Journal of Economic Dynamics & Control 151 (2023) 104667

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Dynamic spending and portfolio decisions with a soft social norm ☆

Knut Anton Mork [a,*], Fabian Andsem Harang [b], Haakon Andreas Trønnes [a], Vegard Skonseng Bjerketvedt [c]

[a] Department of Economics, Norwegian University of Science and Technology (NTNU), Norway
[b] Department of Economics, BI Norwegian Business School, Norway
[c] Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology (NTNU), Norway

ARTICLE INFO

Article history:
Received 29 June 2022
Revised 29 March 2023
Accepted 29 April 2023
Available online 29 April 2023

JEL classification:
C63
E21
G11

Keywords:
Soft social norm
Withdrawal smoothing
Dynamic risk taking
Fund sustainability

ABSTRACT

We explore the implications of a preference ordering for an investor-consumer with a strong preference for keeping consumption above an exogenous social norm, but who is willing to tolerate occasional dips below it. We do this by splicing two CRRA preference orderings, one with high curvature below the norm and the other with low curvature at or above it. We find this formulation appealing for many endowment funds and sovereign wealth funds, including the Norwegian Government Pension Fund Global, which inspired our research. We derive an analytical solution, which we use to describe key properties of the policy functions for consumption and portfolio allocation. We find that annual spending should not only be significantly lower than the expected financial return, but mostly also procyclical. In particular, financial losses should, as a rule, be followed by larger than proportional spending cuts, except when some smoothing is needed to keep spending from falling too far below the social norm. Yet, at very low wealth levels, spending should be kept particularly low in order to build sufficient wealth to raise consumption above the social norm. Financial risk taking should also be modest and procyclical, so that the investor sometimes may want to "buy at the top" and "sell at the bottom." Many of these features are shared by habit-formation models and other models with some lower bound for consumption. However, our specification is more flexible and thus more easily adaptable to actual fund management. The nonlinearity of the policy functions may present challenges regarding delegation to professional managers. However, simpler rules of thumb with constant or slowly moving equity share and consumption-wealth ratio can reach almost the same expected discounted utility. Nevertheless, the constant levels will then look very different from the implications of expected CRRA utility or Epstein–Zin preferences in that consumption is much lower.

© 2023 The Author(s). Published by Elsevier B.V.
This is an open access article under the CC BY license
(http://creativecommons.org/licenses/by/4.0/)

☆ We are indebted to three referees for exceptionally helpful comments and recommendations. Thanks to Snorre Lindset, Alfonso Irrarazabal, and Espen Henriksen for useful comments and suggestions. The research has benefitted from financial assistance from Finansmarkedsfondet, the financial markets research fund under The Research Council of Norway, Project no. 294398.

* Corresponding author.
E-mail address: knut.anton.mork@ntnu.no (K.A. Mork).

https://doi.org/10.1016/j.jedc.2023.104667
0165-1889/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license
(http://creativecommons.org/licenses/by/4.0/)

1. Introduction

This paper is part of an effort to seek sound criteria for the management of sovereign wealth funds (SWF), the level of risk in their portfolios and the use of their proceeds to finance current spending. The importance of this effort has been made apparent by the proliferation of such funds in recent years, as documented, e.g., by Braunstein (2022). Our work has been inspired by the issues facing the owners of the largest of these funds, the Norwegian Government Pension Fund Global (GPFG), currently worth about USD 1.3 trillion. Many endowment funds of universities, museums, and the like, face similar challenges.

Technically speaking, our problem consists of deriving a solution to the Merton (1971) problem and analyzing the implications of such a solution. Merton's seminal article solved this problem for an agent with constant relative risk aversion (CRRA). This specification is mathematically tractable, but also highly restrictive. In particular, it fails to account for the observed magnitudes of the equity premium, which has given rise to the equity premium puzzle, according to the evidence presented by Mehra and Prescott (1985). A number of subsequent authors have explored various modifications of the CRRA assumption as possible solutions to this puzzle, including Constantinides (1990), Campbell and Cochrane (1999), Shirkhande (1997), and Choi et al. (2022).

A common thread in these contributions has been to explore the success of alternative assumptions about investor preferences in matching empirical asset-price data while assuming equity capital to be exogenously given and safe assets in zero net supply, typically in representative-agent general equilibrium models. This approach does not fit the study of SWFs as their objectives and behavior may differ from the majority of other financial agents. As convincingly argued by Gabaix and Koijen (2022), institutional differences among various kinds of agents are important for the proper understanding of short-run as well as long-run movements in global equity markets. Analyses of SWFs and endowment funds must thus instead rely on partial-equilibrium models where asset prices and their stochastic processes are exogenously given. Rather than checking the implications against market data, the analysts must instead check the implications of various specifications against the demands directed at such funds by their owners, managers, and beneficiaries. An important side effect will be for the analysis to reveal necessary tradeoffs and possible inconsistencies among such demands. Conflicts between demands and actual practice thus revealed should be important inputs to fund owners' strategic decisions. These are the kind of questions this paper seeks to answer.

Central among the demands made of SWFs is sustainability, that the value of the fund be preserved over time, typically for future generations. Cochrane (2022) urges long-term investors to focus on long-term returns in the form of coupons, dividends, and share buybacks, rather than current asset values. Others have focused on sustainable consumption, such as Arrow et al. (2004), who, in response to the Brundtland report (World Commission on Environment and Development, 1987), presented and analyzed criteria for sustainable consumption on the global scale. Campbell and Martin (2022) recently added to this line of research by specifying sustainability as a constraint on investment and consumption decisions, requiring that welfare should not be expected to decline over time. This constraint does not distort portfolio choice but does impose an upper bound on the consumption-wealth ratio as well as the rate of time preference.

