ISSN 1503-299X

WORKING PAPER SERIES

No. 2/2021

Soft Habits*

Knut Anton Mork
Department of Economics
Norwegian University of Science and Technology

Vegard Skonseng Bjerketvedt
Department of Industrial Economics and
Technology Management
Norwegian University of Science and Technology

Department of Economics

Norwegian University of Science and Technology
N-7491 Trondheim, Norway
http://www.ntnu.edu/econ/working-papers

*We are indebted to Snorre Lindset for encouraging us to pursue this topic and to Alfonso Irrarazabal for help with numerical solution methods. Espen Henriksen has offered useful comments. The research has benefitted from financial assistance from Finansmarkedsfondet, the financial markets research fund under The Research Council of Norway, Project no. 294398.

Soft Habits

Knut Anton Mork* and Vegard Skonseng Bjerketvedt**

June 2021

*Adjunct Professor, Norwegian University of Science and Technology (NTNU), Department of Economics, knut.anton.mork@ntnu.no

**PhD student, Norwegian University of Science and Technology (NTNU), Department of Industrial Economics and Technology Management, vsb@ntnu.no

JEL classification: C63, E21, G11

Key words: Habit formation
Withdrawal smoothing
Risk taking
Long-term fund preservation

We are indebted to Snorre Lindset for encouraging us to pursue this topic and to Alfonso Irrarazabal for help with numerical solution methods. Espen Henriksen has offered useful comments. The research has benefitted from financial assistance from Finansmarkedsfondet, the financial markets research fund under The Research Council of Norway, Project no. 294398.

ABSTRACT

Models of habit formation in consumption typically specify utility over the excess of consumption above some habit level. This specification is unsatisfactory in settings where agents occasionally have to tolerate consumption below the habit level. More importantly, they often imply infeasible solutions with realistically low riskless rates. We propose an alternative specification, where the curvature of the utility function rises steeply for consumption below the habit level, but without utility falling abruptly to minus infinity. We explore analytically the key features of the implied behavior and present representative numerical solutions of the model in continuous time. We then simulate investor-consumer behavior with these preferences and compare this behavior to simpler rules of thumb. We find that soft habits, like hard habits, imply procyclical risk taking. Soft habits also allow some smoothing, especially in the downward direction. However, its most distinguishing feature takes the form of deliberate efforts to build sufficient capital to limit the probability of consumption having to fall below habits. Because the priority given to the buildup of wealth, the question of smoothing remains mostly moot in practice. The simpler rules of thumb tend to smooth more and save less that the base case and thus lead to insufficient buildup of capital over time. Of the simpler rules, the relatively best result is found for behavior as if preferences were CRRA with risk aversion somewhere between the soft-habit risk aversion for consumption above and below the habit level.

1. Introduction

Habit formation models have been used extensively in the literature as explanation partly of the so-called excess smoothness of consumption in macroeconomics [(Campbell & Deaton, 1989), (Galí, 1990), and others], and partly of the equity premium puzzle in finance [Abel (1990) Constantinides (1990), Campbell & Cochrane (1999), and Cochrane (2017)]. Whereas these contributions seek to explain the behavior of macroeconomic and macrofinancial time series, this paper focuses on the decision making by the owners and managers of individual funds, primarily endowment funds and sovereign wealth funds[1]. Although we believe our model can be fitted to empirical data, our current objective is to provide a consistent framework for normative use by the owners, managers, and advisors of such funds.

Lindset and Mork (2019) used habit formation to explain the tendency for backward smoothing of withdrawals from endowment[2] and sovereign wealth funds[3] and argued that, for such policies to be consistent, they would have to be paired with cautious attitudes towards financial risk. However, their results also reveal some serious shortcomings of the conventional habit model. The first is the specification of the habit as a hard lower bound on acceptable consumption. Consumption below this level is plainly infeasible in that model. That would make sense in a model applied to a true subsistence minimum which, however, is not typically the case in the context of endowment funds or sovereign wealth funds.

---

[1] Our interest in the subject was inspired by the issues facing the Norwegian government in managing the Norwegian Government Pension Fund Global, the Norwegian sovereign wealth fund.
[2] For example, the so-called MIT-Tobin rule, which prescribes withdrawals as a weighted average of last year's withdrawal and the expected fund return, usually with a weight of 0.8 for last year's withdrawal, see http://web.mit.edu/fnl/volume/205/alexander_herring.html, accessed on Dec 1, 2020. See also Tobin (1974) .
[3] For the Norwegian Government Pension Fund Global, a Fiscal Rule limits withdrawals to the expected annual real return, but with an explicit caveat that large changes in the withdrawals "shall be smoothed over several years." See https://www.regjeringen.no/en/topics/the-economy/economic-policy/economic-policy/id418083/, accessed on Dec 1, 2020.

Campbell and Cochrane avoid this problem by modeling the excess of consumption over habit in log form, which forces it to always stay positive, and by forcing habits to shrink fast whenever recessions strike. This paper proposes an alternative formulation where the slope of marginal utility becomes discontinuously steeper if consumption falls below the habit level, but without jumping to infinity. Thus, we replace a discontinuity in marginal utility with a discontinuity in its derivative. We refer to this feature as a "soft habit" as opposed to the "hard habit" in the existing literature.

The other shortcoming of hard-habit models, which we consider more serious, concerns feasibility. As shown by Constantinides (op. cit.), maintaining a hard habit requires a sufficiently large part of the agent's wealth set aside for investment in safe assets whose yield can fully fund the habit. If the agent's total wealth falls below this magnitude, maintaining the habit would not be feasible, as noted by Lax (2002). This problem would obviously arise with non-positive riskless returns, although that may not be a major shortcoming considering that constant, non-positive rates cannot support a general equilibrium with an indefinite time horizon. However, the problem can arise quite easily with positive rates as well. An illustrative example can be based on the Norwegian sovereign wealth fund, the Government Pension Fund Global (GPFG). At the time of this writing, the value of this fund's assets is about USD 1.36 trillion. According to the Government's fiscal rule, 3%, or about USD 40 billion, can be withdrawn each year to help defray government spending. As contributions of this order of magnitude have been available for about a decade, one may reasonably argue that, say, 25 out of the 40 billion have worked their way into fiscal habits. Now, if the real riskless rate is 1.25%—a very generous estimate at this point in time—Constantinides' rule requires USD 2 trillion to be set aside for this purpose. With a total fund value of USD 1.36 trillion, that is obviously not feasible.

A model with soft habits sidesteps this problem. Although even a soft-habit agent would want to avoid consumption below the habit level, the cost of that happening is finite and can thus be weighed against the lower expected return of a less risky portfolio. Thus, optimization with soft habits is feasible at any level of wealth and riskless rates of return. Indeed, we believe that this specification comes close to mimicking the actual way of thinking for governments with sovereign wealth funds as well as many university administrators in charge of endowment funds.

Abel (1990) introduced a distinction between external and internal habits. With internal habits, the decision maker is fully aware of how his or her habits are formed and develop over time, for example, as a slowly moving average of past consumption. Decisions about current consumption then take into account the effects of this decision on future habits under the assumption that habits will continue to be updated the same way. For example, a person might be reluctant to make expensive purchases for fear that (s)he will get trapped in luxurious habits that may prove hard to break.

With external habits, as used e.g., by Campbell and Cochrane (op.cit.), the habit is determined outside the model. Habits influence decisions because they help shape current preferences; but the influence of current decisions of future habits is ignored. In models of general equilibrium, the habit level is typically specified as determined by the past spending by other agents in the economy, which Abel dubs “catching up with the Joneses.” In our context of endowment funds and sovereign wealth funds, we prefer to think of external habits as established by the pattern of past spending decisions, but that current decision makers ignore the effect that their decisions have on that pattern. Put more plainly, we do not believe that decision makers of this kind give much thought to the dangers of getting trapped in expensive habits.

