ISSN 1503-299X

WORKING PAPER SERIES

No. 1/2025

Dynamic Spending and Portfolio Decisions with an Internal Soft Habit

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

Frida Nymark Engelstad
Department of Economics
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

Dynamic Spending and Portfolio Decisions with an Internal Soft Habit

Knut Anton Mork and Frida Nymark Engelstad

Norwegian University of Science and Technology

Department of Economics

March 2025

JEL classufications: C61 E21 G11

Initial research for this paper was carried out by Frida Nymark Engelstad and Ingeborg Eldevik Rusaas for their joint MSc thesis with Mork as supervisor. We thank Fabian Andsem Harang and Silvia Lavagnini for initial discussions and Haakon Andreas Trønnes for helping us get started with Julia. The responsibility for errors and omissions is our own. Comments are welcome. Email addresses: frida.n.engelstad@ntnu.no and knut.anton.mork@ntnu.no.

Abstract

We solve the Merton problem for an agent with a soft internal habit whose utility is defined over the consumption-habit ratio, and where the curvature of the utility function jumps discretely to a higher level when consumption falls below the habit. The habit evolves over time as a weighted average of past consumption, so that current consumption choices have consequences for future habits. We solve the model numerically in continuous time. Optimal behavior is characterized by extensive consumption smoothing and risk taking that is high on average but varies significantly with wealth. Withdrawal rates tend to fall below portfolio rates of return, so that wealth and consumption tend to grow exponentially over time, like in AK macro models. A lack of self-awareness, whereby the agent underestimates the rate at which habits catch up with consumption, leads to significant, though hardly huge, loss of utility. Simpler rules, such as constant withdrawal rates and equity shares over time, lead to utility losses in similar orders of magnitude. However, when withdrawal rates are held constant, risk taking must be significantly lower to be optimal in a constrained sense. We believe the results should be interesting for the management of and spending from funds like the Norwegian Government Fund Global, which motivated our study. Although our results may serve as a defense of its current regulations with a fixed equity share and (roughly) fixed withdrawal rates, our results suggest that risk taking is much too high.

Contents

1 Introduction

A person acting out of habit seeks to maintain past patterns of behavior whenever conditions change. Although adjustments must be made eventually, they are carried out sluggishly. A habit-driven individual seeks to avoid abrupt changes. In economic analysis, habit formation is typically introduced in preference orderings in such a way that the utility of a given amount of consumption depends on what the agent is used to.

Habit formation has been invoked as a modeling device to account for such observed phenomena as the excessive smoothness of consumption in macroeconomics (Campbell and Deaton 1989; Galí 1990) and Mehra and Prescott 1985's equity premium puzzle in finance (Constantinides 1990; Campbell and Cochrane 1999; Shrikhande 1997; Choi et al. 2022). Munck 2008 and Angoshtari et al. 2024 explore some wider implications of habit formation in more complex settings. Related formulations of downward restrictions can be found in Shin et al. 2007; Arun 2012; Jeon and Park 2021; Dybvig 1995; and Dybvig 1999.

Our interest in habit formation arises from the challenges involved in the management of the Norwegian Pension Fund Global (GPFG, popularly referred to as the oil fund) and the regular use of withdrawals from this fund as partial funding for the Norwegian government's annual budget. Because of large, regular deposits from the government's oil and gas revenues, this fund has grown remarkably rapidly from zero in 1996 to USD 1.8 trillion at the time of this writing. Its official mandate stipulates that the real principal value be preserved for future generations. It allows for fairly high risk taking with a 70 percent equity share. Annual withdrawals are permitted up to amounts corresponding to the fund's expected real financial return, further stipulated as 3 percent of assets under management. However, the mandate also allows the withdrawals to be smoothed relative to the vagaries of the global financial markets, as well as some flexibility in response to fiscal needs, such as during recessions.

Satisfying all these requirements at the same time is a far from trivial task. Smoothing in this context means backward smoothing in the sense of avoiding large budget cuts in case of financial adversity. This is different from the forward smoothing implied by the standard Euler-equation models, first studied by Hall 1978, which assume agents to react immediately to news about current or futures income or wealth. Dybvig and Qin 2021 demonstrate how linear smoothing may unintentionally deplete endowment funds. A simulation study by Mork et al. 2022 suggest high probabilities of depleting the GPFG in finite time under its current mandate.

This paper explores ways that mandates such as the one for the GPFG, may be improved by specifying the fund owners' preferences as influenced by habits formed from the consumption experience of the recent past. The habit feature ensures some degree of smoothing in the sense that past spending of fund revenues sets a precedent for subsequent spending. Although smoothing in general risks depleting the fund, that will not happen in the optimal solution.

Our results thus suggest two alternative strategies for an agent that wants to smooth consumption in the face of financial surprises. The optimal strategy allows considerable smoothing and high risk taking on average, but requires very careful management of risk taking so as to avoid depletion. The alternative strategy is to rely on a main rule of a constant consumption-wealth ratio, but with abrupt consumption changes dampened by a linear smoothing rule. In that case, however, the equity share, also constant, must be kept quite low because part of the smoothing then must be done by limiting the volatility of the portfolio financial returns. Expressed in popular terms, strong smoothing and high risk taking don't mix unless managed extremely meticulously.

We specify the habit as a soft rather than a hard constraint in the sense that we assume that the fund owner can endure withdrawal spending at levels below the established habit, albeit with a substantial loss of utility. Technically, we do this by specifying utility as a function of consumption relative to habit rather than the arithmetic difference between the two. Angoshtari et al. 2024 refer to the two types of specification as multiplicative and linear habit formation, respectively. In the literature, the linear formulation, which we refer to as a hard habit, appears to be most common, including in the influential papers by Constantinides 1990 and Campbell and Cochrane 1999. The multiplicative, or soft habit, was introduced by Abel 1990 but not used as much in subsequent studies.

We feel the hard specification is excessively restrictive in our context exactly because it does not allow consumption to fall below the habit. In order for this constraint to be always satisfied, the hard-habit agent needs to set aside a risk-free part of the portfolio that is large enough to fund consumption at the habit level indefinitely. A rebalancing of the portfolio of the Norwegian GPFG according to this principle at the present time would require much larger changes than we believe anyone would find reasonable. At the riskless rates prevailing before the Covid-19 pandemic it might even be infeasible as the riskless set-aside would have to be larger than the fund itself.

A recent paper by some of us (Mork et al. 2023) specified the owners' utility as the ratio of withdrawal spending to a social norm that grows over time at an exogenously given rate. The specification could be interpreted as a requirement that the withdrawals remain stable relative to a growing GDP. Our present specification is equally soft but updates the norm endogenously as a weighted average of past withdrawal spending.

The literature on habit formation makes a further distinction between internal and external habits. An internal habit is based on the agent's own past consumption. The agent is furthermore conscious of this connection, so that, when making decisions about current spending, the agent takes into account the effect that current consumption will have on future habits. An external habit, in contrast, reflects the agent's desire to maintain his or her spending relatively to everybody else's, in what Abel 1990 refers to as "catching up with the Joneses." Naturally, this agent will not perceive any effect of his or her own actions on everybody else's habits. However, in a symmetric model, their habits will nevertheless move the same way in equilibrium. In our context, external habits do not seem very relevant as the Norwegian government has no fund-owning peers to catch up with.

Because the issues facing the GPFG resemble those of the endowment funds for many pension funds and endowment funds for universities and charitable foundations, our paper makes a contribution to the literature on such funds, exemplified by Hansmann 1990, Novy-Marx and Rauh 2011, Campbell 2012, Barber and Wang 2013, Brown et al. 2014, Gilbert and Hrdlicka 2015, Dahiya and Yermack 2018, Campbell and Sigalov 2022, and Campbell et al. 2024. It is also related to the more general literature on sovereign wealth funds, including Baldwin 2012, Alhashel 2015, Arouri et al. 2018), and Braunstein 2022.

We find that the habit feature implies a significant smoothing of consumption relative to wealth changes. This effect is much stronger than with a soft social norm as in Mork et al. 2023. Yet, because utility is optimized on an infinite horizon, this smoothing does not risk depletion of the fund. The mechanism tying these features together is the risk taking, which on average is quite high under a fully conscious soft habit. The reason may seem subtle, namely, that the effect on future habits blunts the loss of utility caused by a financial loss. When the agent is aware that he or she will get used to the lower consumption over time, the utility loss is short-lived, which reduces the discounted sum of utility losses. This context makes risk taking less risky in a utility sense. As a result, this agent can afford to spend more even at quite low wealth levels. On the other hand, risk taking varies considerably with wealth. In particular, strong smoothing of consumption in response to adverse financial developments also requires risk to be taken down as an attempt to avoid further losses at a time when habit-driven consumption eats into the already reduced wealth.

Because we have no strong views on the rate at which habits are or should be updated, we derive solutions for varying degrees of this mechanism. Unsurprisingly, the features of the results just described show up more strongly the faster the updating is assumed to happen.

Although we believe Norwegian policy makers are quite aware of the precedent past spending of fund revenues sets for current spending, we are less convinced that they are equally conscious of the precedent that current spending sets for future spending or of the speed with which their own perception of spending needs is updated. In fact, this level of self-awareness may seem unrealistic in a number of contexts. We nevertheless believe it to be appropriate to include it in a normative analysis like ours. However, we also include a set of analyses of cases where the agent is less self-aware in this sense.

By analyzing cases where agents are not fully aware of the consequences of current actions for future habits, we take a step in the direction of behavioral finance. In fact, habit formation itself can be considered a step in this direction, as argued by Angoshtari et al. 2024.