Sustainability puts restrictions on withdrawals to finance current spending even as beneficiaries want them to be available as a reliable source of budget funding despite shifts in financial fortunes. Such use of the fund as a budget backstop may require spending to be smoothed relative to market movements or even to move in the opposite direction if other revenues decline at the same time as equity prices fall. Thus, governments may want to draw on SWFs to compensate for temporary shortfalls in tax revenues even as universities and museums may want to draw on endowment funds to compensate for shortfalls in tuition or ticket revenues, respectively. And finally, beneficiaries will expect revenue distributions to grow along with the rest of the economy, which is the focus of this paper.

In response to the various and sometimes competing demands, the owners of SWFs and endowment funds have typically set up a number of rules. The Norwegian GPFG is governed by the Fiscal Rule of 2001[1]. This rule seeks to maintain sustainability by limiting annual withdrawals to amounts corresponding to the real financial return on the fund. This return was first stipulated as 4 pct of the fund balance; but in view of the subsequent long decline in global interest rates lowered to 3 pct effective in 2018[2]. However, in case of large financial losses, the rule allows the implied budget cuts to be implemented gradually over time. It also allows temporary tax shortfalls (sometimes referred to as automatic fiscal stabilizers) as well as discretionary, coundercyclical fiscal policy to be funded by extraordinary draws on the GPFG[3]. Many endowment funds employ similar rules. A popular formulation is the so-called MIT – Tobin rule, which stipulates spending as a weighted average of the expected return and last year's withdrawal, often with a weight of as much as 0.8 for the latter[4].

Reasonable as such rules may seem, they tend to suffer from unintended consequences and internal inconsistencies. In the presence of uncertainty, the limitation of withdrawals to expected real returns ensures sustainability only in expecta-

[1] https://www.regjeringen.no/no/dokumenter/stmeld-nr-29-2000-2001-/id194346/, accessed October 5, 2022.

[2] https://www.regjeringen.no/contentassets/6fc0451c5069408791d67ca2fdcc51eb/no/pdfs/stm201720180001000dddpdfs.pdf, p42, accessed March, 2023.

[3] It is worth noting that the Act governing the GPFG does not allow temporary borrowing to be used to smooth spending or fund discretionary fiscal measures as long as there is money left in the fund. Thus, all smoothing and extraordinary spending is applied directly to the withdrawals. See https://www.regjeringen.no/contentassets/9d68c55c272c41e99f0bf45d24397d8c/government-pension-fund-act-01.01.2020.pdf, accessed October 5, 2022.

[4] http://web.mit.edu/fnl/volume/205/alexander_herring.html, accessed on Dec 1, 2020. See also Tobin (1974).

tion[5] and, as shown by Dybvig and Qin (2021) actually implies eventual depletion because of the skewness of the lognormal distribution. Although the latter problem can be overcome by tying withdrawals to geometric rather than arithmetic returns, the simulation study by Mork et al. (2022) demonstrates how active use of fiscal policy makes the depletion of the Norwegian GPFG in finite time a high-probability event.

Although allowing for smoothing may dampen withdrawal fluctuations when stock prices move, smoothing after a price decline unavoidably means raising withdrawals relative to fund value, which in turn raises the risk of early depletion, as pointed out by Dybvig and Qin (op. cit.). An alternative would be to limit fluctuations in fund value by taking down financial risk. Paradoxically, none of the fund rules of which we are aware mention this connection. Decisions about withdrawals and risk taking seem to be made separately without consideration of their interconnectedness. In fact, Campbell and Sigalov (2022) present evidence indicating that rules tying withdrawals to financial returns have made a number of large endowment funds and SWFs, including the GPFG, raise risk in response to falling risk-free rates. While searching for yield might indeed seem attractive for politicians and institutional managers seeking to maintain spending, it hardly seems prudent from a sustainability point of view.

This list of inconsistencies and unintended consequences demonstrates the need for a logically consistent framework to guide the decisions surrounding the kind of funds we are considering here. The route we take toward this end is to propose a set of preferences based on which decision can be made subject to the constraints of the global financial markets. Although, for the reasons mentioned above, we cannot check the implications of our specification against market data in general equilibrium, we make similar checks against the actual behavior of fund owners. As an equally important addition, our method allows us to uncover inconsistences in existing rules, such as when withdrawals and risk taking are determined separately.

The specification explored in this paper is a set of preferences that we refer to as optimization subject to a soft social norm. The norm is exogenous and thought of as a consequence of expectations formed by beneficiaries used to steady growth in the general standard of living. It is similar to habits but not directly related to the person's own past experience or that of others. It is soft in the sense that consumption below the norm is possible but especially painful. Technically, we do this by splicing two CRRA instantaneous utility functions, one with high and one with low curvature, depending as consumption is below or above the social norm, respectively. Marginal utility is continuous everywhere but has a kink at the social norm.

This specification was designed as a way to incorporate some of the expectations and demands discussed above. Financial sustainability is ensured by optimization on an infinite horizon. The soft norm was introduced as a way to model aversion against large cuts in consumption. In this sense, it favors some smoothing. Perhaps surprisingly, however, explicit smoothing of consumption relative to market volatility is only a minor part of the solution. A more efficient way to avoid consumption below the social norm involves sharp cuts in risk taking when wealth declines, which dampens the impact of market volatility. At the same time, extreme thriftiness at low wealth levels helps build wealth and thus limits the probability of having to consume less than the social norm in the future. These findings are part of the main contributions of this paper and should be useful for fund owners when reviewing their rules for fund management and withdrawals.

The specification introduced in this paper represents a first step in our research agenda and does not address all the expectations, demands, and possible inconsistencies referred to above. Subsequent steps will seek to widen our approach to address further concerns. We are already working on a model where the exogenous soft social norm is replaced by the formation of soft habits, internal or external, so that the concept of lower than acceptable consumption is formed by the agent's own experience or that of the society at large. A further step will be to introduce business cycles and a corresponding cyclical pattern of withdrawals, with business cycles modeled as a stochastic process that may be correlated with, but distinct from financial volatility.