For this reason, and in this sense, we treat external habits as our main case. We have furthermore found that a model with internal soft habits is very hard to solve, even numerically, as explained in a footnote to Section 4. We thus present solutions for external habits only. That is, we assume that decision makers are aware of habits, but perceive them not only as exogenous, but also as constant over time, so that expectations of future habit changes are ignored as well. We believe this specification is reasonable when habits change slowly over time and consider it similar in spirit to Alan Greenspan's famous definition[4] of price stability as "that state in which expected changes in the general price level do not effectively alter business or household decisions." More generally, it is related to the literature on rational inattention, starting with Sims (2003), which Gabaix (2020) recently developed into a more complete macroeconomic model[5].

In a wider context, our research is related to the research on behavioral finance, dating back at least to Kahneman and Tversky (1976). In particular, our concept of habit can be interpreted as a reference point in the sense of Barberis, Huang, and Santos (2001) in their use of prospect theory as part of their explanation of the equity premium. However, our focus differs from theirs as we are more interested in the actual behavior of institutional fund managers. Also, and partly for the same reason, we maintain the assumption of rational expectations, except for the updating of external habits.

Finally, our analysis makes a contribution to the literature on sovereign wealth funds and institutional endowment funds. The latter have been studied empirically by, for example, Barber and Wang (2013), Brown, et al. (2014), and Dahiya & Yermack (2018).

[4] Minutes of the Federal Open Market Committee, July 2 – 3, 1996, https://www.federalreserve.gov/monetarypolicy/files/FOMC19960703meeting.pdf, accessed January 13, 2021.
[5] Alternatively, because our model ignores other (e.g. labor) income, the perception of constancy can also be interpreted as the implicitly zero drift of other income.

Overviews of sovereign wealth funds have been written by Baldwin (2012) and Alhashel (2015); and their investment strategies and performance have been studied by Bernstein et al. (2013), Paltrinieri & Pichler (2013), Dreassi, et al. (2017) and Johan et al. (2013). A special issue has been taken up by van der Bremer et al. (2016), who discuss the joint decision of exhaustible-resource extraction and portfolio choice for a sovereign wealth fund established to safeguard the revenues for that extraction for future generations. We bypass that issue and study instead the joint decision of portfolio management and revenue spending for an already established fund without consideration of other possible non-tradeable assets.

The next section presents the utility function that we use to represent soft habits. Section 3 puts it into a dynamic context and derives the first-order optimality conditions. Section 4 explores analytically the key properties of the solution. Section 4 present the numerical solution; Section 5 simulates the preformance of the solution over time; and Section 6 compares them to simpler rules of thumb. Section 7 discusses the results and offers some concluding remarks.

2. The Model

We start by considering two CRRA utility functions, differing only in the attitude towards risk:

$$u_j(c) = \frac{c^{1-\gamma_j}}{1-\gamma_j}, j = 1{,}2; \quad \gamma_1 > \gamma_2. \tag{1}$$

We refer to these preference orderings as $CRRA_1$ and $CRRA_2$, respectively. Assuming two assets, one risky (equity) and one riskless (bonds), the optimal equity share is, as shown by Merton (1969),

$$\omega_j^* = \frac{\pi}{\gamma_j \sigma^2}, \tag{2}$$

where π is the equity premium, σ the standard deviation of equity returns; and we have implicitly assumed that the riskless rate is constant and equity prices follow a Brownian motion in continuous time.

Optimal consumption is similarly given by the Keynes-Ramsey rule:

$$c_j^*/A = \eta_j \equiv (1/\gamma_j)\rho + (1 - 1/\gamma_j)\left(r + \frac{1}{2}\omega_j^* \pi\right), j = 1,2, \tag{3}$$

where ρ is the subjective discount rate.

2.1. Spliced utility function

We now use the two preference orderings in (1) as building blocks in the specification of the preferences of our consumer-investor agent. That is, we splice these orderings such that the curvature jumps discontinuously at the point where consumption equals the soft habit level, which we denote h:

$$u(c,h) = \begin{cases} \dfrac{h^{\gamma_1 - \gamma_2} c^{1-\gamma_1}}{1-\gamma_1}, c < h \\ \\ \dfrac{(1-\gamma_1)c^{1-\gamma_2} + (\gamma_1 - \gamma_2)h^{1-\gamma_2}}{(1-\gamma_1)(1-\gamma_2)}, c \geq h. \end{cases} \tag{4}$$

This function is continuous and concave everywhere. It is increasing in consumption and decreasing in habits. Its first derivatives are continuous everywhere and differentiable for all values of c and h except where $c = h$, at which point the first-order derivatives have kinks.

The second-order derivatives are continuous everywhere except for jumps at this point. At $c = h$, the standard measure of relative risk aversion jumps from γ_1 to γ_2.

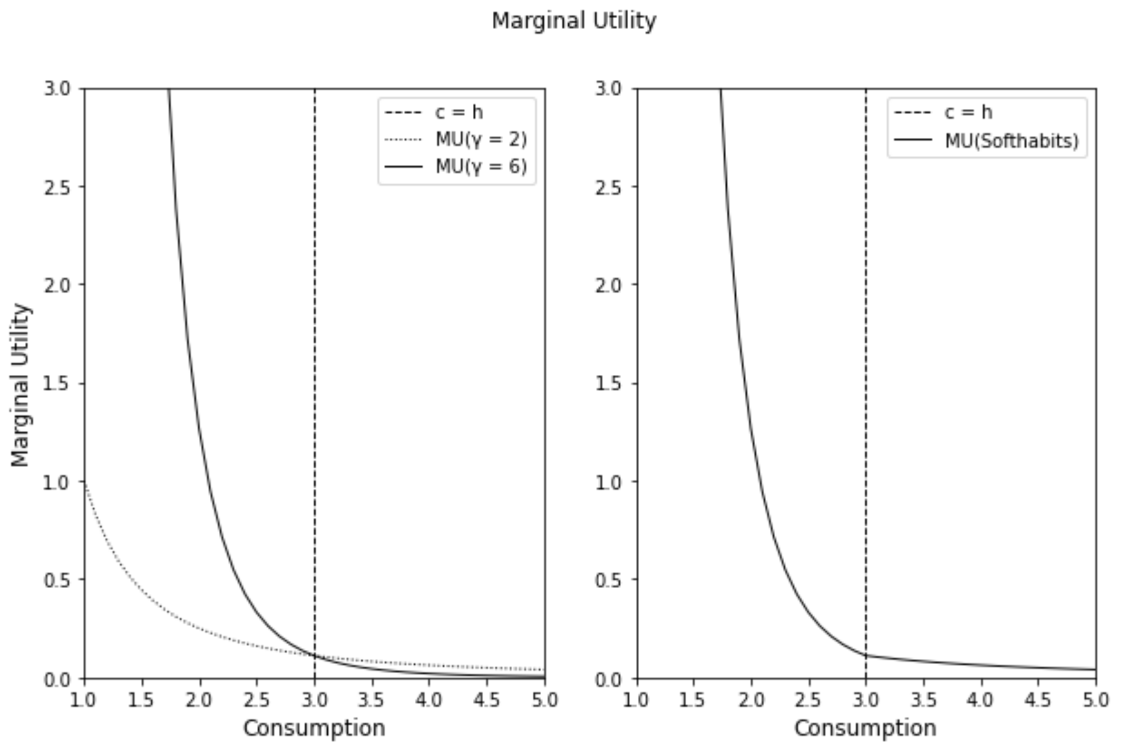

Figure 1: Marginal utility with soft habit

Figure 1 displays the marginal utility of an agent with a soft habit of $h = 3$. The panel on the left shows the marginal-utility graphs for the two CRRA functions separately, whereas the panel on the right shows the corresponding graph for the spliced function. The graph is drawn on the assumption that $\gamma_1 = 6$ and $\gamma_2 = 2$. Thus, this agent is extremely averse to consumption falling below this point. Yet, (s)he can accept such an outcome if the alternatives are worse.