On the other hand, in the version of our model where the agent is conscious of the future consequences of current actions can be considered a case of rational addiction, as analyzed by Becker and Murphy 1988. Indeed, a number of popular observers have described Norwegian policy makers or even the country's population as a whole as addicted to the country's oil-based wealth.

Our analysis of behavior with less than full self-awareness in regard to habit development can be used to estimate the utility loss that arise from such behavior. We do this by running simulations of behavior based on a belief that the rate of adjustment is lower than it actually is. The utility loss can be expressed in units of initial wealth. In the most extreme case, we find a utility loss corresponding to a 25 percent loss of initial wealth.

Because the optimal variations in withdrawals and equity shares may be hard to communicate effectively, we also examine the implications of following simpler rules, such as constant equity shares and constant consumption-wealth ratios, with and without some ad hoc linear smoothing of consumption. We apply a grid search to determine the optimal equity shares and consumption-wealth ratios thus constrained. The utility losses from following such simple rules are generally similar in order of magnitude to those of being less than fully aware of one's own habit generation.

However, the constrained-optimal values of the consumption-wealth ratio are significantly lower than the corresponding expected rates of financial return, which governs the Fiscal Rule currently governing the Norwegian government's annual withdrawals from the GPFG. And, importantly, the constrained-optimal equity share is significantly lower than the central tendency in the fully optimized case. This is especially true when we add yet another constraint, namely that the withdrawal rate must equal the portfolio rate of return.

The rest of the paper is organized as follows. The next section presents the model and its solution. Section 3 presents and discusses the implied policy functions for consumption/withdrawals and the equity share. Section 4 presents the dynamic behavior of the main variables, based on simulation of the model over a 50-year period, starting from a point where consumption just equals the habit. Section 5 considers the cost of being less than fully self-aware, in the sense of underestimating the habit catch-up rate; and Section 6 presents similar results from following simpler rules, such as a fixed withdrawal rate and a constant equity share. Section 7 discusses the results and concludes.

2 Model

2.1 Soft habit

We consider a decision maker whose preferences are described by the same instantaneous utility function as in Mork et al. 2023:

$$u(c,x) = \begin{cases} \frac{(c/x)^{1-\gamma_1}-1}{1-\gamma_1}, & \text{if } c < x, \\ \frac{(c/x)^{1-\gamma_2}-1}{1-\gamma_2}, & \text{if } c \geq x, \end{cases} \qquad \gamma_1 \geq \gamma_2 \geq 1, \tag{1}$$

where c denotes consumption and x the habit. This utility function is obviously homogeneous of degree zero in its two variables. The parameters γ_1 and γ_2, which capture the curvature of the utility function, are assumed to at least equal unity, with the limiting case of unity as usual represented by the logarithmic function. This stronger-than-usual assumption, as opposed to the more common requirement that γ_1 and γ_2 simply be positive, ensures that consumption and habit are complements. This complementarity aligns with the definition of addiction proposed by Becker and Murphy 1988 [1].

The utility function is spliced, consisting of two segments with different risk-aversion parameters, γ_1 and γ_2, depending on whether consumption is below or above the habit, as illustrated in Figure 1. When $c < x$, the marginal-utility curve is especially steep, so that the experience of consumption falling below this level is particularly painful. Conversely, for $c \geq x$, the relative flatness of the curve makes reducing consumption within this interval less of a sacrifice. Whereas the utility function and its first derivative both are continuous at $c = x$, the second derivative jumps at that point.

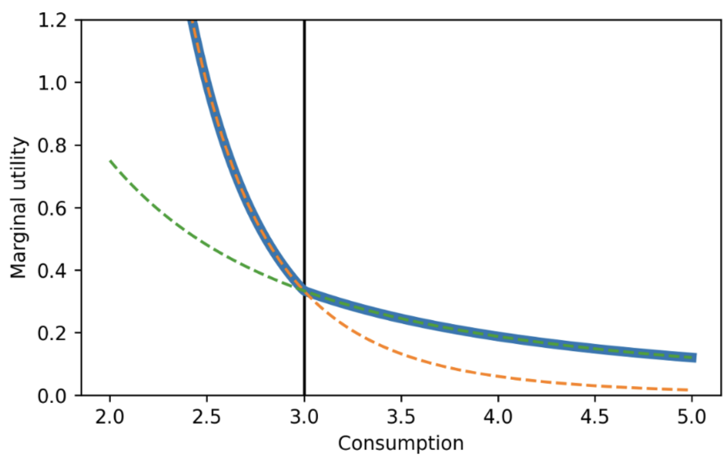

Figure 1: Marginal utility with a soft habit
The red dashed curve shows marginal utility for a CRRA utility funciton with risk aversion parameter $\gamma_1 = 6$ and the green dashed curve the marginal utility for $\gamma_1 = 2$. The solid green curve represents the marginal utility for the spliced utility function.

We refer to this structure as a "soft habit" because it allows consumption to fall below the habit, albeit at a considerable subjective cost. Abel 1990 analyzes the simpler case of a soft habit with constant relative risk aversion, i.e. $\gamma_1 = \gamma_2$. Because the habit is continually updated based on the agent's past consumption, Abel's formulation amounts to more than a constant scaling of the utility function. We nevertheless ad the jump feature because it allows us to study the case of an agent who erroneously believes that his or her habits are static. That agent is aware of the special pain of having consumption fall below the habit threshold despite the erroneous belief that this threshold is unaffected by current consumption choices.

[1] For the limiting case where $\gamma_1 = \gamma_2$, this requirement also ensures the existence of a feasible interior solution for the special case where the habit does not change over time. Then, the model is reduced to the standard CRRA case, for which optimal consumption-wealth ratio is known to be, in our notation, $c/w = p/\gamma + (1 - 1/\gamma)(r + \frac{1}{2}(\frac{1}{\gamma})(\frac{\pi}{\sigma}))^2$

2.2 Optimization setup

We specify the habit as a weighted average of past consumption in continuous time,

$$x_t = e^{-\theta t}x_0 + \theta \int_0^t e^{-\theta(t-s)}c_s ds, \theta \geq 0, \tag{2}$$

which implies the updating condition

$$dx_t = \theta(c_t - x_t)dt. \tag{3}$$

We note that the habit is internal in that it is update on the agent's own past consumption. For a large investment fund, which has motivated our work, we find this specification more plausible than the external-habit formation used elsewhere in the literature, for example Campbell and Cochrane 1999. Whereas external habit formulations may make sense when analyzing equilibrium in a market of many investors with similar characteristics, our interest lies in the precedence that actual withdrawal decisions set for future decisions. That seems better approximated by an internal habit.

We assume that the agent has the choice between two assets, one risky, referred to as equity, and one riskless. The agent's wealth, w evolves according to:

$$dw_t = (r + \alpha_t \pi)w_t dt + \sigma \alpha_t w_t dB_t - c_t dt, \tag{4}$$

where B_t is a Brownian motion governing equity returns; r is the risk-less rate of return, π the equity premium, and σ the standard error of equity price fluctuations, all asssumed constant; and α_t is the portfolio's equity share.

The agent's objective is to maximize the expected, subjectively discounted integral of current and future consumption given present wealth and habit, expressed as the following payoff function:

$$P(t, w_t, x_t) = \mathbb{E} \int_t^\infty e^{-\rho(s-t)}u(c_s, x_s)ds,$$

where ρ is the subjective discount rate.

The HJB equation for this optimization problem in continuous time, as derived by e.g. Rogers 2013 takes the following form:

$$
\begin{aligned}
0 = \max_{c_t,\alpha_t} \Big\{ & u(c_t, x_t) + V_t(t, w_t, x_t) + V_w(t, w_t, x_t)\big[(r + \alpha_t \pi) w_t - c_t\big] \\
& + V_x(t, w_t, x_t)\theta(c_t - x_t) + \frac{1}{2}\sigma^2 \alpha_t^2 w_t^2 V_{ww}(t, w_t, x_t)\Big\}.
\end{aligned} \tag{5}
$$

Here, V denotes the value function, which represents the value, in utility units, at time t, of owning the initial wealth w_t and having the habit x_t. The subscripts of V denote its partial derivatives.

2.3 First-order conditions and transformation

Optimization of (5) with respect to consumption, making use of the homogeneity of u, yields the first-order condition

$$
u_c(c_t, x_t) = u'(c_t/x_t)(1/x_t) = V_w(t, w_t, x_t) + \theta V_x(t, w_t, x_t). \tag{6}
$$

This equation makes it clear that the value function must share the homogeneity property of the utility function, in other words, that it is homogeneous of degree zero in w_t and x_t.

Equation 6 can be solved for optimal consumption c_t^* as

$$
c_t^* = x_t u'^{-1}\{[V_w(t, w_t, x_t) + \theta V_x(t, w_t, x_t)]\}. \tag{7}
$$

Because x_t is an average of past consumption, equation (7) defines a fixed-point problem. Proving the existence of its solution turns out to be far from trivial. However, a forthcoming paper by Harang and Irrarazabal n.d. presents a solution to an analogous problem. We thus proceed under the assumption that the condition in equation (7) can be satisfied under our conditions.

The policy function for the equity share has the familiar form

$$
\alpha^* = -\frac{V_w(t, w_t, x_t)}{w_t V_w w(t, w_t, x_t)}, \tag{8}
$$

which is clearly homogeneous of degree zero in w_t and x_t.