Our specification is closely related to the literature on habit formation, which builds penalties into the utility function in the form of fast rising marginal utility if spending is cut relative to an established norm. Habits can be formed internally as an aggregation of the agents' own consumption (Abel, 1990; Constantinides, 1990), or as an attempt to keep up with habits formed in society at large (Abel, op.cit.; Campbell and Cochrane, 1999). Such models have been used in the macroeconomics literature to account for the so-called excess smoothness of consumption (e.g. Campbell and Deaton 1989, Galí 1990). Our specification is simpler in the sense that our social norm is completely exogenous, but significantly more general in the sense that the penalty it puts on consumption below the norm is finite rather than the infinite penalty implicit in the habit-formation models. Downward constraints on consumption are even stricter in models with drawdown constraints, such as Shin et al. (2007), Arun (2012), and Jeon and Park (2021). Dybvig (1995, 1999) goes further still by constraining consumption to be non-decreasing over time, in the spirit of Duisenberry's (1949) hypothesis about household behavior.

Such downward constraints do not limit the feasibility of the various models in equilibrium because the agents in question will save enough in good times and invest enough of their portfolios in riskless assets to be able to maintain the respective lower-bound consumption levels in bad times. However, the agent's initial wealth must be high enough to support consumption at the lower bound from riskless investments. This constraint would become a real issue for real-world

[5] Mork and Trønnes (2023) show that the future fund value will not even be preserved in expectation if the financial returns are negatively serially correlated. This is likely to be true for a fund domiciled in a small country that is invested in the global market in foreign currencies with spending withdrawals based on the fund's value in the domestic currency if the exchange rate obeys long-term purchasing parity. Mork, Trønnes, and Bjerketvedt (op. cit.) find this effect to be quantitatively important for the GPFG.

fund owners who consider transitioning to the rules implied by a habit-formation or similar model. Suppose, for example, that the Norwegian GPFG, whose current rule is to withdraw 3% of the fund's value annually, were to switch to rules corresponding to a model such as Constantinides (1990) or Dybvig (1995). The fund, whose value at this writing is about USD 1.3 trillion, contributed USD 30 billion to the 2022 government budget. To prevent this contribution from declining, as in Dybvig's case, the entire fund would have to be invested without risk if the real riskless rate is 2.5%. With Constantinides' model, the habit level would likely be lower than the actual spending, so that the riskless share could be somewhat less than 100%. Even so, the implication for risk taking would contrast sharply with the fund's current mandate, according to which 70% of the fund can be invested in equity or real estate and another 9% in corporate bonds.

A specification that is more similar to ours in this respect can be found in the loss-aversion models of Shirkhande (1997), Watson and Scott (2015), and Choi et al. (2022). Like us, they assume that the utility loss for lowering consumption below the relevant barrier is finite rather than infinite. The barrier in their models is not exogenous like in ours, but equal to the agent's recent consumption. Shirkhande refers to this specification as a "non-addictive habit." The endogeneity of the barrier is a strength of this model. It comes at the cost of two weaknesses, however. First, habits are formed extremely fast in that the agent instantaneously gets used to whatever the current consumption level might be. Second, as a consequence of the model setup, changes in consumption come only as jumps at discrete points in time. In between these jumps, consumption is constant. At every moment, the agent weighs the costs (in terms of current or future penalties) of changing consumption against the potential benefit, in analogy with staggered price setting in models where price changes are costly, such as Rotemberg (1982) and Mankiw (1985). This feature makes the model tractable and is consistent with previous models of adjustment costs for consumption changes, going back to Grossman and Laroque (1990). However, although loss aversion has been supported by a number of experiments, it represents a deviation from rational decision making as this concept is commonly understood, which we feel makes it less suitable in a normative context. Perhaps not surprisingly, then, we find that the implications of our model are quite different from those of the loss-aversion models.

Methodologically, our main contribution lies in our solution of the Merton problem for an agent whose preferences can be described by a splicing of two CRRA utility functions that are identical except for the value of the curvature parameter γ, so that the curvature of the spliced function is much stronger for low levels of consumption than high ones. This feature allows us to model softness in the self-imposed constraint without loss aversion. Our utility function has the standard characteristics of monotonicity and concavity; the only unusual trait is that the second derivative makes a jump at the point of splicing. This formulation helps us avoid the assumptions of additional restrictions that are seen in much of the literature, as in Dybvig's (op.cit.) constraint that consumption never decrease; the drawdown constraints of Shin et al. (op. cit.), Arun (op. cit), and Jeon and Park (op. cit); and the jumpiness constraint of Shirkhande (op. cit.) and Choi et al. (op. cit.). One of the problems of such constraints is that it often is unclear whether they are part of the agent's preferences or imposed by the environment in the same way as budget constraints. By placing our critical assumption squarely in the utility function, we avoid that confusion.

Our analysis finally adds to the literature on sovereign wealth funds and institutional endowment funds. Among recent studies, Braunstein (2022) discusses how the purpose of the various funds differ depending on the political context. Earlier surveys have been written by Baldwin (2012) and Alhashel (2015). Arouri et al. (2018) similarly discusses their varying investment strategies, as did Bernstein et al. (2013), Paltrinieri and Pichler (2013), Dreassi et al. (2017) and Johan et al. (2013). Empricial studies of endowment funds include Barber and Wang (2013) Brown et al. (2014), and Dahiya & Yermack (2018). A special issue has been taken up by van der Bremer et al. (2016) and Irarrazabal et al. (2023), who discuss the joint decision of exhaustible-resource extraction and portfolio choice for a sovereign wealth fund established to safeguard the revenues of 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 spliced utility function that we use to represent the soft social norm and sets up the dynamic optimization problem. Section 3 derives an analytical solution. Although this solution proves too complicated to be very informative, it nevertheless allows derivation of key properties of the implied policy functions for risk taking and consumption, which are presented and illustrated in Section 4. Section 5 simulates the model's dynamic behavior. Section 6 discusses the results, partly by comparing them to related models in the literature, partly by highlighting the possible implications for sovereign wealth funds, and partly by considering how the actions suggested by our model can be delegated and carried out in practice. Section 7 summarizes the results and concludes.