2.2. Dynamic setup

In continuous time, the Bellman equation can be written as:

$$\max_{c,\omega}\{u(c,h) - \rho V(A,h) + \mathbb{E}\, dV(A,h)/dt\} = 0. \tag{5}$$

The maximization is done subject to the capital constraint. Assuming that the value of the risky asset follows a Brownian motion, we can write this constraint as

$$dA = [(r + \omega\pi)A - c]dt + \omega\sigma A dB, \tag{6}$$

where $\{B\}$ is a Wiener process.

Itô's lemma now implies

$$dV(A,h) = \left\{V_A[(r + \omega\pi)A - c] + \frac{1}{2}V_{AA}\omega^2\sigma^2A^2\right\}dt + V_A\omega\sigma A dB,$$

where subscripts denote the partial derivatives of the value function. Thus,

$$\mathbb{E}\left(\frac{dV}{dt}\right) = V_A[(r + \omega\pi)A - c] + \frac{1}{2}V_{AA}\omega^2\sigma^2.$$

Substituting, we can then write the Bellman equation as follows:

$$\max_{c,\omega}\left\{u(c,h) - \rho V(A,h) + V_A[(r + \omega\pi)A - c] + \frac{1}{2}V_{AA}A^2\omega^2\sigma^2\right\} = 0.$$

After partial differentiation of this function with respect to consumption and equating the partial derivative to zero, it is now easily seen that the following first-order conditions for consumption and the equity share are,

$$u_c(c^*,h) = V_A(A,h) \tag{7}$$

and

$$\omega^* = -\left(\frac{V_A}{V_{AA}A}\right)\left(\frac{\pi}{\sigma^2}\right), \tag{8}$$

respectively.

Given these choices, we can write the Bellman equation on compact form:

$$u(c^*, h) - \rho V(A, h) + V_A\left[\left(r + \frac{1}{2}\omega^*\pi\right)A - c^*\right] = 0. \tag{9}$$

Equation (9) is obviously an ordinary differential equation in V as a function of A for given h. However, it is of the second order, as can be seen from the presence of the second-order derivative V_{AA} in (8). It is non-stochastic because the stochastic elements have been taken care of by the combination of Itô's lemma and the laws of expectation for Wiener processes. On the other hand, it is highly non-linear. We thus need to rely on numerical methods for its solution. Yet, some important features of the solution can be obtained analytically, which we do next.

3. Key Properties

In this section, we start by showing that the value function is homogeneous in wealth and habit, which gives the policy functions for consumption and the equity share a nice homotheticity property. Next, we show that the policy functions are continuous in wealth, that both functions are non-decreasing monotonic, and that the policy function for consumption has a kink at the point where $c^* = h$. Finally, we derive the limits of both functions as $A \to \infty$ and $h \to \infty$. Taken together, these properties imply a special shape for the consumption policy function, which implies a willingness to dip into savings when $c^* < h$

even as the agent then has a special incentive to save in order to build sufficient capital to allow $c^* \geq h$ in future periods.

3.1. Homotheticity

The policy functions are homothetic in the sense that the optimal consumption-wealth ratio and the optimal equity share are both invariant to a simultaneous multiplication of habit and wealth by the same positive factor. To see this, note first from (4) that the utility function is homogeneous of the degree $(1-\gamma_2)$:

$$u(\mu c, \mu h) = \mu^{1-\gamma_2} u(c,h), \mu > 0, \forall\ c, h.$$

Thus, marginal utility is homogeneous of the degree $-\gamma_2$. From (7) it then follows that $V_A(A,h)$ must be homogeneous of the same degree in A and h. This makes the policy function for consumption homogeneous of degree one in wealth and habit, so that the corresponding function for the consumption-wealth ratio must be homogeneous of degree zero. Furthermore, the homogeneity of $V_A(A,h)$ implies homogeneity of the same degree for the product $V_{AA}(A,h)A$. Thus, the ratio

$$\frac{V_A(A,h)}{V_{AA}(A,h)A}$$

must be homogeneous of degree zero, making the policy function for the equity share homogeneous of degree zero in wealth and habit.

3.2. Continuity and monotonicity

For the consumption policy function, continuity as well as monotonicity follows directly from (7). For the same reason, V_A is continuous as well.

For the equity share, we start by writing the first-order condition in continuous time as

$$0 = \lim_{\Delta t \searrow 0} \mathbb{E}_t V_{A,t+\Delta t}\big(R_{e,t+\Delta t} - R_f\big) = \lim_{\Delta t \searrow 0} \mathbb{E}_t u_{c,t+\Delta t}\big(R_{e,t+\Delta t} - R_f\big),$$

where the second inequality follows from (7). The special form of the utility function (4) furthermore means that, if A_t rises by a small amount, the formula for $u_{c,t+\Delta t}$ would shift from $h^{(\gamma_1-\gamma_2)/\gamma_1} c^{-\gamma_1}$ to the smaller $c^{-\gamma_2}$ for some of the realizations of $R_{e,t+\Delta t}$. However, because the distribution of the rate of return is continuous, the effect of this shift is merely to move a small part of the probability mass of the product $u_{c,t+\Delta t}\big(R_{e,t+\Delta t} - R_f\big)$ to the left, thus reducing its expectation by a small amount as well. For the above equation to hold, the equity share would then need to be raised by a similarly small amount. Replacing the term "small" by "infinitesimal" in this argument establishes the continuity of ω^* as a non-decreasing function of A. From (8) and the fact that V_A is continuous, it then follows that the second-order derivative V_{AA} must be continuous as well.

3.3. Kink in the consumption function

Although the policy function for consumption is continuous everywhere, its derivative makes a jump at the point where $c^* = h$. To see this, differentiate first both sides of equation (7) with respect to A and solve for the derivative of optimal consumption with respect to wealth. On elasticity form, it becomes

$$\frac{\partial \ln c^*}{\partial \ln A} = \frac{V_{AA} A}{u_{cc} c^*}.$$

Although, as just found, V_{AA} is continuous, u_{cc} makes a jump at the point where consumption equals habit. Let A_h denote the wealth level at which this choice is optimal. We then see that the consumption policy function must have a kink at this wealth level.

To explore this issue further, note from (8) and the corresponding condition (2) with CRRA$_j$ preferences, $j = 1,2$, that

$$\frac{V_{AA}A}{u_{cc}c} = \frac{\omega_j}{\omega^*}.$$

From the monotonicity of ω^*, we thus have

$$\frac{\partial \ln[c^*(A,h)/A]}{\partial \ln A} = \begin{cases} \dfrac{\omega_1 - \omega^*}{\omega^*} < 0, A < A_h \\ \text{undefined}, A = A_h \\ \dfrac{\omega_2 - \omega^*}{\omega^*} > 0, A > A_h \end{cases} \tag{10}$$

3.4. Limiting properties

For a given h, $A \to \infty$ means that the probability of optimal consumption falling below the habit level is extremely unlikely. In the limit, then, the agent will act as if her or his preferences were CRRA$_2$, so that $\omega^* \to \omega_2$ and $c^*/A \to \eta_2$. The monotonicity of ω^* and the positive sign of (10) for $c^* > h$ furthermore warrant that both limits will be approached from below.

If, in contrast, $h \to \infty$ for given A, the agent will behave as if preferences were CRRA$_1$. Monotonicity implies that ω^* then will approach ω_1 from above. However, the negative sign of (10) for $c^* < h$ implies that c^*/A then will approach η_1 from below. We use these results to approximate the solution when $A \to 0$ for given h. Thus, for given (but finite) h, the optimal equity share will start just above ω_1 for low A and rise monotonically to ω_2 as A

increases. The asymptotic properties indicate that the slope of this function will be rather flat on both ends, meaning that it will resemble an *S* curve.

The policy function for consumption will be more complicated. Although monotonic, it will start just below the optimal level for $CRRA_1$ for low A and rise more slowly than it until $A = A_h$. There, the consumption function will have a kink and to the right of this point rise faster than the function for $CRRA_2$. Similarly, consumption as a share of wealth will start just below η_1, decline further until $A = A_h$, and then rise rather steeply until asymptotically levelling out at η_2.