The homogeneity of V in w_t and x_t furthermore implies

$$
V_x(t, w_t, x_t) = -\frac{w_t}{x_t} V_w(t, w_t, x_t).
$$

Moreover, from the time homogeneity of the payoff function, $V_t = -\rho V(w_t, x_t)$. So for fixed t, we may write $V(t, w_t, x_t) = V(w_t/x_t)$. With these simplifications, we can write the optimized HJB equation as the following ODE in w_t/x_t:

$$u(c_t^*/x_t)-\rho V(w_t/x_t)+V'(w_t/x_t)\big[(r+\alpha_t^*\pi-\theta(c_t^*/x_t-1))w_t/x_t-c_t^*/x_t\big]+\frac{1}{2}\sigma^2\alpha_t^{*2}(w_t/x_t)^2V''(w_t/x_t)=0. \tag{9}$$

Define now the transformed variables $y_t := c_t/x_t$ and $\omega_t := w_t/x_t$. Then, we can write the optimized HJB equation (9) as

$$u(y_t^*)-\rho V(\omega_t)+V'(\omega_t)\big[(r+\alpha_t^*\pi-\theta(y_t^*-1))\omega_t-y_t^*\big]+\frac{1}{2}\sigma^2\alpha_t^{*2}\omega_t^2V''(\omega_t)=0. \tag{10}$$

With this transformation, we readily see that our original optimization problem is equivalent to maximizing the payoff function $\tilde{P}_t = \mathbb{E}\int_t^\infty e^{-\rho(s-t)}u(y_s)ds$ subject to the transformed wealth constraint

$$d\omega_t = (r+\theta+\alpha_t\pi)\omega_t dt + \sigma\alpha_t\omega_t dB_t - (1+\theta\omega_t)y_t dt. \tag{11}$$

This transformation moves the updating of the habit from the utility function to the capital constraint. The result is an optimization based on a CRRA utility function (albeit with a shifting risk parameter γ) where the agent faces a tax-and-subsidy scheme on consumption at the amount $\theta(y_t-1)\omega_t$. The agent receives a subsidy if $y_t < 1$, but is taxed whenever $y_t \geq 1$. In both cases, the agent faces a marginal consumption tax at the rate $\theta\omega_t$. This tax is highly progressive.

2.4 Solution

Because our prior notion of the rate of habit updating θ is highly uncertain we solve the model for four different values: $\theta = (0, 0.1, 0.3, 0.7)$. The updating condition (3) implies a half life of closing a given gap between consumption and habit of $\ln 2/\theta$. Because $\ln 2 \approx 0.7$, the case of $\theta = 0.7$ represents a case where it takes about a year to half-way close such a gap. With $\theta = 0.3$, it takes between two and a half years, with $\theta = 0.1$, it takes seven, and with $\theta = 0$ it never happens.

The jump in the curvature of the utility function at $y = 1$ may raise questions about the continuity of the value function and its derivatives. However, for the case of $\theta = 0$, an analytical solution is available from the model in Mork et al. 2023 for the case of a zero growth rate for the exogenous social norm in their model. In that solution, the value function has continuous first and second order derivatives. Because the addition to the model of $\theta > 0$

contains no new features that could change this result, we conclude that the value function must have continuous first and second order derivatives in our case also.

Except for the case where $\theta = 0$, we have not been able to solve our model analytically. Because the marginal value of one additional unit of the consumption-investment good differs from its marginal utility, the dual function method recommended by e.g. Rogers 2013 does not work. However, the following results is useful in the search for a numerical solution. From (7) and (8), we find, after some algebra,

$$\frac{d\ln y^*}{d\ln\omega} = \frac{\frac{\pi}{\alpha^*\sigma^2} - \frac{\theta\omega}{1+\theta\omega}}{\gamma_j}, \tag{12}$$

where $j = 1$ if $y^* < 1$ and $j = 2$ otherwise. The expression on the right must obviously be nonnegative for y^* to be optimal, which places an upper limit on the optimal equity share:

$$\alpha^* \leq \left(\frac{1+\theta\omega}{\theta\omega}\right)\frac{\pi}{\sigma^2}. \tag{13}$$

On the other hand, if the equity share is too high, consumption will grow explosively until all wealth is depleted, which would also be not optimal. This observation suggests a saddle-point path for the optimal equity share.

This insight becomes useful for our numerical solution, which is based on the parameter values in Table 1. The habit-formation mechanism makes sure that transformed wealth $\omega = w/x$ will rarely take on very different values from the one ensuring $y = 1$, at which optimal consumption just equals the habit. Thus, it will be sufficient for the numerical solution to cover a relatively narrow interval $(\omega_{min}, \omega_{max})$, where $\omega_{min} = 10$ and $\omega_{max} = 200$, whereas values of ω around 35 correspond to $y = 1$, varying a little depending on the rate of habit catch-up θ.

r	0.025
ρ	0.04
π	0.045
σ	0.19
γ_1	6
γ_2	2
θ	0, 0.1, 0.3, 0.7

Table 1: Parameter values used in numerical solutions

Given guesses for the optimal values y_0 and α_0 for $\omega = \omega_{min}$, we use the routine available in *DifferentialEquations.jl*, a package programmed in Julia, to solve the ODE in (10) forward[2]. Our problem is then reduced to choosing the values for y_0 and α_0 that maximize

[2]Documentation of this method can be found in Rackauckas and Nie 2017. Our problem turns out to be

the value function.

We carry out this search in a double loop, an inner loop for α_0 and an outer loop for y_0. In the inner loop, we use binary search to identify a unique value of α_0 that, given our guess for y_0, produces non-explosive paths for both variables. In the outer loop, we search for the highest value of y_0 for which we are able to find a non-explosive path in the inner loop. This procedure produces a value function whose graph lies above the corresponding graphs for other feasible y_0 guesses.

3 Policy Functions

3.1 Consumption

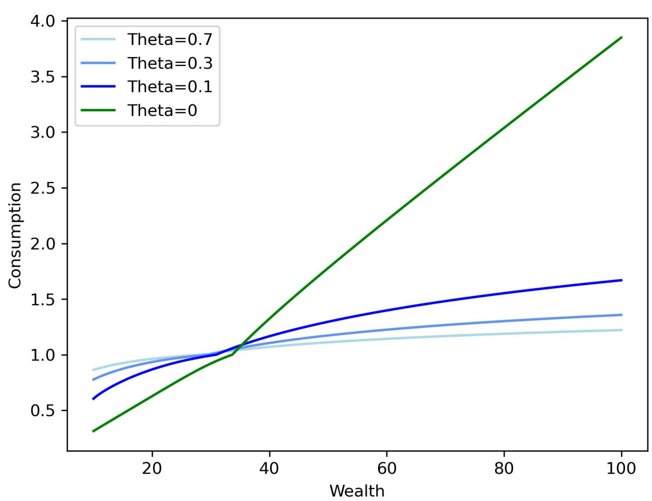

Figure 2: Policy functions for consumption for varying values of the habit updating rate θ.

Figure 2 presents a graph of the policy functions for consumption given habit $x = 1$[3] for $\theta = (0, 0.1, 0.3, 0.7)$. As a reflection of the spliced nature of the utility function (1), each graph consists of two spliced, concave curves. The splicing occurs at the wealth level for which the optimal consumption just equals the habit, which varies a little depending on the value of θ. The slope is somewhat flatter in the segment to the left of this point, which means that the agent seeks to especially smooth the effect on consumption if wealth falls below this level. This is even the case for $\theta = 0$. Recall that the jump in the risk-aversion parameter makes the case of $\theta = 0$ different from the straight CRRA case. However, the second segment of each graph is concave as well, reflecting how habit formation induces consumption smoothing even if consumption stays well above the habit. This kind of backward smoothing, where consumption is allowed

a "stiff ODE", meaning that there is a high degree of numerical instability when attempting to solve the model. We use the algorithm "Rodas5" in Rackauckas and Nie 2020, which is recommended for problems of this kind.

[3]By fixing the habit at $x = 1$, Figure 2 simultaneously graphs c against w and y against ω.

to change slowly in the case of financial adversity, is different from the forward smoothing that is implied by forward-looking behavior in general, as originally pointed out by Hall 1978.

The concavity of both segments is quite pronounced in the cases where $\theta > 0$. To see that, the reader may want to pay attention to the scale of the vertical axis. For all $\theta > 0$, as wealth rises tenfold from 10 to 100, consumption only increases from a range between 0.8 to 1.3, to one between 1.0 and 1.5. The reason is the agent's realization that habits always tend to catch up with any increase in consumption over time, so that the enjoyment of any rise in consumption will be short lived. This realization softens the incentive to raise consumption when wealth rises. Not surprisingly, this effect is stronger the faster habits adjust to actual consumption, i.e. the larger is θ. The contrast with the case of $\theta = 0$ is substantial, even for θ as low as 0.1. Although the model is obviously continuous in θ, the effect of a small increase in θ from zero is apparently very large. Updating of habits matters even if the updating is very slow. Even so, consumption only quadruples if wealth rises by a factor of 10 in the case where habits aren't updated at all ($\theta = 0$).