2. The model

Our idea of a soft social norm involves a certain norm of the consumption level, denoted by $x > 0$, and a utility function whose curvature becomes flatter when consumption rises above this norm. A natural starting point might be a specification of marginal utility of the form

$$u'(c) = c^{-\gamma(c)},$$

where, γ is a decreasing function, possibly depending on the norm x. If this function is differentiable, the standard curvature measure becomes

$$-\frac{u''(c)c}{u'(c)} = \gamma(c) + \gamma'(c)c\ln c,$$

which may be negative for certain levels of consumption, depending on the structure of the chosen γ, thus violating a fundamental feature of utility functions.

To overcome this issue while still keeping the model tractable, we propose instead to choose the risk-aversion function γ as a step function, falling discontinuously from a higher to a lower value when consumption falls below a certain level we identify as the social norm. That is, we start by considering two CRRA utility functions, differing only in the attitude towards risk:

$$u_i(c) = \frac{c^{1-\gamma_i} - 1}{1-\gamma_i}, \;\; i = 1, 2; \;\; \gamma_1 > \gamma_2. \tag{1}$$

We refer to these preference orderings as CRRA_1 and CRRA_2, respectively[6]. We assume that the agent has access to two assets, one risky asset that we refer to as equity and that moves according to a geometric Brownian motion with mean μ and volatility σ, and one riskless asset with constant return r. These two assets are combined by investing a share ω of wealth in the risky asset and $1-\omega$ in the riskless asset. The portfolio constructed by the investor is self-financing. The investor is allowed to consume at rate c from total wealth, and thus the wealth of the investor is described by the stochastic differential equation

$$dW_t = W_t(r + \omega_t\pi)dt + W_t\sigma dB_t - c_t dt, \;\; W_0 = w, \tag{2}$$

where, $\{B_t\}_{t\geq 0}$ describes Brownian motion, and $\pi :- \mu - r$ is the equity premium, which we assume to be positive for the optimization problem to be well posed. Given an instantaneous utility function u, the objective of the agent is to maximize expected utility of discounted consumption over an infinite time horizon, i.e. to maximize the following payoff function over any admissible c and ω

$$P(t, w|c, \omega) = \mathbb{E}\left[\int_t^{\infty} e^{-\rho(s-t)} u(c_s)ds | W_t = w\right]. \tag{3}$$

If we choose the utility function $u = u_j$ for $j = 1, 2$, the optimal equity share would be, as shown by Merton (1971),

$$\omega_i^* = \frac{\pi}{\gamma_i\sigma^2}. \tag{4}$$

The optimal rate of consumption would similarly be given by the Keynes–Ramsey rule:

$$c_i^*/w = \eta_i :- (1/\gamma_i)\rho + (1 - 1/\gamma_i)\left(r + \frac{1}{2}\omega_i^*\pi\right), \;\; j = 1, 2, \tag{5}$$

where, ρ is the subjective discount rate, as in Eq. (3), and w represents the value of the portfolio.

2.1. Spliced utility function

We now use the two preference orderings in Eq. (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 social norm x:

$$u(c, x) = \begin{cases} \frac{1}{1-\gamma_1}\left[\left(\frac{c}{x}\right)^{1-\gamma_1} - 1\right], & c < x \\ \frac{1}{1-\gamma_2}\left[\left(\frac{c}{x}\right)^{1-\gamma_2} - 1\right], & c \geq x, \end{cases} \tag{6}$$

where,[7] $\gamma_1 > \gamma_2$. This function is continuous and concave everywhere in its first argument. It is increasing in consumption and decreasing in the social norm; and the mapping $\alpha \mapsto u(\alpha c, \alpha x)$ can easily be seen to be homogeneous of degree zero.

Let u_c denote the derivative in the first variable of the function u and u_x the derivative in the second variable. One can readily check that u is such that the first partial derivative with respect to c satisfies $u_c(c, x) = c^{-\gamma(c,x)}$, where

$$\gamma(c, x) = [\gamma_1 + (1-\gamma_1)log_c(x)]\mathbb{1}_{c<x} + [\gamma_2 + (1-\gamma_2)log_c(x)]\mathbb{1}_{c\geq x}.$$

Furthermore, the first derivative in the first variable of u is continuous on the positive real numbers and continuously differentiable everywhere. At the point $c = x$ the first-order derivative has a kink. The second-order derivatives are continuous everywhere except in the point $c = x$, where a jump happens. The main attractive feature of this utility function is that at $c = x$, the standard curvature measure jumps from γ_1 to γ_2. It may also be worth noting that utility has the same sign as $c - x$.

Fig. 1 displays the marginal utility of an agent with a soft social norm of $x = 3$, and of the two associated CRRA agents. The graph is drawn on the assumption that $\gamma_1 = 6$ and $\gamma_2 = 2$. Thus, this agent is extremely averse to consumption c falling below $x = 3$. Yet, the agent can accept such an outcome if the alternatives are worse.

[6] Formula (1) implicitly assumes $\gamma_i \neq 1$. If $\gamma_i = 1$, this function is replaced by $u_i(c) = \ln(c)$.