4. Numerical solution

Our solution method transforms the second-order ODE into a two-equation first-order ODE system. We used the solution software for such problems in DifferentialEquations.jl, a package developed by Rackauckas and Nie (2017) using the Julia programming language[6]. Figure 3 offers a clearer picture of the solution of this model[7]. The parameter values are

[6] Because the value function is known to be Lipschitz continuous a unique solution is known to exist. This property allowed us to use a binary search algorithm to search for the initial conditions and final conditions that also satisfy the properties derived in the text. The computations were performed on an HP EliteBook 840 G5 with a Intel Core i7-8550U CPU 1.80Ghz 4 Core processor, using the Rosenbrock23() algorithm with a real and absolute tolerance level of 1e-12 and a maximum iterations of 10e6, the solution to the base case was obtained within 750 seconds.

[7] We also tried to find a numerical solution to the model with internal soft habits, but without success. We believe the problem lies in the effects that the internal habits have on the curvature of the value function. By construction, our instantaneous utility is more highly convex when consumption is low. Normally, one would expect the value function to share this property. However, if low wealth places optimal consumption far below habits, internal habits will be expected to adjust fast downward, which would have the opposite effect on the convexity of the value function. By the same token, the expectation of falling habits could justify higher consumption at low wealth levels than with external habits. The same would be true, with the opposite sign, for high incomes and consumption. This complication thus invalidates the end-point conditions derived in Section 3.4., on which our solution method relies. Forcing habits to adjust very slowly does not mitigate the problem because the adjustment process for habits would then extend so much further into the future that the overall effect on the value function would be qualitatively the same. We suspect that a local solution near the habit level nevertheless can be found with a projection method. However, because we consider the internal-habit case less interesting, we leave this task to future research.

Variable	Symbol	Value
Equity premium	π	4.8 pp
Equity return std. dev.	σ	17.89%
Riskless return (real)	r	1%
Subjective discount rate	ρ	2%
RRA for $c^* < h$	γ_1	6
RRA for $c^* \geq h$	γ_2	2
Implied expected portfolio return, $CRRA_1$	$\bar{r}_1 = r + \omega_1 \pi$	2.2%
Implied expected portfolio return, $CRRA_2$	$\bar{r}_2 = r + \omega_2 \pi$	4.6%
Implied c^*/A, $CRRA_1$	η_1	1,67%
Implied c^*/A, $CRRA_2$	η_2	2.4%
Implied equity share, $CRRA_1$	ω_1	25%
Implied equity share, $CRRA_2$	ω_2	75%

Table 1: Parameter values used in solution (annual rates)

listed in Table 1. The two spliced preference orderings $CRRA_1$ and $CRRA_2$ have been calibrated so as to serve as reasonable representations of the preferences of a highly and less highly risk-averse investor, respectively. In particular, the risk-aversion parameters $\gamma_j, j = 1{,}2$ and the equity volatility parameter σ have been chosen to make the optimal equity shares 75% and 25% for $CRRA_1$ and $CRRA_2$, respectively. The riskless rate and the equity premium have been calibrated so as to imply withdrawal rates $\eta_j, j = 1{,}2$ of 1.67% and 2.4%, respectively. As is typical for CRRA preferences, both withdrawal rates are lower than the expected rates of portfolio return, which are 2.2% and 4.6%, respectively. The solutions presented in Figure 2 are for a habit level of $h = 2$. The homotheticity property makes the solutions for other habit levels analogous.

The two upper panels display the solutions for the equity share (on the left) and the consumption-wealth ratio (on the right) for wealth varying between 0 and 1,000. The middle panel shows the same two graphs, except that the scale on the wealth axis has been limited to the more realistic 50 to 200. The equity share is seen to follow the continuous S-curve implied by the analytical results above, starting at the 25% share for the highly risk-averse preference ordering $CRRA_1$ for very low wealth levels and approaching the 75% of $CRRA_2$ for very high ones.

The upper and middle panels on the right show the solution for the consumption-wealth ratio. As indicated above, this curve starts just below the corresponding ratio η_1 for