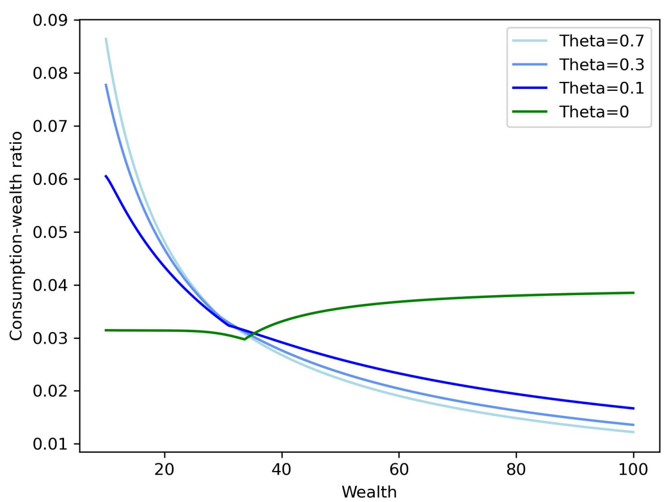

Figure 3: Optimal withdrawal rates (consumption-wealth ratios) as functions of wealth, for varying values of the wealth-updating rate θ.

Figure 3 provides an alternative perspective on the effect of habit updating by displaying the implied consumption-wealth ratios — or withdrawal rates — for the same set of values for θ. Although an agent whose habits don't adapt at all ($\theta = 0$) may allow the withdrawal rate to rise slightly if wealth falls below the level implying consumption equal to habit, this effect is much, much stronger for agents whose habits adapt, even for adaptias slow as with $\theta = 0.1$.

The willingness to dip into wealth to maintain consumption in the face of financial losses in this model may seem quite remarkable. It may raise questions about the extent to which the agent risks running out of reserves, which would hardly be optimal. The answer is that this agent, while dipping into reserves, at the same time takes down risk to a significant extent, an issue to which we now turn.

3.2 Equity share

Figure 4 presents the optimal equity shares for different wealth levels for each alternative value of θ. As pointed out in Mork et al. 2023, the equity share for the case of no habit updating ($\theta = 0$) rises from $\pi/\gamma_1\sigma^2 \approx 0.2$ for very low wealth levels, and approaches the level of $\pi/\gamma_2\sigma^2 \approx 0.6$ asymptotitcally. For $\theta > 0$, the graph lies uniformly above the one for $\theta = 0$, and increasingly so the higher the value of θ. This result is somewhat surprising considering that habit formation in the literature has been advanced as an explanation of the equity premium puzzle, i.e. as an explanation of the apparently low risk appetite of real-world investors, as in Constantinides 1990 or Campbell and Cochrane 1999. The reason given there is that behavior influenced by habit requires a sufficiently large riskless part of the portfolio needs to be set aside to make sure that the habit can be maintained at all times.

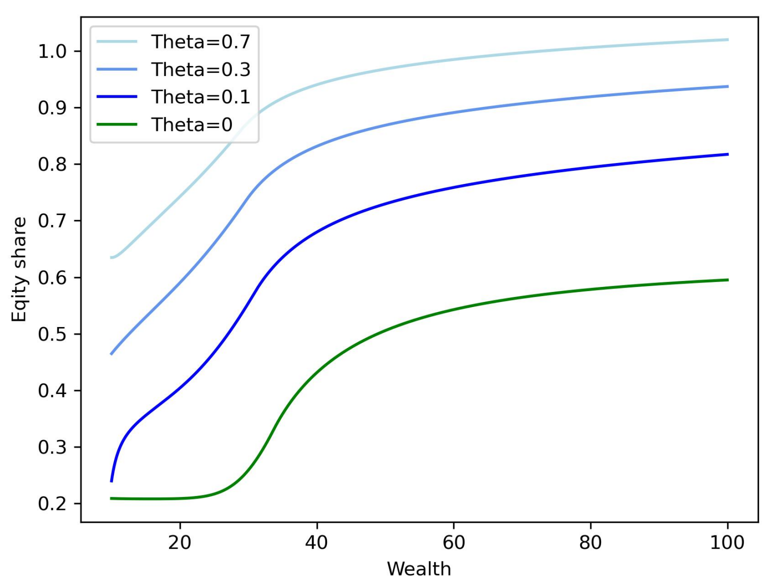

Figure 4: Policy functions for the equity share as a function of wealth, for varying values of the wealth-updating rate θ.

However, when the habit is soft, as in our case, that mechanism is weakened. Furthermore, we get an oppsite effect from the agent's realization that current behavior will affect future habits. For example, although a financial loss will make the agent cut consumption, over time, the agent will get used to the lower consumption level, so that the pain caused by the financial loss will be somewhat short-lived. The effects of a financial gain will be symmetric, of course. Taken together, this means that the habit updating mechanism effectively reduces risk, and more so the faster the updating.

In contrast, optimal behavior requires the agent to take down risk rather aggressively if wealth is reduced below the level for which optimal consumption equals the habit. This feature significantly reduces the probability of wealth falling further. It should make intuitive sense for all investors who feel that financial adversity should be followed by caution in regard to risk taking so as to limit further losses. Needless to say, it contrasts rather sharply with the standard advice of keeping the equity share stable over time, as in the simple CRRA case.

A common consequence of rules requiring stable equity shares is that the portfolio must be rebalanced after large changes in relative asset prices, regardless of direction. Such rebalancing is normally costly. Although our model ignores such costs, we note that it may involve lower transaction costs by reducing the need for rebalancing.

Although we have not been able to prove so rigorously, it seems clear from the above arguments that the equity share in our model uniformly increases with wealth. That is indeed what we find numerically. When considering the model with the transformed variables $y = c/x$ and $\omega = w/x$, it similarly seems to follow from the fact that the implied consumption tax in (11) is progressive, so that the the after-tax value of an additional unit of wealth change is progressively lower as wealth rises. In this way, the tax lowers the effective risk more the higher the wealth, which makes the risk appetite grow with wealth.

These remarks should make it clear that our model offers no help in explaining the equity premium puzzle. But then that is not the purpose of our study as we instead seek to derive the logical implications for the owner of a large, public fund, whose preferences are defined by the objectives given in the introduction. Our work is normative and partial, and thus does not seek to explain any implications of market equilibrium.

4 Dynamic behavior

We use the policy functions implied by our model solution to simulate the dynamics of our agent's behavior. The simulations are based on 10,000 draws of pseudo-random, 50-year histories of stochastic returns. For this exercise, we discretize our continuous-time solution by approximating dt as $\Delta t = 1/2$, corresponding to months. All simulations start at the wealth level at which consumption just equals the initial habit, which we in turn normalize to unity. We make separate simulations for each of our four alternative values of the habit adaption parameter θ.

We start by examining the behavior of equity shares. As we saw in Figure 4, these shares

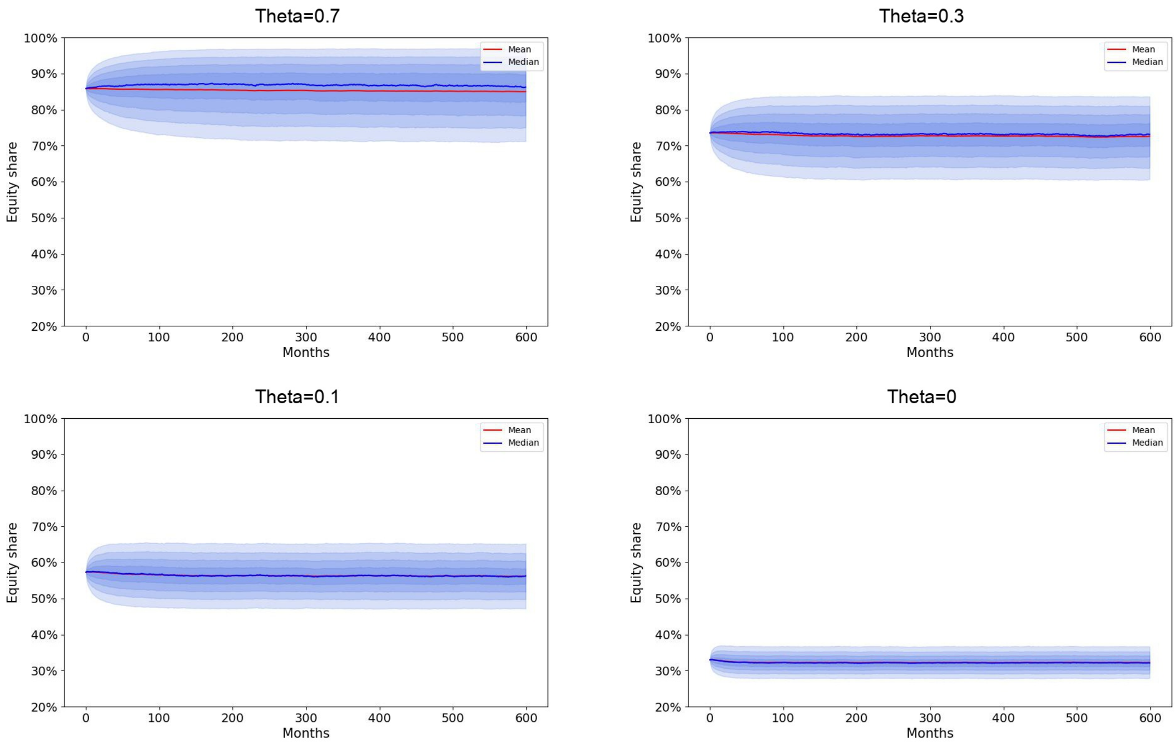

Figure 5: Fan charts for equity shares over time for varying varying rates θ of habit catch-up.

vary significantly with wealth as well as with the rate of habit adjustment. This pattern is equally reflected in the dynamic behavior, as shown in the fan graphs of Figure 5. The figure shows separate graphs for each value of the habit catch-up rate θ. The graphs are all drawn with the same scale on the vertical axis, which helps illustrate how this rate influences both the variability and the central tendency of the respective distributions. As suggested by the policy functions in Figure 4, a higher catch-up rate clearly raises the central tendency as well as the variability. We furthermore note that the distributions appear to be stationary, slightly skewed to the left, and approaching a steady state after at most 25 years.