[7] As in (1), this formula implicitly assumes $\gamma_i \neq 1, \;\; i = 1, 2$. If $\gamma_i = 1$, the relevant function is replaced by $\ln(c/x)$.

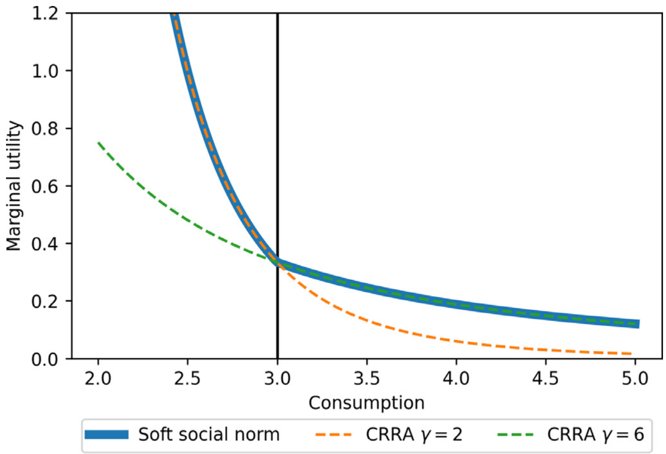

Fig. 1. Marginal utility with a soft social norm.

For comparison, we note that Constantinides' (op. cit.) habit utility function can be written in our notation as

$$u^H(c, x) = \frac{(c-x)^{1-\gamma} - 1}{1-\gamma}, \tag{7}$$

where, the function is only defined for $c \geq x$. Although this is not a special case of our specification, it is somewhat similar in that both specifications imply sharply larger marginal utility for low levels of consumption. In the loss-aversion specification used by Choi (op. cit.) and Shrikhande (op. cit.), consumption moves only in discrete jumps. The instantaneous utility function is

$$u^{LA}(c) = \frac{c^{1-\gamma} - 1}{1-\gamma} - \alpha \Delta u(c)^+ - \beta \Delta u(c)^-, \tag{8}$$

where, $\Delta u(c)^+$ and $\Delta u(c)^-$ denote the changes in utility induced by positive and negative jumps in consumption, respectively, and $\alpha \geq 0$ and $\beta > 0$ are penalties for positive and negative jumps, respectively. Our assumption of a social norm is somewhat similar to this specification's penalty for downward jumps in that the utility loss when consumption falls below the social norm is greater than a similar rise from the same starting point. However, consumption in our model does not move in jumps; and the starting point, given by the immediately past consumption level, does not play the same significant role in our model as it naturally does in a loss-aversion model. Finally, the penalty for letting consumption fall below the social norm follows directly from the form of our utility function without any ad hoc constraints.

2.2. Dynamic setup

For a fixed habit level x, let the value function V be described by the optimization $V(t, w, x) = \max_{c,\omega} P(t, w|c, \omega)$, where the payoff function P is built with the utility function defined in Eq. (6). Because our model is partial, we assume that the social norm is dynamic, exogenously given, and that it moves with a constant drift g, so that

$$dx_t = x_t g dt, \quad x_0 = x. \tag{9}$$

In continuous time, the associated Hamilton–Jacobi–Bellman (HJB) equation describing the dynamics of $V(t, w, x)$ can then be written as:

$$\max_{c,\omega} \left\{ u(c, x) + V_t + V_w(r + \omega\pi)w - V_w c + \frac{1}{2}(\sigma\omega w)^2 V_{ww} + V_x x g \right\} = 0.$$

The reader is referred to, for example, Rogers (2013) for a detailed derivation of this equation. Furthermore, due to the time-homogeneity of the payoff-function, one can easily check that $V(t, w, x) = e^{-\rho t} V(w, x)$, where $V(w, x)$ then must satisfy

$$\max_{c,\omega} \left\{ u(c, x) + V_w(r + \omega\pi)w - V_w c + \frac{1}{2}(\sigma\omega w)^2 V_{ww} + V_x x g - \rho V \right\} = 0. \tag{11}$$

Partial differentiation with respect to consumption c and the equity share ω yields the first-order conditions

$$u_c(c^*, x) = V_w(w, x) \tag{12}$$

and

$$\omega^* = -\left(\frac{V_w}{V_{ww} w}\right)\left(\frac{\pi}{\sigma^2}\right), \tag{13}$$

respectively.

When Eq. (12) and Eq. (13) are substituted into the HJB equation, it becomes apparent that our model is scalable in the sense that the value function must share the homogeneity property of the utility function and thus be homogeneous of degree zero in w and x. It then follows from Eq. (13) that the optimal equity share must be homogeneous of degree zero in w and x as well. Furthermore, the policy function for consumption, defined implicitly by Eq. (12), must be homogeneous of degree one in w and x, sharing this property with the consumption policy function of the classical Merton problem with a standard CRRA utility function. Indeed, combining Eq. (12)) with the fact that u_c is homogenuous of degree -1, it follows that

$$u_c\left(\frac{c^*(\alpha w, \alpha x)}{\alpha}, x\right) = V_w(w, x),$$

and therefore,

$$c^*(\alpha w, \alpha x) = \alpha u_c^{-1}(V_w(w, x), x) = \alpha c^*(w, x).$$

By extension, the implied policy function for the consumption-wealth ratio is homogeneous of degree zero in the same two variables. In other words, the policy functions for the equity share and the consumption-wealth ratio depend on the wealth-norm ratio w/x only.