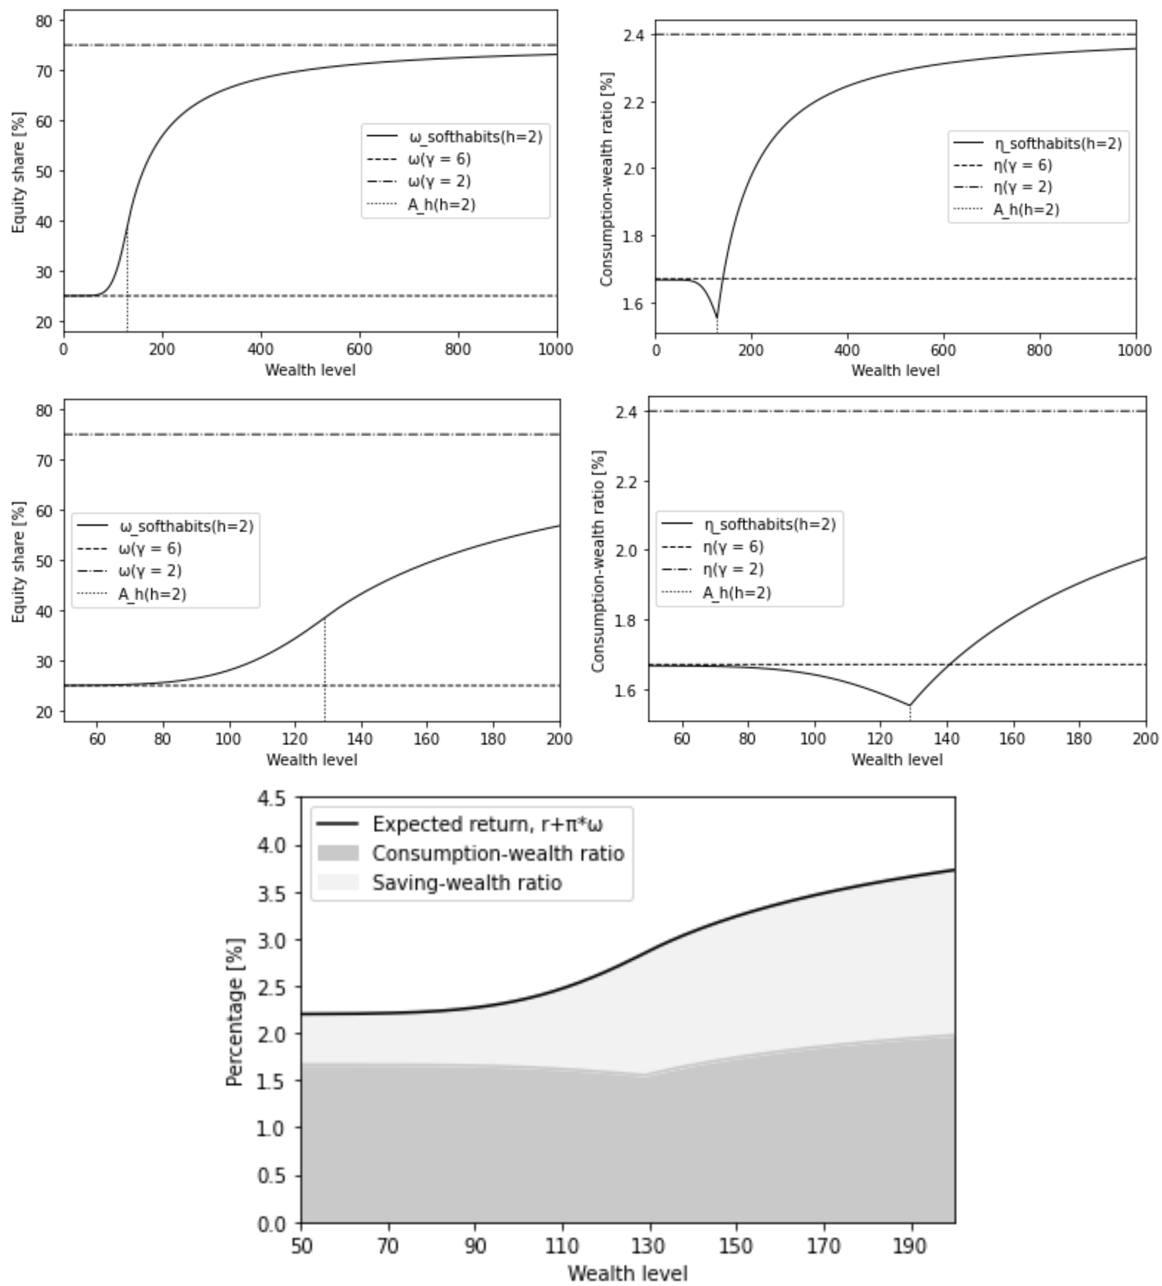

Figure 2: Optimal equity share (top and middle left panels) and consumption-wealth ratio (top and middle right panels) for the habit level $h = 2$. Bottom panel: Saving and consumption-wealth ratios. For illustrative purposes, saving is defined as the difference between the expected portfolio return and the consumption-wealth ratio.

the highly risk-averse preference ordering CRRA_1. However, in contrast to the equity share, this graph actually curves down from there until it reaches the kink at A_h, where the slope discontinuously turns positive. As wealth grows, the graph continues to rise and eventually reaches the low-risk-averse spending ratio η_2, albeit only for very high wealth levels

The bottom panel takes a closer look at the behavior of consumption around the kink in the curve. At the kink, optimal consumption just equals the habit level even though, at the same wealth level, higher consumption would have been chosen if preferences were given by even the highly risk-averse CRRA_1. The reason is that our agent saves more so as to stave off the misfortune of having to accept $c < h$ in the future. The saving is furthermore helped by the fact that the soft-habit agent takes more risk than the one with CRRA_1 preferences, so that the expected portfolio return is higher. The lower-panel graph illustrates this point by including a separate graph for this return, defined as $r + \omega^*\pi$, which naturally rises with ω^*. Note that this graph starts rising significantly at wealth levels below A_h.

Should wealth drop below A_h, the agent will have to accept consumption in the subjectively painful region below the habit. However, (s)he will seek to soften the pain by dipping into savings, thus smoothing the downward movement, as the consumption-wealth ratio is raised above the level at A_h. However, if wealth instead rises above A_h, there is no upward smoothing. Instead, consumption is allowed to rise faster than wealth. Even so, the saving rate, defined as the difference between the expected rate of return and the consumption-wealth ratio, also rises as the agent takes on more risk when the probability of wealth falling below A_h is lower. Saving is an important channel for this agent to reduce the risk of consumption falling below habits.

5. Soft-Habit Behavior over Time

We now use the numerically solved policy functions to study the time-series behavior of a consumer-investor with soft-habit preferences. For this purpose, we discretize the continuous-time dynamic movements with 12 time steps per year, corresponding to months. Although we solved the model under the assumption that the agent ignores habit changes, we now assume that the external habits actually adjust to changing spending patterns as a weighted average of past consumption with geometrically declining weights, which implies the following linear updating function:

$$\frac{dh}{dt} = \theta(c - h). \tag{11}$$

We assume an updating velocity of $\theta = 0.48$ per year, implying $\theta\Delta t = 0.48/12 = 0.04$ per monthly step. This value lies in the upper range of the ones considered by Constantinides (1990) in his model of hard habits in continuous time; but we believe a much slower speed of adjustment would be less realistic.

The simulations use 10,000 duplications and cover a period of 30 years (360 months), with initial wealth of $100 and the initial habit level equal to actual consumption. The top left panel of Figure 3 shows, as a fan chart, the stochastic performance of consumption as a share of wealth, c^*/A. The accompanying line diagram shows, as an illustrative example, the actual performance of an arbitrarily chosen scenario.

Consistent with the pattern in Figure 2, the initial consumption-wealth ratio is lower than even the 1.67% implied by the very risk averse preference ordering CRRA_1. The resulting saving serves to build wealth high enough to help avoid subsequent consumption levels below habit. As more wealth is built this way, the tendency is for the consumption-wealth ratio to gradually increase and approach a kind of steady state. However, even after

30 years, the consumption-wealth ratio has, in 95% of the duplications, risen to much less than the 2.4% implied by the $CRRA_2$ preferences because, as the habit level catches up, the incentive to save persists even though wealth has increased.

This dynamic process is far from monotonic, as illustrated by the fairly large width of the fan chart and especially the rapid movements in the illustrative scenario. These

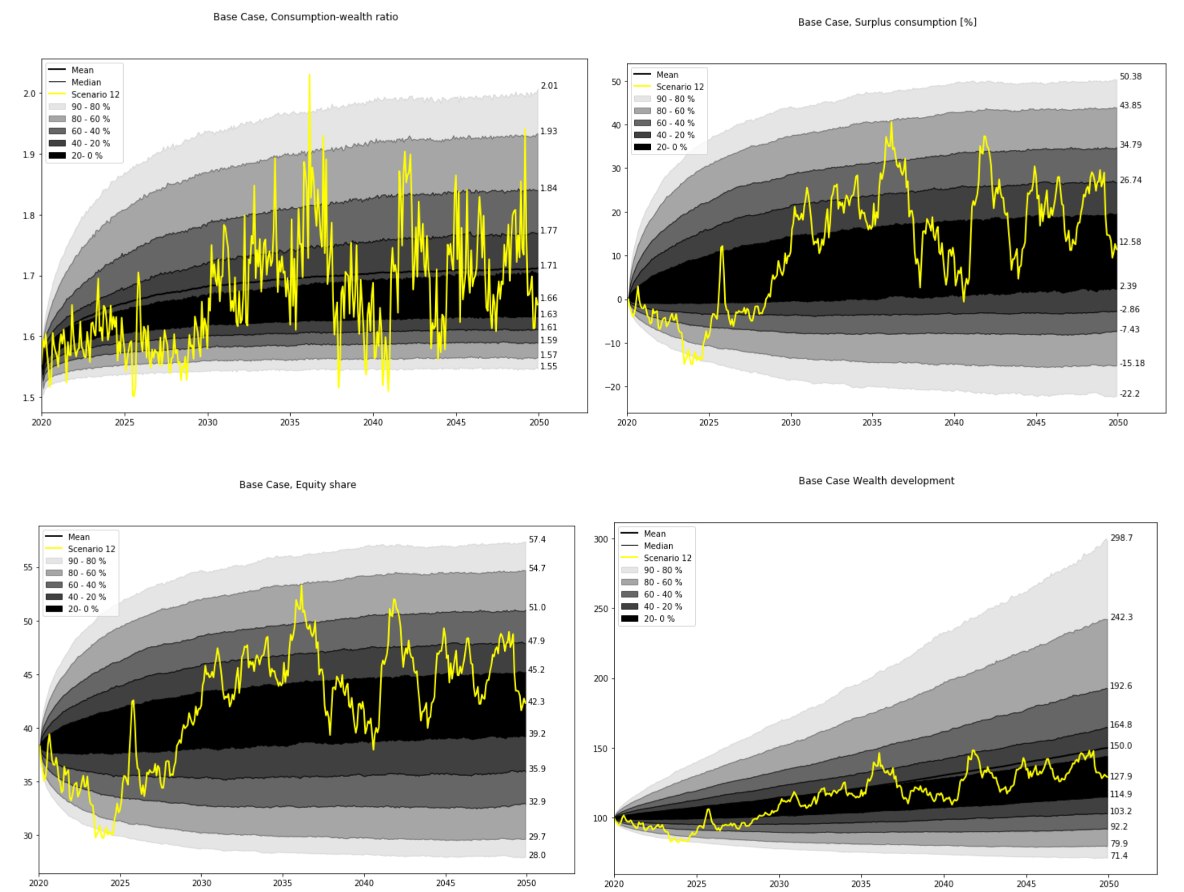

Figure 3: Simulation results, base case. Panels: Top left: Consumption-wealth ratio (annual rates). Top right: Surplus consumption ratio. Bottom left: Equity share. Bottom right: Total wealth.