Next, we consider similar fan graphs for the stochastic behavior of the withdrawal rate, i.e. the consumption-wealth ratio, c/w, over time. In figure 6, we again draw the graphs for the different catch-up rates on the same scales for easy comparison of variability as well as central tendency. The central tendency is not very different for the respective cases, although slightly higher the faster the rate of catch-up θ. All specifications imply some smoothing in the sense that the withdrawal rate varies depending on the agent's fortunes. However, the extent of smoothing clearly depends on the rate of habit updating θ. For $\theta = 0$, i.e. no habit updating, the variability is barely visible, whereas, for $\theta = 0.7$, in which case any gap between consumption and habit has a half life of just a year, the variability is substantial.

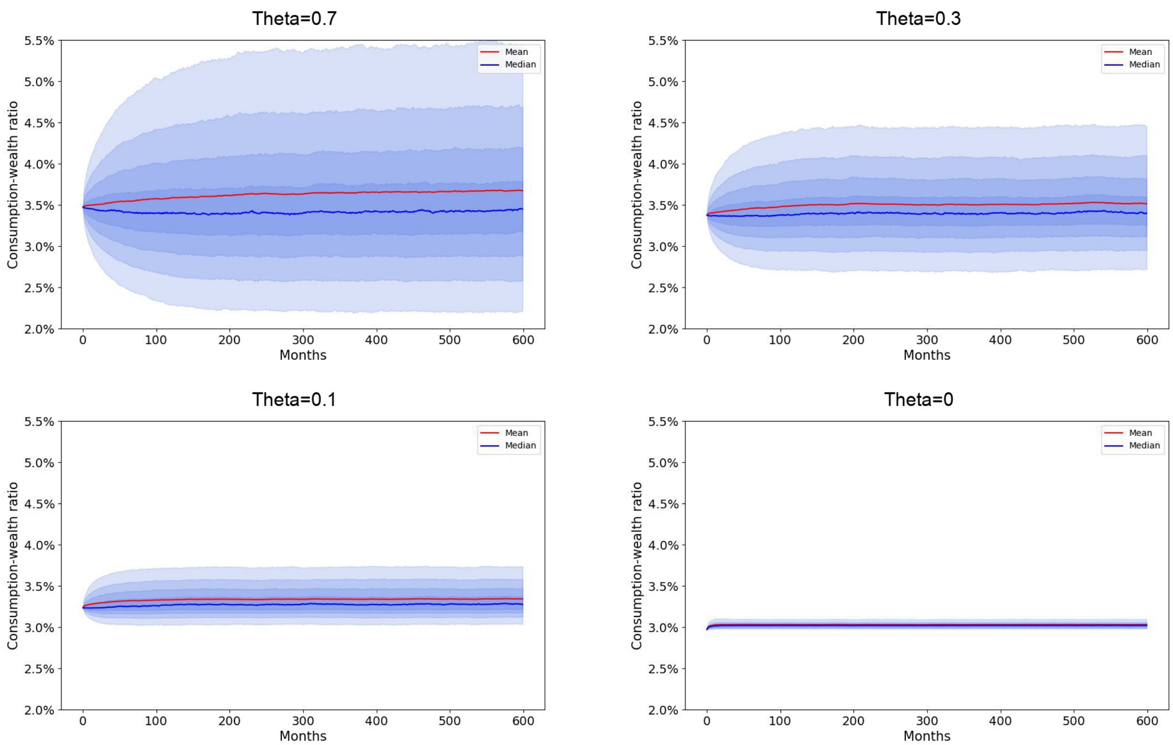

Figure 6: Fan charts for the withdrawal rate (consumption-wealth ratio, c/w), over time.

The difference among the three cases where $\theta > 0$ may seem surprisingly large considering the close similarity of the consumption policy functions in Figures 2 and 3. However, if habits catch up fast, consumption can also react faster to a given financial surprise. This dynamic element then enhances the variability of consumption relative to wealth.

Like the graphs in Figure 5, the ones in Figure 6 all appear to be stationary and to approach a steady state over time. The corresponding distributions for the consumption-habit ratio show a similar pattern, as seen in Figure 7. However, that does not at all mean that consumption itself is stationary. Figure 7 gives a hint in this direction by showing that consumption tends to exceed the habit on average. This feature is partly the result of the asymmetry of the preferences in (1). Because a drop in consumption below the habit level is particularly painful, the agent will seek to avoid it; yet, habits will always tend to catch up with actual consumption. So for consumption to stay ahead, it will need to rise faster than the habit on average. However, a positive endogenous mean growth rate can arise even without habit. In the CRRA version of our model, i.e. with $\theta = 0$ and $\gamma_1 = \gamma_2 = 2$, the endogenous mean growth rate is 1.35 percent.

This "rat race" is clearly illustrated in Figure 9, which shows the mean paths of consumption and habit over time. The mean growth rate is higher the faster the habit catch-up,

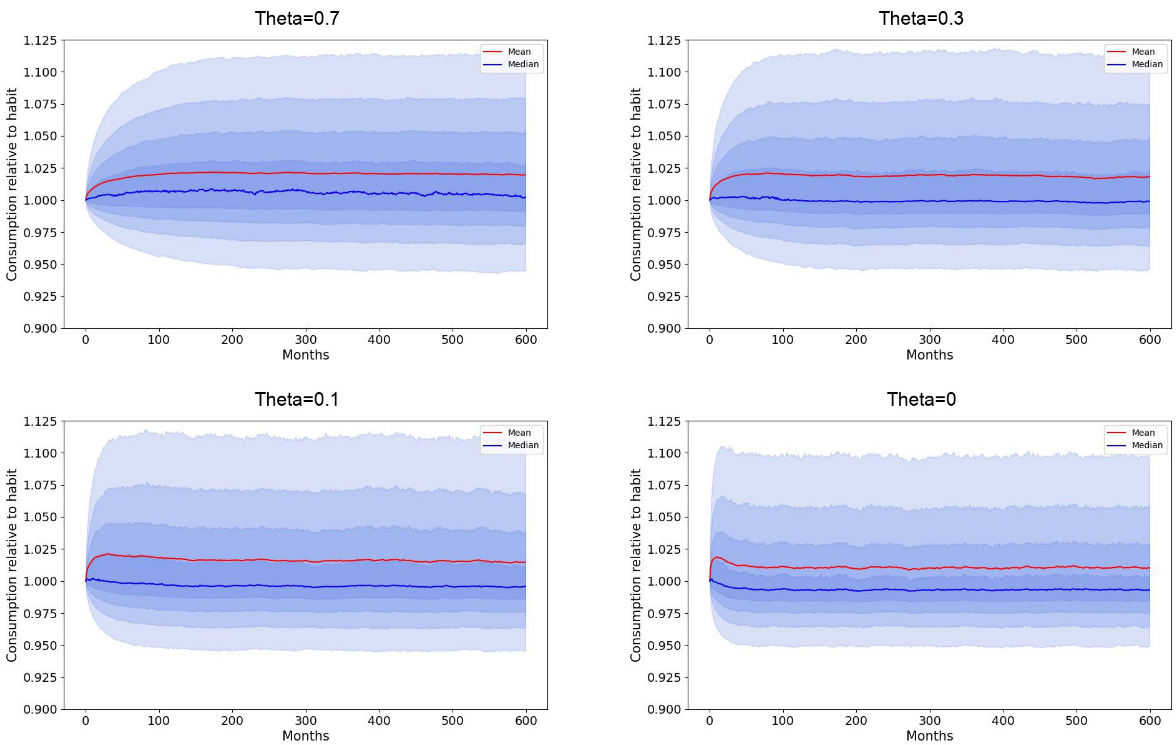

Figure 7: Fan charts for consumption relative to habit (c/x).

as is the mean difference between consumption and habit. As all the graphs in the figure are drawn on the same logarithmic scale, we also see that the mean growth rates appear to asymptotically become constant for consumption as well as the habit.

This persistent tendency for consumption to grow over time may seem surprising because our model has no built-in growth rate. However, because we specify financial returns as exogenous and thus independent of wealth, our model is technically equivalent to a stochastic version of the so-called AK growth models that assume a constant social return to capital, as in Romer 1986. Our model thus displays endogenous growth, albeit in a stochastic sense. For the mean growth rate to be positive, initial consumption must be low enough to allow wealth to grow on average. In particular, the withdrawal rate must, on average, be lower than the expected rate of financial return $r + \alpha\pi$. That rate is not constant because the equity share is highly variable in our model. This variability is clearly implied by the policy functions in Figures 4 and 5.

