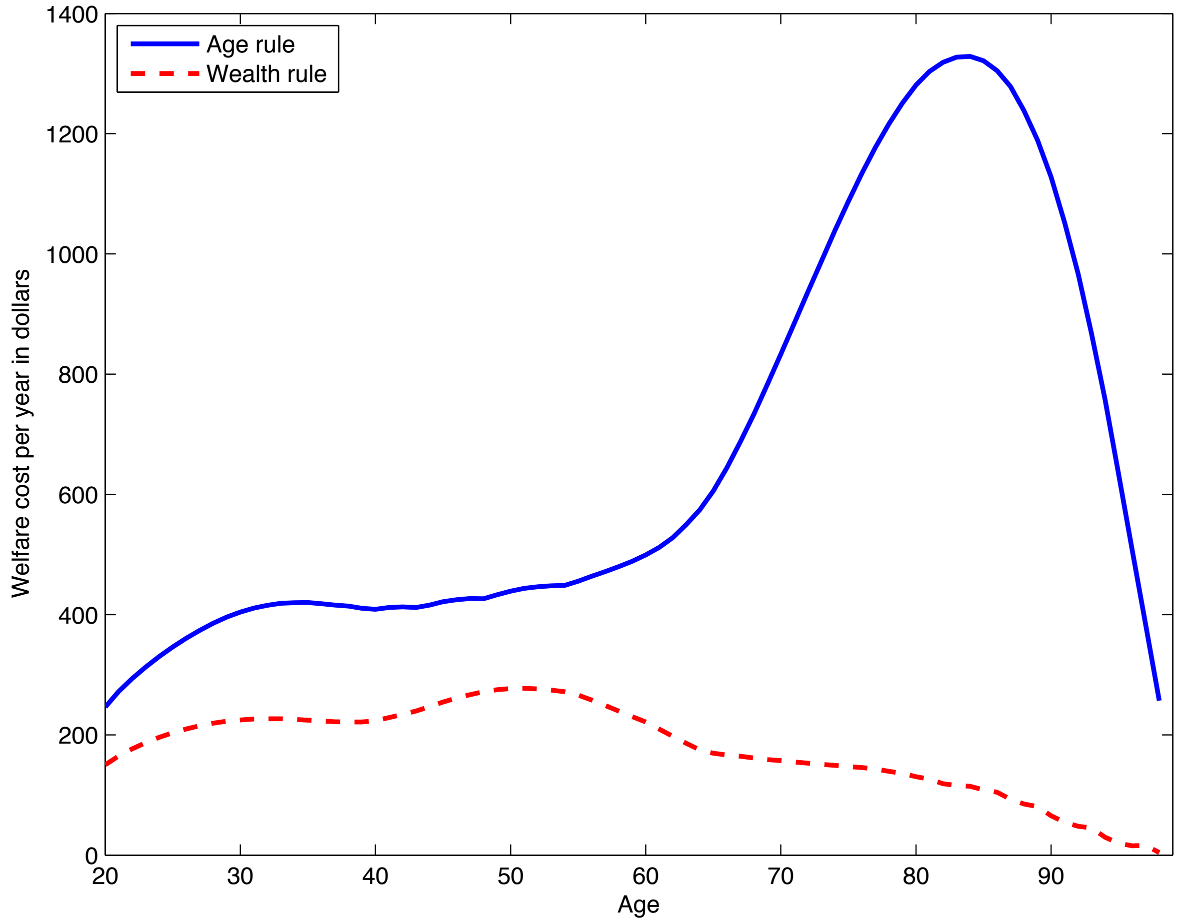

Figure 6
Welfare comparison: This figure compares the welfare costs of using the optimal linear age rule to the welfare costs of using the optimal linear wealth rule in the baseline specification for college graduates. The optimal rules for a 20-year-old are maintained for the entire life cycle. For each age, the welfare costs represent the *annual* dollar amount of additional consumption that would make the individual indifferent between using the optimal rule and maintaining the rule of thumb for the rest of life.