Legend: Fan charts show the stochastic performance of the variable. The upper and lower 5% of the respective distributions are ignored for clarity of illustration. The light-colored solid lines show the respective paths of one scenario, arbitrarily chosen, for the purpose of illustration.

movements reflect the tug-of-war between saving to stay relatively safely away from the habit level in good times, drawing down wealth to minimize consumption below habit in bad times, and seeing habits catch up with actual consumption once some new wealth has been built up. This process is illustrated further in the top right panel of Figure 3, which displays the ratio that Campbell and Cochrane (op. cit.) call the surplus consumption ratio, $(c-h)/c$. The graph shows this variable starting at zero (by design) but rising gradually over time as wealth is built up. However, in contrast with Campbell and Cochrane's model, it may be negative. In fact, the fan chart reaches consistently below zero; and in the illustrative scenario it falls down to -13% before rising rapidly as the financial luck turns positive. During the difficult period, consumption is lowered a little more than the decline in wealth. This method to get out from under a situation with consumption falling short of habits is absent from the hard-habit literature.

This habit-driven individual is not only thrifty, but also cautious about risk most of the time. The lower left panel in Figure 3 shows the dynamic development of the equity share. As suggested by the left-hand panel of Figure 2, it starts higher than with $CRRA_1$. Even though wealth is build up over time, it tends to decline just a little because habits catch up as well. In fact, no realized equity shares within the 90% percentile get anywhere close to the $\omega_2 = 75\%$ implied by the $CRRA_2$ preferences. Like for the surplus consumption ratio, the movement is far from monotonic in individual scenarios. In particular, we note that risk taking is procyclical as it is taken down during the early period of financial adversity in the illustrative scenario and then increased once consumption has been raised to well above the habit level.

The combination of high saving and procyclical risk taking pays off over time. After 30 years, mean wealth has grown by more than 50%, as shown in the second-to-last row of Table 2. The dispersion of final wealth is also rather high, as illustrated in the lower right panel of Figure 3. The distribution is highly skewed to the right, so that most of the uncertainty implies a chance of much higher wealth than the expectation. The part of the distribution to the left of the median is highly concentrated, implying a virtual zero probability of wealth falling below half of its initial value. This limitation of the downside risk is the result of the careful strategy that not only seeks to avoid consumption below the habit level, but also exploits opportunities for improving upside risks. In particular, when wealth is high enough to allow consumption well above the habit level, the strategy calls for higher risk taking, which significantly improves the prospect of building additional wealth over time.

6. Comparison with Simpler Rules

Because the soft-habit strategy may seem complicated it is natural to compare it with rules of thumb that may achieve approximately the same objectives. Rules of thumb are easier and thus cheaper to administer. In particular, they are better suited to be delegated. We consider the following alternatives

- $\overline{\mathbf{CRRA}}$: Behavior as if preferences were CRRA with risk parameter $\bar{\gamma},\ \gamma_1 > \bar{\gamma} > \gamma_2$, calibrated such that the constant equity share of $\bar{\omega} = \pi/(\bar{\gamma}\sigma^2)$ equals 42%, the average simulated base-case equity share in the final year of the simulations. The consumption-wealth ratio is constant as

$$\bar{\eta} = (1/\bar{\gamma})\rho + (1 - 1/\bar{\gamma})\left(r + \frac{1}{2}\bar{\omega}\pi\right).$$

- **Smoothed** $\overline{\mathbf{CRRA}}$: Same equity share as $\overline{\mathbf{CRRA}}$, $c_t = \lambda c_{t-1} + (1-\lambda)\bar{\eta}A_t, \lambda = 0.45$.
- $\overline{\mathbf{CRRA}}$ **smoothed downwards**: Same equity share as $\overline{\mathbf{CRRA}}$,

$$c_t = \begin{cases} \lambda c_{t-1} + (1-\lambda)\bar{\eta}A_t \ if \ \bar{\eta}A_t < c_{t-1} \\ \bar{\eta}A_t \ otherwise \end{cases}$$

All four rules were simulated using the same set of draws for the stochastic wealth process.

The simulations showed remarkably similar results for the three alternative rules of thumb, while at the same time differing sharply from the base case. Figure 4 compares the consumption-surplus ratio for all four strategies. With identical initial habits and initial

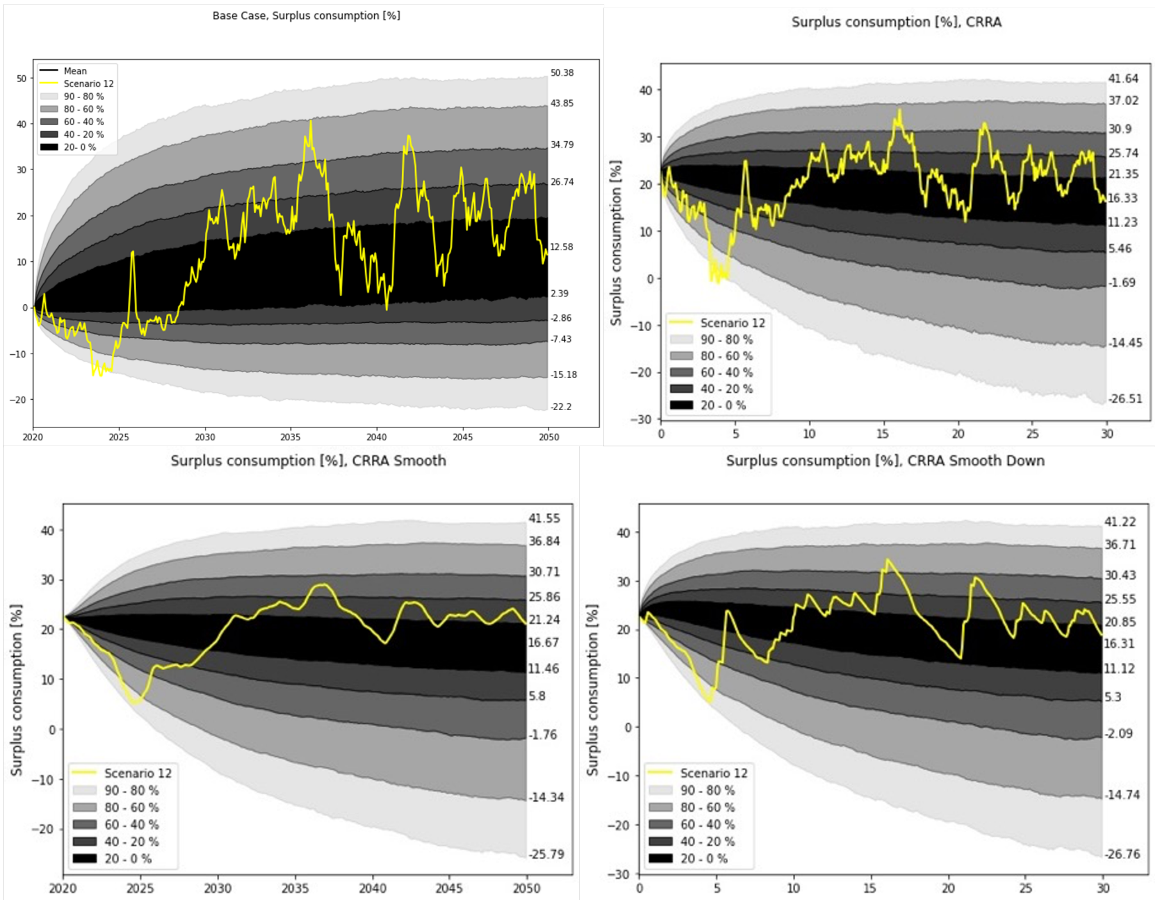

Figure 4: Surplus consumption ratios under alternative rules

wealth as the base case, the alternative rules all start with higher consumption levels—and hence higher consumption surplus ratios—because they ignore the dynamic benefit of saving to build up capital. Over time, however, the tables are slowly turned. The surplus consumption of the alternative rules tends to decline, however slightly, whereas the one in the base case clearly tends to rise.