The variability of the difference between the two rates is somewhat limited by the fact that the withdrawal rate decreases and the equity share increases with wealth. It is nevertheless quite variable, as illustrated in Figure 8. However, the central tendency is clearly for this difference to be negative. The result is that wealth grows on average and thus allows

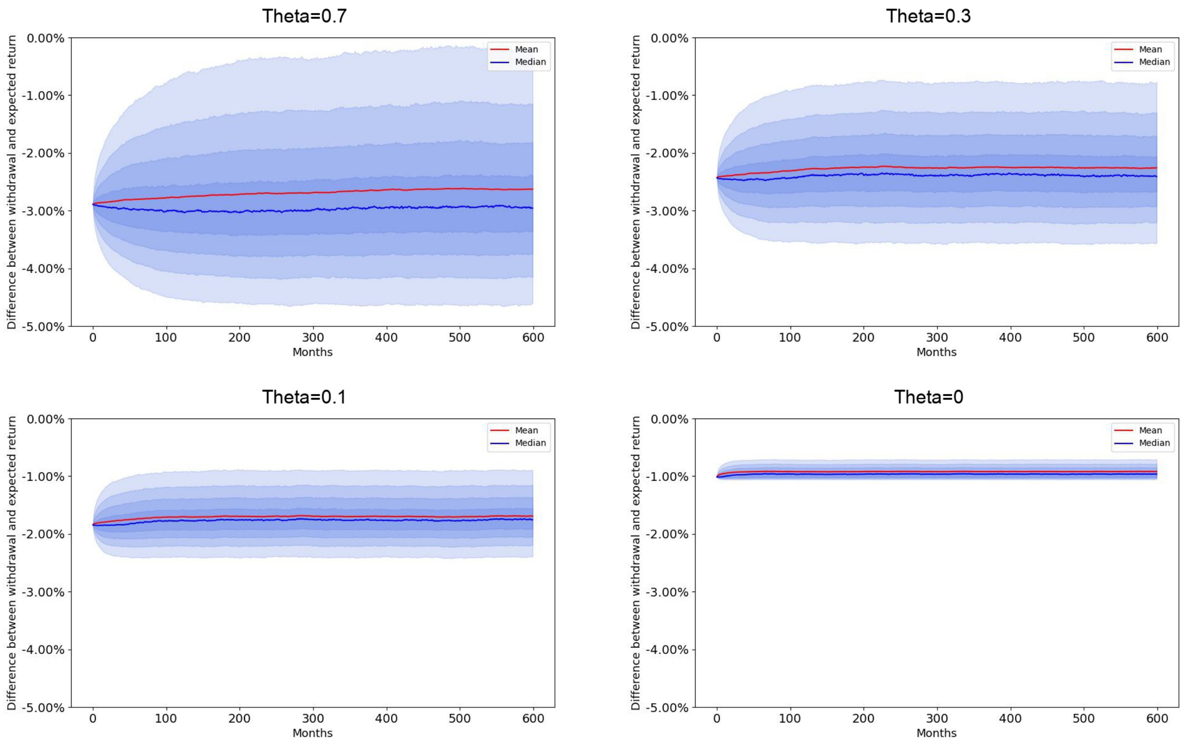

Figure 8: Fan charts for difference between withdrawal rate and expected return

consumption to tend to grow as well.

As a final illustration of the smoothing implied by our model, Figure 10 presents fan charts for the marginal propensity to consume (MPC) out of wealth. These graphs show not only that habit formation tends to lower the MPC, but also that habit updating enhances this effect significantly, even for a catch-up rate as low as $\theta = 0.1$. A higher catch-up rate lowers the MPC further. It also lowers the MPC variability. This result may seem surprising in view of the opposite pattern for the withdrawal rate, as seen in Figure 6. It should be remembered, however, that, a fast catch-up means that the distance between consumption and habit will never be large. As a result, the MPC will be closely concentrated around its value at that point.[4].

[4]All of the graphs in Figure 10 show some curious jumps at the beginning of the simulation period. This is an insignificant side effect of our convention of starting the simulations at a point where optimal consumption just equals the habit. Because this also is the point where the curvature of the utility function shifts, the MPC jumps accordingly.

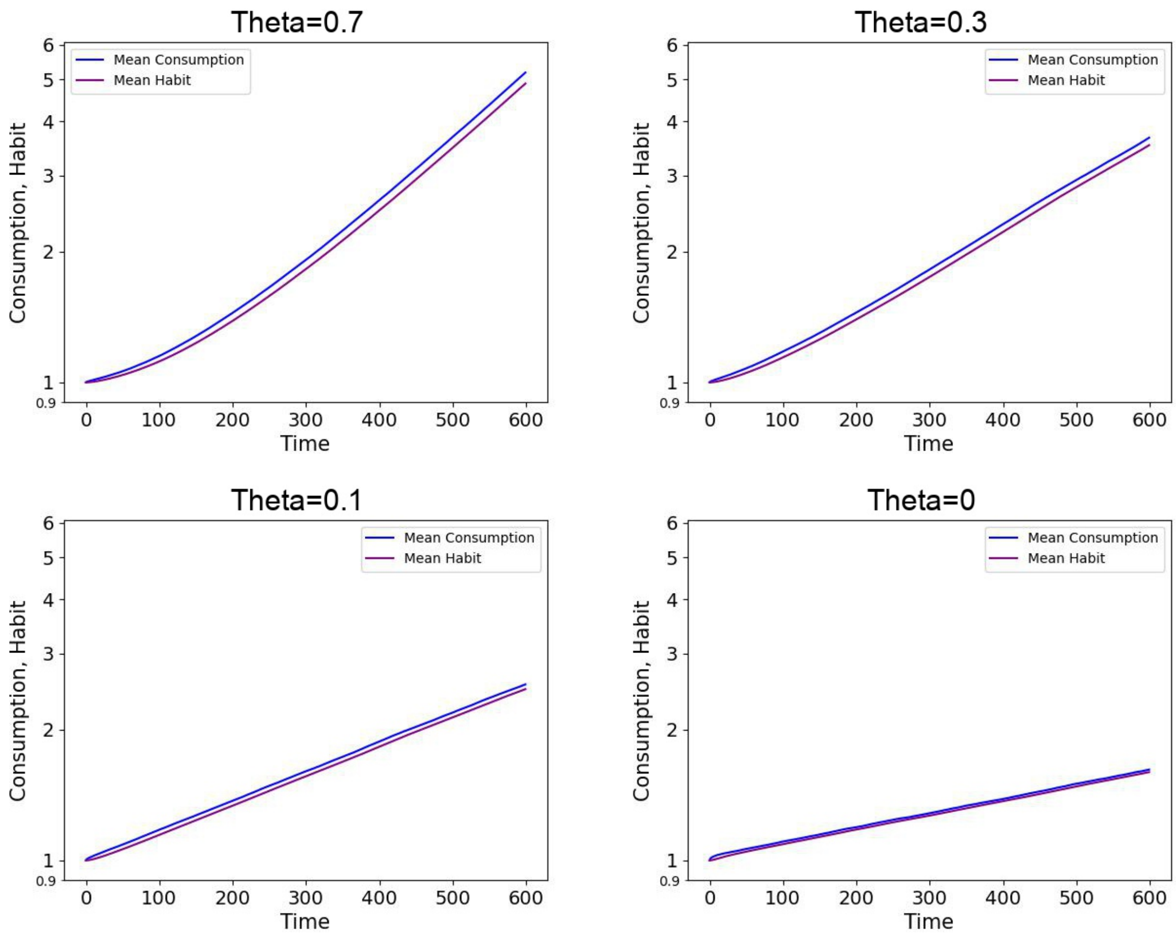

Figure 9: Mean consumption and habit over time a varying rates θ of habit catch-up. The graphs in all the respective panels have been drawn at the same logarithmic scale.

5 The cost of underestimating the habit catch-up rate

Although our agent's behavior is influenced by habit, he or she is sufficiently self-aware to understand, not only the role of the habit, but also how the habit adjusts to recent experience. It may be reasonably argued that this degree of rationality is too high to be likely observed in practice. In a normative setting like ours, we believe it nevertheless should be part of optimal decision making. Even so, our model can be used to estimate the implied cost of being less than fully self-aware. In particular, we study the loss of utility that will result if the agent underestimates the speed of habit catch-up. In that exercise, we continue to assume that the current habit level matters for current decision making but allow the agent's actions to be based on a belief that the catch-up rate θ is lower than its true value.

Technically, we simulate behavior over time based on the policy functions for the cases where the catch-up rate is lower than the one used to update the habit in the simulations.

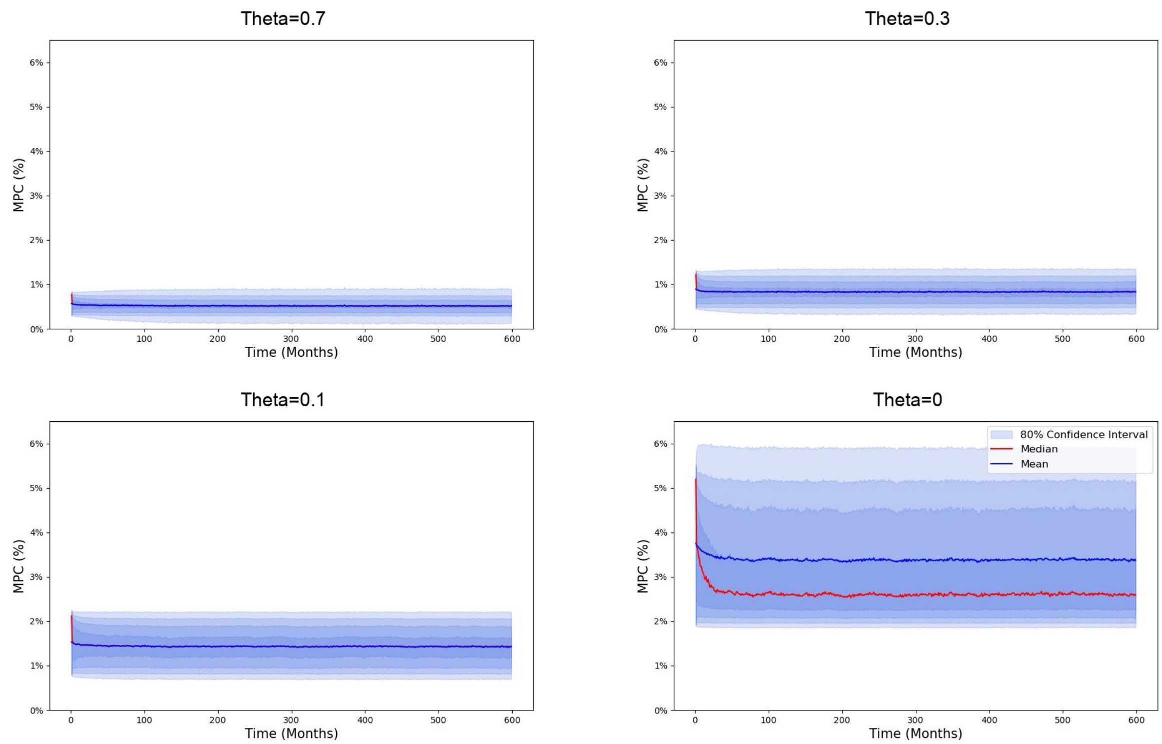

Figure 10: Fan chart of Marginal Propensity to Consume over time.