Substituting the optimal values transforms the HJB equation into a partial differential equation (PDE) for the value function V in the variables w and x, we obtain:

$$u(c^*(w, x), x) - \rho V(w, x) + V_w(w, x)\left[\left(r + \frac{1}{2}\omega^*(w, x)\pi\right)w - c^*(w, x)\right] + V_x(w, x)xg = 0$$

However, the homogeneity property that we just derived transforms it into an ordinary differential equation in w only, for given fixed $x > 0$. That is, Euler's homogeneous function theorem implies that

$$V_x(w, x)x = -V_w(w, x)w.$$

Suppressing the arguments of V and V_w, we can then write the HJB equation in compact form as

$$u(c^*, x) - \rho V + V_w\left[\left(r + \frac{1}{2}\omega^*\pi - g\right)w - c^*\right] = 0. \tag{14}$$

The presence of the second-order derivative V_{ww} in the expression Eq. (12) for the optimal equity share ω^* makes this a non-linear ODE of the second order in w/x.

3. Analytic solution

Although the solution to the nonlinear ODE Eq. (14) is not simple it is intuitively appealing. In this section, we derive a solution formula for the value function, which implies informative key characteristics of the policy functions for consumption and the equity share. Even though the explicit form of the solution formula is rather involved and might not be economically intuitive in itself, we provide a brief discussion of this formula and complement it with a numerical analysis of the implications.

Given a general utility function u to be used in the infinite horizon Merton problem, one would not in general expect a closed form solution. However, our proposed utility function can be seen as a state dependent combination of two CRRA functions. It is well known that a closed form solution to the Merton problem can be found when the utility function is a simple CRRA function, and it is therefore natural to think that the same methods could be applied to our case. By following for example Rogers (2013) Section 1.3, we see that the dual value function approach to solving the Merton problem will allow us to find certain closed form expressions. We derive this method in full in the Appendix and give a brief overview of the findings in this section.

As seen from expression Eq. (6) in the preceding section, the utility function is divided into two regimes depending on the social norm x, namely through the indicators $\mathbb{1}_{c<x}$ and $\mathbb{1}_{c\geq x}$. We begin to fix a level of the social norm at a value $x > 0$. By doing a change of variables, we will let z implicitly represent the marginal utility of consumption, that is, we define

$$(t, z) = (t, V_w(t, w, x)),$$

where, $(t, z) \in A :- [0, \infty) \times (V_w(t, \infty, x), V_w(t, 0, x))$. Define a new function $J(\cdot|x) : A \to \mathbb{R}$ by $J(t, z|x) = V(t, w, x) - wz$, and then through certain manipulations (explicitly proven in Appendix 1) we find that the dual value function can be represented by

$$J(t, z|x) = F_1(t, x)(xz)^{1-1/\gamma_1} + F_2(t, x)(xz)^{1-1/\gamma_2} + zG(t, x). \tag{15}$$

Here F_1 and F_2 are two functions capturing the weighted expectation of the indefinite integral of the fractional moment of a geometric Brownian motion conditioned on the event of being above or below the soft social norm x; and the function G captures the weighted probabilities of a geometric Brownian motion being above or below the norm x.

With this expression for $J(t,z|x)$, we may now derive the quantities of interest based on $V(t,w,x)$, namely the consumption and equity share policy functions. We begin to observe that, by the definition of $J(t,z|x)$, we have $J_z(t,z|x) = -w$. Using this fact, differentiating the dual function we see that

$$(1-1/\gamma_1)F_1(t,x)z^{-1/\gamma_1}x^{1-1/\gamma_1} + (1-1/\gamma_2)F_2(t,x)z^{-1/\gamma_2}x^{1-1/\gamma_2} + G(t,x) = -w. \tag{16}$$

Recalling that z represents the marginal value of wealth $V_w(t,w,x)$, which at the optimum must equal the marginal utility of consumption, one can invert the above relation in order to obtain an equation of the form

$$V_w(t,w,x) = z = H(t,w,x), \tag{17}$$

for some suitable function H. Note that the function H will be positive as $z > 0$ by definition.

From Eq. (17)(and the first-order condition Eq. (12) for consumption, we clearly have

$$c^*(t,w) = (u_c)^{-1}(H(t,w,x)).$$

To make this expression a little clearer, note that from Eq. (6), we can write it as

$$c^*(t,w) = \begin{cases} H(t,w,x)^{-1/\gamma_1}x^{1-1/\gamma_1} \; if \; H(t,w,x) > 1/x \\ \\ H(t,w,x)^{-1/\gamma_2}x^{1-1/\gamma_2} \; if \; H(t,w,x) \le 1/x. \end{cases} \tag{18}$$

Note that $u_c(x,x) = 1/x$, and recall that the marginal value of wealth equals marginal utility at the optimum. Thus, Eq. (18) shows that the formula for optimal consumption differs depending on whether wealth is such that its marginal value is higher or lower than the marginal utility of just consuming the social norm.

This observation makes it tempting to conclude that behavior with our spliced utility function will be a simple combination of behavior under the two CRRA component utility functions. However, that is not the case because optimal choices depend not only on the current state (i.e. whether $H(t,w,x)$ is greater than or less than $1/x$), but also on the conditional probability of ending up in another state at various times in the future, as expressed in a highly summary way by the dependence of H on x. In a very broad sense, this means that the solution will behave in a manner similar to a combination of the solutions for the component preferences CRRA_1 and CRRA_2, respectively. For example, the wealthier the agent is relative to the social norm, the more likely it is that the marginal value of wealth in future periods will fall short of $1/x$, which in turn would make future behavior more like what would be observed under the component utility function that we have called CRRA_2. However, this similarity is much more involved than a linear combination. Its nature will be made clearer in the next subsection.

We should note that the above method for finding a closed form solution in the case where the social norm is fixed at a level x, can easily be extended to the general case of a growing social norm at a rate $g > 0$. In Appendix 2 we show that if we denote by $c_t(w,x_0)$ the optimal consumption policy with wealth level w and fixed social norm x_0, then the optimal consumption policy $c_t^g(w,x_0)$ when the social norm is dynamically growing and starting in x_0, i.e. with social norm given by $x_t = x_0e^{g\,t}$, is given by $c_t^g(w,x_0) = e^{gt}c_t(w,x_0)$.