The dynamic consequences of the differing saving behavior become apparent in Figure 5, which compares the stochastic development of wealth for the four respective rules. In the base case, wealth growth faster on average, reaches greater highs when lucky, and is able to limit the downside better is the draws are consistently on the bad side of luck. The thriftiness of the base case involves early sacrifices but pays off over time by building more wealth.

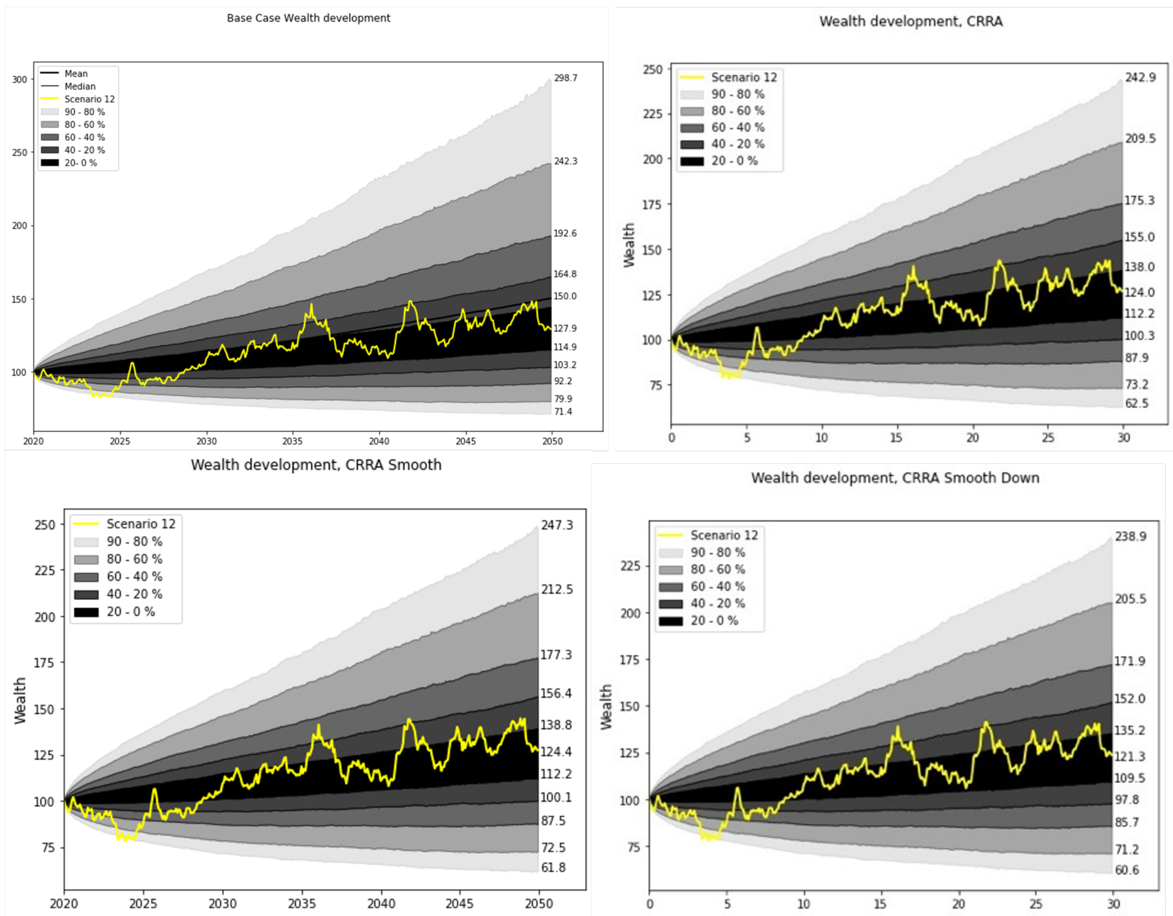

Figure 5: Wealth after 30 years under alternative rules

Table 2 summarizes the results. The top three lines present our main yardstick for comparing the respective alternative rules of thumb with the base case, namely, the discounted expected utility of following the respective rules. With a subjective discount rate of $\rho = 2\%$ per year, the discount factors don't become negligible until 200 years hence. To avoid having to average all the 10,000 realizations for 2,400 months, we instead started by computing computed the simulated discounted average for the 30 years that we simulate. These are the numbers in the first row of the table. The second row shows the the discounted average value function valuated 30 years hence:

$$DEU_i^{30} = \sum_{k=1}^{n} \frac{1}{n} \sum_{t=1}^{360} e^{-(\rho/12)t} u(c_{tki}, h_{tki}),$$

where $i =$ Base case, $\overline{\text{CRRA}}$, smothed $\overline{\text{CRRA}}$, and downward smoothed $\overline{\text{CRRA}}$, t indicates months, k indicates duplications, and $n = 10{,}000$.

The value function assumes behavior according to the base case of soft habits, but starting, 30 years hence, with the wealth level and habit accumulated each of the alternative cases:

$$DVF_i^{30} = \sum_{k=1}^{n} e^{-30\rho} V\big(A_{360,k,i}, h_{360,k,i}\big).$$

The sum then shows the discounted expected utility of following the rule in question for 30 years and then revert to the rules of the base case, but with the wealth and habit level inherited from the rule followed the preceding 30 years. Needless to say, this feature biases the comparison in favor of the alternatives.

On the horizon of the first 30 years, it pays to be frivolous, in the sense that the alternative rules, which ignore the dynamic benefit of thriftiness implied by the soft habit, produce slightly higher expected utility. This benefit is also illustrated in the fourth row of the table, which shows the average surplus consumption ratio over the first 30 years.

However, this advantage comes at the significant dynamic cost of starting the subsequent future with much lower wealth. Thus, on long horizons, meaning more than 30 years, the expected utility loss of following an alternative rule becomes substantial.

Of the three alternative rules, we are somewhat surprised to see that the straight $\overline{\text{CRRA}}$ case comes out best. The explanation again appears to lie in the fact that the main feature of soft-habit behavior is thriftiness rather than smoothing. Limiting the smoothing to downward movements is actually the worst alternative because it represents the most spendthrift behavior.

Smoothing turns out to make the biggest difference for the marginal propensities to consume (MPC). The highest average MPC is found in the base case. When the attention is limited to the spending reaction to wealth declines, the difference from the alternatives is actually greater. We believe the explanation to this finding lies in the particularly steep slope of the policy function for the consumption-wealth ratio for wealth just a little higher than A_h in Figure 2, which is the region where most of the action takes place in the base case.