For example, we use to policy functions for the case of $\theta = 0.3$ to simulate behavior in the case where the habit is actually updated at rate $\theta = 0.7$. We then compute the value of acting under this belief as the mean, subjectively discounted utility from the simulations. We further check the size of the initial wealth that would have implied this magnitude of the value function if the agent's decisions had been based on the true catch-up rate $\theta = 0.7$. The difference between the two can be interpreted as the cost, in terms of initial wealth, of underestimating the catch-up rate.

In Table 2 we report it as a percentage of the inital wealth in the full-awareness case. Each column in the table represents a case of true catch-up, and each row a case of perceived catch-up. For example, the column headed 0.7 shows the losses suffered by an agent with an actual catch-up rate of 0.7, from believing that this rate is 0.7, 0.3, 0.1, and 0, respectively. For the case where the belief rate equals the true rate, the loss is obviously zero. If, however, the perceived rate is lower than the true one, losses occur. For modest amounts of underestimation, such as believeing that $\theta = 0.3$ when actually $\theta = 0.7$, or believing that $\theta = 0.1$ when $\theta = 0.3$, the losses are modest. However, ignoring habit updating all together, i.e. believing $\theta = 0$ when it really is positive, carries a sizeable cost, even if the true catch-up rate is as low as $\theta = 0.1$.

Truth theta / Belief theta	0.7	0.3	0.1	0
0.7	0	–	–	–
0.3	3.4%	0	–	–
0.1	10.8%	5.3%	0	–
0	25.3%	19.3%	11.2%	0

Table 2: Utility loss of underestimating the rate θ of habit catch-up, in units of equivalent percentage loss of initial wealth

6 Simpler rules

The optimal solutions derived in Section 2.4 can be difficult to effectively communicate to managers and/or policy makers. We believe that is one of the reasons that simpler rules typically are used in practice. For the Norwegian GPFG, the mandate to the fund management team stipulates a constant share of equities and/or real estate. Regular rebalancing is required to maintain this share. Similarly, the Storting (Parliament) is permitted to withdraw a fixed percentage of the fund's assets under management each year, although the rule also allows large changes to be spread over two or several years in order to avoid abrupt changes in government services.

In this section, we evaluate the usefulness of such simpler rules, using the same method as in the preceding section. Because this is an optimization in only two dimensions, we use simple grid search. We start by reoptimizing the model for the same set of habit catch-up rates θ under the constraints that the equity share and the withdrawal rate be constant over time. We then repeat the same exercise with the addition of a linear smoothing mechanism for the withdrawal rate of the form

$$c_t = c_{t-\Delta t} + \lambda(c_t^* - c_{t-\Delta t})\Delta t, \tag{14}$$

where $c_t^* = \eta w_t$, and η is the withdrawal rate that would be constant in the absence of smoothing. We subsequently measure the utility loss from these restrictions with and without smoothing. For the smoothing scenarios, we assume $\lambda = 0.5$. The results are presented in Table 3.

Four features are apparent from this table. First, the utility losses are again substantial, of similar orders of magnitudes to those of underestimating the speed of habit catch-up. Second, adding some linear smoothing helps. Some of the literature, like Tobin 1974, have recommended more aggressive smoothing, for example by raising the weight on the most recent withdrawal to 80 percent ($\lambda = 0.2$ in our parametrization). We experimented with

	Equity share	Withdrawal rate	Exp. rate of return	Utility loss
$\theta = 0.7$				
No smoothing	39.0%	3.58%	4.26%	21.6%
Smoothing, $\lambda = 0.5$	65.0%	3.58%	5.43%	10.3%
$\theta = 0.3$				
No smoothing	38.4%	3.55%	4.23%	16.2%
Smoothing, $\lambda = 0.5$	50.8%	3.50%	4.79%	10.6%
$\theta = 0.1$				
No smoothing	36.5%	3.45%	4.14%	9.12%
Smoothing, $\lambda = 0.5$	40.0%	3.40%	4.30%	7.5%

Table 3: Results of constrained optimization with withdrawal rates and equity shares constant over time, with and without smoothing for varying rates of habit catch-up θ. Utility losses are relative to the respective first-best solutions.

this value and found that the utility gain was little more than marginal. In practice, this gain should be weighed against the risk of a losing spiral, as pointed out by Dybvig and Qin 2021. Of course, no such spiral can occur in the optimal solution, not even a constrained one. We furthermore observed no such spiral in our search for the constrained optimum. However, it is worth noting that Mork et al. 2022 found rather disquieting high probabilities of depletion in a simulation study of a sovereign wealth fund charged with funding discretionary fiscal policy.

The third feature of Table 3 is that all the equity shares are significantly lowered than the ones displayed in Figure 5. The reason is that the burden of smoothing, which the habit-driven behavior calls for, now mainly falls on the risk taking, by limiting the volatility of wealth, even when linear smoothing is added. In the first-best solution, the agent has the opportunity to cut down on risk if wealth is in danger of being depleted. Because that feature is lost in the constrained case, risk must instead be kept low throughout.

The fourth and last feature is that the implied withdrawal rates are consistently lower than the corresponding expected returns. This feature contrasts with the widespread rule of equating the two. To shed some light on that practice, we carried out a second set of constrained optimizations where the withdrawal rates are required to equal the expected rates of return. Specifically, we imposed the constraint $\eta = r + \alpha\pi$, where η denotes the withdrawal rate c_t/w_t in the case of no smoothing and c_t^*/w_t in the case of smoothing.

Because the expected rate of return depends on the equity share, the optimization problem is now reduced to picking the optimal constant value for this variable, as in Campbell and Sigalov 2022, who derived the implications of this constrained optimization for the CRRA case without habit formation. The results for our case are reported in Table 4.

	Equity share	Withdrawal rate	Utility loss
$\theta = 0.7$			
No smoothing	31.3%	3.90%	22.9%
Smoothing, $\lambda = 0.5$	39.6%	3.96%	17.8%
$\theta = 0.3$			
No smoothing	30.5%	3.87%	17.7%
Smoothing, $\lambda = 0.5$	34.1%	4.03%	15.1%
$\theta = 0.1$			
No smoothing	28.5%	3.78%	10.9%
Smoothing, $\lambda = 0.5$	28.5%	3.78%	10.4%

Table 4: Results of constrained optimization with withdrawal rates and equity shares constant over time, and withdrawal rates required to equal expected rates of return, with and without smoothing for varying rates of habit catch-up θ. Utility losses are relative to the respective first-best solutions.

The imposition of this requirement does not imply substantial increases in the utility loss compared to the first-best cases. A striking result, however, is that the equity shares now are even lower, and significantly so, especially for the cases with smoothing . This result, which is consistent with the findings of Campbell and Sigalov 2022, may serve as a word of caution to owners of funds such as the Norwegian GPFG, for which withdrawals up to amounts corresponding to the normal real financial return are combined with a 70 percent equity share.

As is well known, when withdrawals are equated to expected returns, the expected future wealth equals current wealth. Our simulation results are naturally consistent with that feature. However, as shown by Dybvig and Qin 2021, the fund value will nevertheless approach zero with probability one when viewed at increasingly long horizons. The fact that these results are both true at the same time is a feature of the lognormal distribution.

7 Discussion and Conclusions

The analysis in the preceding sections have sought to shed some light on what the owners and beneficiaries of a large fund ought to do in response to the realization that past use of fund wealth sets a precedence for future use. Our way of operationalizing this this phenomenon is to specify a set of preferences where the instantaneous utility of current consumption depends on the agent's consumption habit, which in turn depends on past consumption. The habit is soft in the sense that the agent can tolerate consumption that falls short of the habit; it's just especially painful. We furthermore assume that the habit is internal in the sense that it is a function of the agent's own past consumption. That is what sets the precedent.

Our analysis of optimal behavior with such preferences is not easily described in simple terms. One implication is entirely unsurprising, however, namely, that such an agent will want to smooth consumption relative to the vagaries of the financial markets. In fact, the marginal propensity to consume is typically lower than one percent in a setting where the expected rate of portfolio return tends to be about six percent.