The optimal portfolio weights can be derived from the dual value function in a similar manner to consumption. Inserting the derived function $H(t,w,x)$ for the marginal value of wealth in the first-order condition Eq. (13) for the equity share, we see that

$$\omega^* = -\left(\frac{H(t,w,x)}{H_w(t,w,x)w}\right)\left(\frac{\pi}{\sigma^2}\right). \tag{19}$$

4. Policy functions

Because the left-hand side of Eq. (16) consists of a sum of two different powers of z in addition to the additive term $G(t,x)$, the expression $H(t,w,x)$ cannot be expressed in a simple closed form. For this reason, the solution is not very informative by itself. However, it allows us to derive some key analytical features of the policy functions for consumption and the equity share and numerical computation of these functions for reasonable values of γ_1 and γ_2.

As a general observation, we note from Eq. (15) that the dual value function J has smooth first and second order derivatives with respect to z, which represents the marginal value of wealth. From Eq. (16) and the inversion implicit in Eq. (17), we then see that the function H is smooth and positive and that its partial derivative H_w must be smooth and negative. It then follows from Eq. (19) that the policy function for the equity share also must be smooth and positive. Similarly, we see from Eq. (18) that the policy function for consumption is continuous but has a kink at the point where $H(t,w,x) = 1/x$. For given t and x, we denote the corresponding wealth level as w^x, so that $H(t,w^x,x) = 1/x$.

4.1. Policy function for the equity share

As seen in Eq. (13) and Eq. (19), the equity share is proportional to the reciprocal of the curvature of the value function. The latter is a decreasing function of wealth because higher wealth implies lower conditional probabilities of having to live with consumption below the social norm. Thus, the optimal equity share ω^* increases with wealth. As $w \to \infty$, the

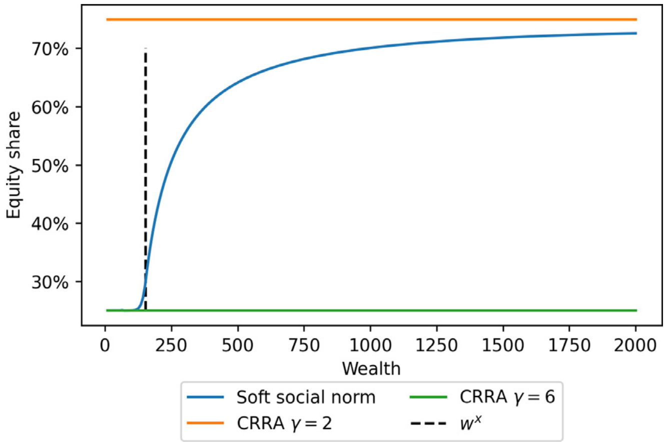

Fig. 2. Policy function for the equity share.
The red and green lines show the optimal equity shares for the CRRA preference orderings with relative risk aversion of 2 and 6, respectively. The dashed vertical line denotes the wealth level w^x at which the optimal consumption equals the social norm. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 1
Parameter values (annual rates).

Variable	Symbol	Value
Equity premium	π	4.8 pp
Equity return std. dev.	σ	17.89%
Riskless return (real)	r	2.5%
Subjective discount rate	ρ	4%
RRA for $c^* < h$	γ_1	6
RRA for $c^* \geq h$	γ_2	2
Social norm growth rate	g	1.9%

curvature approaches the curvature measure γ_2 of the preference ordering CRRA$_2$. Consequently, as $w \to \infty$, $\omega^* \to \omega_2$ as defined in Eq. (4). Because of monotonicity it approaches this limit from below. And because it is smooth it approaches this limit asymptotically.

Fig. 2 displays the graph of this policy function for wealth levels ranging from 10 to 2000, which seems sufficient for characterizing behavior with $x = 3$. The solution is computed[8] under the assumption of the parameter values[9] displayed in Table 1. As expected, the optimal equity share starts out as the same level as the one with the highly risk averse preference ordering CRRA$_1$ at very low wealth levels. From there, it increases monotonically and continuously with the approximate shape of a sigmoidal curve and approaches the 75% level of the less risk averse preference ordering CRRA$_2$ when wealth is high enough for the probability to be very low of consumption falling below the norm. The rise from the bottom level starts at wealth levels somewhat below the level w^x where optimal consumption just equals the social norm. The optimal equity share rises rather steeply around this level, reflecting the diminishing probability of having to restrict consumption below the social norm as wealth increases.

The procyclicality[10] of risk taking is an important feature of our model. It may cause the agent to want to sell equity after prices drop because the loss of wealth makes the agent more risk averse. Similarly, the agent may want to "buy at the top" because risk aversion is lower in good times.

[8] For ease of computation, this solution is derived as a direct numerical solution of the non-linear ODE (14) rather than the more complicated formulae of Section 3.1. We transformed the second-order ODE into a two-equation first-order ODE system, which we solved over the range $[w_{min}, w_{max}]$, where $w_{min} = 10$ and $w_{max} = 4000$. For given guesses of the initial equity share and consumption-growth ratio, $\omega(w_{min}, x)$ and $\eta(w_{min}, x)$, we solved this system forward using the solution software for such problems in DifferentialEquations.jl, a package developed in the Julia programming language by Rackauckas and Nie (2017). Our initial guesses for $\omega(w_{min}, x)$ and $\eta(w_{min}, x)$ were derived from the corresponding values for the highly risk averse preference ordering CRRA$_1$. As our convergence criterion, we required the value of $V_w(w_{max}, x)$ to be close to the corresponding value for the preference ordering CRRA$_2$ with low risk aversion. By trial and error, we found the equity share for CRRA$_1$ of 25% to be a highly robust guess for the solution value $\omega(w_{min}, x)$, which allowed us to limit the search for initial guesses to one dimension, the consumption-wealth ratio. We were then able to solve the model as a saddle-point problem.