7. Discussion and Conclusion

The behavioral implications of the soft habits analyzed in this paper differ from those of all the other cases we have considered. We note in particular that soft habits do not simply lead to smoother consumption than, say CRRA preferences. In fact, consumption levels may move up and down quite a bit as this agent focuses on his or her main concern, namely, to avoid consumption below the habit level. An element of downward smoothing is indeed allowed for when consumption falls below the habit level because the agent then will sell some assets to prevent consumption from falling too far. A relatively low equity share adds some further smoothing. However, when wealth is higher than A_h, yet so low as to barely

	Base case	$\overline{\text{CRRA}}$ case	Smoothed $\overline{\text{CRRA}}$	Downward smoothed $\overline{\text{CRRA}}$
Expected discounted sum of utility over 30 years (in utils*)	-1.86	-0.68	-0.66	-0.48
Expected discounted value function after 30 years**	0.19	-6.62	-7.37	-9,23
Sum of the two preceding lines	-1.67	-7.3	-8.03	-9.71
Mean $\left(\frac{c-h}{c}\right)$	9.9 %	16.8 %	16.9 %	17.9 %
Std $\left(\frac{c-c_{-1}}{c_{-1}}\right)$, annual rate	3.07 %	2.17 %	0.39 %	0.95 %
Mean c/A, annual rate	1.68%	2.01 %	2.00 %	2.10 %
Mean MPC (trimmed)***	2.16 %	2.01 %	0.13 %	0.32 %
Std MPC (trimmed)***	0.04	0	0.15	0.20
Mean MPC upward (trimmed)***	1.99 %	2.01 %	0.34 %	0.33 %
Mean MPC downward (trimmed)***	2.34 %	2.01 %	-0.10 %	0.31 %
Density after trim***	100 %	100 %	99.98 %	99.96 %
Wealth after 30 years 5th percentile	71.4	62.5	61.8	60.6
Wealth after 30 years 25th percentile	97.8	94.5	94.2	92.1
Wealth after 30 years 50th percentile	127.9	124.0	124.4	121.3
Wealth after 30 years 75th percentile	177.8	164.5	166.1	161.3
Wealth after 30 years 95th percentile	298.7	242.9	247.3	238.9
Mean wealth after 30 years	150.0	135.1	136.0	132.3
Std (wealth after 30 years)	82.1	57.9	59.8	57.2

Table 2: Simulation results.

**Expected utility computed as the discounted sum of the average simulated realizations over 30 years. Because utility in this model may take on positive as well as negative values, we have not attempedt to normalize the units.*

***Assuming behavior reverting to the base case after 30 years*

Sample moments calculated from the last 10 years of the 30 years simulated.

****MPC distribution trimmed to exclude cases of* $|\Delta c/\Delta A| > 10$*. These extreme cases occur when consumption changes even though wealth does not, e.g. because habits change. These extreme cases are not caused by extreme rates of financial returns.*

permit consumption above habit, the agent will seek to save as much as possible (without lowering consumption below habit) so as to build enough wealth to keep some assets to prevent consumption from falling too far below the habit level. This is made possible in part by raising risk taking so as to enable the agent to harvest a larger risk premium. As a result, risk taking is procyclical, a feature shared with hard-habit models. However, soft habits differ by displaying less smoothing.

We believe our model offers important lessons for endowment funds and sovereign wealth funds. The recommendation to keep withdrawals lower on average than the expected return on the fund follows on the heels of similar recommendations from a wide range of other models but is worth noting. Despite the tendency found by Brown et al. (op.cit.) towards “endowment hoarding” by university administrators, we see a tendency for many fund owners to actually withdraw the expected real return. The Norwegian Government Pension Fund Global is a point in fact. The Norwegian government’s fiscal rule is even more generous by also allowing for smoothing, similar in spirit to the Tobin-MIT rule.

The recommendation to let risk taking vary procyclically is also at odds with common practice as well as a number of models but is shared with the hard-habit literature. However, the recommendation for asymmetric variation in the consumption-wealth ratio is unique to our soft-habit model. While allowing for some downward smoothing, it moves the emphasis to the building of sufficient wealth to avoid having to accept spending below the habit. This emphasis requires active management of spending as well as risk taking as low spending is combined with higher risk taking so as to facilitate saving when wealth is just high enough to allow consumption above the habit level.

Following these recommendations obviously requires highly active management of spending as well as portfolio composition. Delegating a rule this complicated may not be simple. The simpler rules that we look at produce deceivingly nice results even if the investment horizon is kept as long as 30 years. However, ignoring the longer horizons leaves future periods or generations with significantly lower wealth.

Our model may possibly be useful empirically to help explain the equity premium puzzle and to refine the macroeconomic modeling of consumption. However, our principal aim in this paper has been to derive useful insights for owners of endowment funds and sovereign wealth funds. Operational models for such purposes need to also take into account other revenues and their behavior over time, including their correlation, if any, with the movements of financial returns. That will be the agenda for our next paper.

References:

Abel, A. B., 1990. Asset Prices under HabitFormation and Catching up with the Joneses. *American Economic Review,* 80(2, Papers and Proceedings), pp. 38-42.

Alhashel, B., 2015. Sovereign Wealth funds: A literature review. *Journal of Economics and Business,* Volume 78, pp. 1-13.

Baldwin, C., 2012. *Sovereign Wealth funds: The New Intersection of Money and Politics.* Oxford: Oxford University Press.

Barber, B. M. & Wang, G., 2013. Do (some) university endowments earn alpha?. *Financial Analysts Journal,* 69(5), pp. 26-44.

Barberis, N., Huang, M. & Santos, T., 2001. Prospect Theory and Asset Prices. *Quarterly Journal of Economics,* CXVI(1), pp. 1-53.

Bernstein, S., Lerner, J. & Schoar, A., 2013. The Investment Strategies of Sovereign Wealth Funds. *Journal of Economic Perspectives,* Volume 27, pp. 219-237.

Brown, J. R., Dimmock, S. G., Kang, J.-K. & Weisbrenner, S. J., 2014. How University Endowments Respond to Financial Market Shocks: Evidence and Implications. *American Economic Review,* Volume 104, pp. 931-962.

Campbell, J. & Deaton, A., 1989. Why is Consumption So Smooth?. *Review of Economic Studies,* Volume 56, pp. 357-374.

Campbell, J. Y. & Cochrane, J. H., 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy,* 107(2), pp. 205-251.

Cochrane, J. H., 2017. Macro-Finance. *Review of Finance,* pp. 945-985.

Constantinides, G. M., 1990. Habit Formation: A Resolution of the Equity Premium Puzzle. *Journal of Political Economy,* 98(3), pp. 519-543.

Dahiya, S. & Yermack, D., 2018. *Investment returns and distribution policies of non-profit endowment funds.* s.l.:Working Paper 25323, National Bureau of Economic Research.

Dreassi, A., Miani, S. & Paltrinieri, A., 2017. Sovereign Pension and Social Security Reserve Funds: A Portfolio analysis. *Globsl Financial Journal,* Volume 34, pp. 43-53.

Gabaix, X., 2020. A Behavioral New Keynesian Model. *American Economic Review,* 110(8), pp. 2271-2327.

Galí, J., 1990. Finite horizons, life-cycle savings, and time-series evidence on consumption. *Journal of Monetarh Economics,* Volume 26, pp. 433-452.

Johan, S., Knill, A. & Mauck, N., 2013. Determinants of Sovereign Wealth Fund Investment in Private Equity vs. Public Equity. *Journal of International Business Studies,* Volume 44, pp. 155-172.

Kahneman, D. & Tversky, A., 1976. Prospect Theory: An Analysis of Decision under Risk. *Econometrica,* 47(2), pp. 263-291.

Lax, Y., 2002. *Habit Formation and Lifetime Portfolio Selection.* Available at: https://ssrn.com/abstract=325043 [Accessed 13 January 2021].

Lindset, S. & Mork, K. A., 2019. Risk Taking and Fiscal Smoothing with Sovereign Wealth Funds in Advanced Economies. *International Journal of Financial Studies,* 7(4), pp. 1-24.

Merton, R. C., 1969. Lifetime portfolio selection under uncertainty: The continuous time case. *Review of Economics and Statistics,* Volume 51, pp. 247-257.

Paltrinieri, A. & Pichler, F., 2013. Asset Management Issues in Sovereign Wealth Funds: An Empirical Analysis. In: J. Falzon, ed. *Bank Performance, Risk and Securitization.* London: Palgrave Macmillan.

Rackauckas, C. & Nie, Q., 2017. DifferentialEquations.jl-2017 – A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia. *The Journal of Open Research Software,* Volume 5.

Sims, C. A., 2003. Implications of rational inattention. *Journal of Monetary Theory,* 50(3), pp. 665-690.

Tobin, J., 1974. What is permanent endowment income?. *American Economic Review,* 64(2), pp. 427-432.

van der Bremer, T., van der Ploeg, F. & Willis, S., 2016. The Elephant in the Ground: Managing Oil and Sovereign Wealth. *European Economic Review,* Volume 82, pp. 113-131.