A much more surprising result is that risk taking tends to be higher on average than for an agent with CRRA preferences. It contrasts rather sharply with the literature that uses habit formation in attempts to explain the equity premium puzzle. The main reason for the difference is that most of that literature specifies habits as hard, i.e. that the agent will not tolerate consumption that falls short of the habit. In that case, the agent needs to set aside an extra part of the portfolio in riskless assets so at to make sure that the habit can be sustained at all times. With our treatment of the habit as soft, this effect is considerably weakened, while another, opposite effect dominates.

This effect arises from the agent's realization that habits tend to catch up with recent consumption. Specifically, if financial losses make wealth shrink, the agent will realize that he or she will in time will get used to a lower standard of living. Thus, the implied drop in utility will be short-lived and hence in sum more modest. This effect naturally works in reverse as well, thus effectively reducing the implied volatility of utility in response to a given degree of market volatility, thus making it prudent for the agent to take more risk.

However, if wealth shrinks a lot, our agent will want to take down risk significantly. That becomes necessary to prevent wealth from being depleted when smoothing calls for high withdrawal rates when wealth falls to very low levels relative to the habit. In fact, risk taking becomes an increasing function of wealth (relative to habit) in general because the damping of volatility in utility terms that is caused by the habit becomes larger the higher the wealth.

Part of these results derive from our assumption that the agent is sufficiently self-aware to take account of the effect that current consumption decisions will have for future habits. In a normative setting, which is our objective here, we believe this assumption is appropriate.

However, our model can also be used to estimate the utility loss from ignoring this effect partially or completely. Our estimates suggest that such losses may be similar to the one that would arise if the initial wealth were up to 25 percent lower. This effect is naturally stronger the greater the contrast between truth and self awareness.

This result raises the question of whether this agent could have done equally well or better by basing decisions about risk taking and withdrawals on simpler rules rather than the full optimization. Real-world examples include constant equity shares and withdrawal rates over time, although some additional smoothing often is allowed for withdrawals. We use our model to estimate the utility loss from using such simpler rules, where the equity shares and withdrawal rates are determined by constrained optimization. The estimates turn out to be similar in orders of magnitude to the ones obtained by assuming less than full awareness of the the effects of current consumption on future habits. However, the implied equity shares are much lower because part of the desired smoothing now must be carried out in the form of less risk taking, so that wealth fluctuates less.

We also find that the constrained-optimal withdrawal rates are consistently lower than the corresponding expected rates of portfolio return, which contrasts with the rather common rule of equating the two. Because the expected portfolio return depends on the equity share, the constrained decision then boils down to choosing this parameter. Adding this requirement to the simple rules just mentioned, does not carry a heavy additional utility cost. We find it worth mentioning, however, that this result requires even lower risk taking. And although this rule will preserve wealth in expectation, the probability distribution of future wealth will eventually collapse into a spike next to zero.

The main lesson from our analysis thus seems to be that consumption smoothing carries a loss regardless of how it is done. Our first-best solution is superior in the sense that it allows smoothing to be combined with high risk taking on average; however, it requires risk taking to be monitored and adjusted meticulously in response to financial surprises. Linear smoothing on top of a simple rule may seem attractive because it will be much simpler to communicate and implement. But that would require risk to be much more curtailed.

The lessons for the Norwegian GPFG, which motivated this research, could be taken to read that the current Norwegian Fiscal Rule of withdrawal may serve as an acceptable approximation to optimal behavior, considering the tendency for recently past withdrawal to serve as precedence for future withdrawals. However, this rule is in practice combined with the rather high norm of 70 percent for equity and real-estate investments. In addition, the remaining 30 percent investment in bonds carry a fair measure of risk as well. Although continuing new deposits of oil and gas revenues have made this fund grow tremendously, we believe our analysis raises serious questions about its sustainability when this stream of deposits eventually dries up.

For a fuller analysis of the issues surrounding the GPFG, our model needs to be expanded to include government revenues other than the return on the fund, as well as the government's total spending. That is the next item on our research agenda.

References

Abel, A. 1990, 'Asset prices under habit formation and catching up with the joneses', *American Economic Review* **80**(2, Papers and Proceedings), 38–42.

Alhashel, B. 2015, 'Sovereign wealth funds: a literature review', *Journal of Economics and Business* **78**, 1–13.

Angoshtari, B., Yu, X. and Yuan, F. 2024, 'Optimal consumption under loss-averse multiplicative habit-formation preferences'. Available on https://www.researchgate.net/publication/381769548.

Arouri, M., Boubaker, S., Grais, W. and Grira, J. 2018, 'Rationality or politics? the color of black gold money', *Quarterly Review of Economics and Finance* **70**, 62–76.

Arun, T. 2012, 'The merton problem with a drawdown constraint on consumption'. https://doi.org/10.48550/arXiv:1210.5205.

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

Barber, B. and Wang, G. 2013, 'Do (some) university endowments earn alpha?', *Financial Analysis Journal* **69**(5), 26–44.

Becker, G. and Murphy, K. 1988, 'A theory of rational addiction', *Journal of Political Economy* **96**(4), 675–700.

Braunstein, J. 2022, *Capital Choices: Sectoral Politics and the Variation of Sovereign Wealth*, University of Michigan Press, Ann Arbor.

Brown, J., Dimmock, S., Kang, J.-K. and Weisbrenner, S. 2014, 'How university endowments respond to financial market shocks: evidence and implications', *American Economic Review* **104**, 931–962.

Campbell, J. and Cochrane, J. 1999, 'By force of habit: a consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy* **107**(2), 205–251.

Campbell, J. and Deaton, A. 1989, 'Why is consumption so smooth?', *Review of Economic Studies* **56**, 357–374.

Campbell, J. and Sigalov, R. 2022, 'Portfolio choice with sustainable spending: a model of reaching for yield', *Journal of Financial Economics* **143**, 188–206.

Campbell, J., Steins, J. and Wu, A. 2024, Economic budgeting for endowment-dependent universities, *in* 'NBER Conference on the Financing of Higher Education'.

Campbell, J. Y. 2012, Investing and spending: The twin challenges of university endowment management, *in* 'Forum Futures 2012, Forum for the Future of Higher Education: Cambridge, MA', pp. 32–40.

Choi, K., Junkee, J. and Koo, H. 2022, 'Intertemporal preference with loss aversion: consumption and risk-attitude', *Journal of Economic Theory* **200**, 10538.

Constantinides, G. 1990, 'Habit formation: a resolution of the equity premium puzzle', *Journal of Political Economics* **98**(3), 519–543.

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

Dybvig, P. 1995, 'Duesenberry's racheting of consumption: optimal dynamic consumption and investment given intolerance for any decline in standard of living', *Review of Economic Studies* **62**(2), 287–313.

Dybvig, P. 1999, 'Using asset allocation to protect spending', *Financial Analysis Journal* **55**, 49–62.

Dybvig, P. and Qin, Z. 2021, 'How to squander your endowment: pitfalls and remedies'. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3939984.

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

Gilbert, T. and Hrdlicka, C. 2015, 'Why are university endowments large and risky?', *Review of Financial Studies* **28**, 2643–2686.

Hall, R. E. 1978, 'Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence', *Journal of Political Economy* **86**(6), 971–987. Accessed 12 Dec. 2024.
URL: *http://www.jstor.org/stable/1840393*

Hansmann, H. 1990, 'Why do universities have endowments?', *Journal of Legal Studies* **19**, 3–42.

Harang, F. and Irrarazabal, A. n.d., Tba.

Jeon, J. and Park, K. 2021, ‘Portfolio selection with drawdown constraint on consumption: a generalization model’, *Mathematical Methods of Operations Research* **93**, 243–289.

Mehra, R. and Prescott, E. 1985, ‘The equity premium: A puzzle’, *Journal of Monetary Economics* **15**, 145–161.

Mork, K., Harang, F., Trønnes, H. and Bjerketvedt, V. 2023, ‘Dynamic spending and portfolio decisions with a soft social norm’, *Journal of Economic Dynamics and Control* **151**, 104667.

Mork, K., Trønnes, H. and Bjerketvedt, V. 2022, ‘Capital preservation and current spending with sovereign wealth funds and endowment funds: a simulation study’, *International Journal of Financial Studies* **10**, 1–24.

Munck, C. 2008, ‘Portfolio and consumption choice with stochastic investment opportunities and habit formation in preferences’, *Journal of Economic Dynamics and Control* **32**, 3560–3589.

Novy-Marx, R. and Rauh, J. 2011, ‘Public pension promises: How big are they and what are they worth?’, *Journal of Finance* **66**, 1211–1249.

Rackauckas, C. and Nie, Q. 2017, ‘DifferentialEquations.jl–a performant and feature-rich ecosystem for solving differential equations in Julia’, *Journal of Open Research Software* **5**(1).

Rackauckas, C. and Nie, Q. 2020, ‘Stability-optimized high order methods and stiffness detection for pathwise stiff stochastic differential equations’, pp. 1–8.

Rogers, L. 2013, *Optimal Investment*, SpringerBriefs in Quantitative Finance, Springer Berlin Heidelberg.
URL: *https://books.google.no/books?id=zZ5EAAAAQBAJ*

Romer, P. 1986, ‘Increasing returns and long-run growth’, *Journal of Political Economy* **94**.

Shin, Y., Lim, B. and Choi, U. J. 2007, ‘Optimal consumption and portfolio selection problem with downside consumption constraints’, *Applied Mathematics and Computation* **188**(2), 1801–1811.

Shrikhande, M. 1997, ‘Non-addictive habit formation and the equity premium puzzle’, *European Financial Management* **3**(3), 293–319.

Tobin, J. 1974, ‘What is permanent endowment income’, *American Economic Review* **64**(2), 427–432.