[9] Because our model is partial, it contains no productivity-driven natural growth rate that can be used to anchor the growth rate g of the social norm. However, as optimal consumption in our model consistently stays below the expected portfolio return, wealth keeps rising in expectation, which makes consumption grow as well. We then calibrated the norm growth rate g at 1.9% so as to roughly match the expected growth rate of consumption and wealth in the model.

[10] Cyclicality in this context is not related to business cycles in the macroeconomic sense, but to the fluctuations of financial returns.

4.2. Policy function for consumption

The monotonicity of the policy function for consumption follows directly from its first-order condition, as stated in Eq. (12) and Eq. (18). It is also continuous everywhere; however, its derivative makes a jump at the point where $c^* = x$. Although this point follows from Eq. (18), a more informative route is to differentiate both sides of Eq. (12) with respect to w and solve for the derivative of optimal consumption with respect to wealth. On elasticity form, it becomes

$$\frac{\partial \ln c^*}{\partial \ln w} = \frac{V_{ww} w}{u_{cc} c^*}.$$

Although V_{ww} is continuous, u_{cc} makes a jump at the point where consumption equals the norm. We then see that the consumption policy function must have a kink at wealth w^x where $c^* = x$. To explore this issue further, note first from Eq. (13) that

$$V_{ww} w = -\left(\frac{V_w}{\omega^*}\right)\left(\frac{\pi}{\sigma^2}\right).$$

Next, differentiating the utility function Eq. (6) twice, we find

$$u_c = -c^{-\gamma_i} x^{\gamma_i - 1}, \;\; u_{cc} = -\gamma_i c^{-\gamma_i - 1} x^{\gamma_i - 1},$$

so that

$$u_{cc} c = -\gamma_i c^{-\gamma_i} x^{\gamma_i - 1} = -\gamma_i u_c,$$

where, $i = 1, 2$ depending whether $c < x$ or $c \geq x$, respectively. Furthermore, substituting for γ_i from the optimal equity share for the preference ordering $\mathrm{CRRA_i}$, as given in Eq. (4), we find

$$u_{cc} c = -\left(\frac{u_c}{\omega_i}\right)\left(\frac{\pi}{\sigma^2}\right).$$

Noting from Eq. (12) that $u_c = V_w$ in equilibrium, we the then find

$$\frac{V_{ww} w}{u_{cc} c} = \frac{\omega_i}{\omega^*}.$$

From the monotonicity of ω^*, we thus have

$$\frac{\partial \ln [c^*(w, x)/w]}{\partial \ln w} = \begin{cases} \frac{\omega_1 - \omega^*}{\omega^*} < 0, & w < w^x \\ \textit{undefined}, & w = w^x \\ \frac{\omega_2 - \omega^*}{\omega^*} > 0, & w > w^x \end{cases} \tag{20}$$

For $w \to \infty$, $c^* \to \eta_2 w$ from below, for similar reasons that $\omega^* \to \omega_2$ from below. This result is reinforced by the observation that behavior according to the preference ordering $\mathrm{CRRA_i}$, i = 1, 2 results in a dynamic process for wealth described by the geometric stochastic motion

$$\frac{dw_t}{w_t} = (r + \omega_i \pi) dt + \omega_i \sigma dB_t - d\left(\frac{c_t}{w_t}\right) = (r + \omega_i \pi - \eta_i) dt + \omega_i \sigma dB_t,$$

where, η_i is defined as in Eq. (5). Thus, the optimal wealth process under $\mathrm{CRRA_i}$ has drift

$$r + \omega_i \pi - \eta_i = (1/\gamma_i)\left[r - \rho + (1/2)(1 + 1/\gamma_i)(\pi/\sigma)^2\right],$$

where, we have made use of Eq. (4) and Eq. (5). For the parameter values in Table 1, this formula implies a drift of 1.95% for i = 2, which is slightly higher than the growth rate g of 1.9% for the social norm. Thus, when wealth is very high relative to the social norm, consumption levels similar to the ones optimal under $\mathrm{CRRA_2}$ are sustainable in our case in the sense that consumption will not tend to outgrow wealth.

If wealth declines below w^x, the second inequality in Eq. (20) tells us that the agent will seek to smooth consumption somewhat by allowing the consumption decline to be slower than the one for wealth. In fact, consumption will then be lower than what would have been optimal under $\mathrm{CRRA_1}$. To see this, note that, for our parameter values, the implied drift for wealth when $i = 1$ is only 0.45%. Thus, consumption as if with $\mathrm{CRRA_1}$ would make wealth tend to shrink relative to the social norm x_t, which would not be optimal because the agent has the option of breaking out of this vicious circle by temporarily consuming less than what would be optimal under $\mathrm{CRRA_1}$. That would indeed be optimal for our agent.

These considerations help explain the graph for the policy function for consumption shown as the blue curve in Fig. 3. For easier reading, this graph stops at $w = 500$. The green and red lines show, for comparison, the corresponding policy functions for $\mathrm{CRRA_1}$ and $\mathrm{CRRA_2}$, respectively. Whereas these are both linear, the policy function with the soft social norm consists of two slightly concave segments, joined together with a kink at $w = w^x$. Although the graph rises everywhere, it is flatter when consumption falls below the social norm. In fact, as wealth rises towards this level, the slope becomes even a little bit flatter until suddenly becoming steeper at $w = w^x$. From there on, it is steeper than both alternatives, crosses the graph for $\mathrm{CRRA_1}$, and ultimately approaches the one for $\mathrm{CRRA_2}$, albeit at much higher wealth levels than the ones shown in this graph.

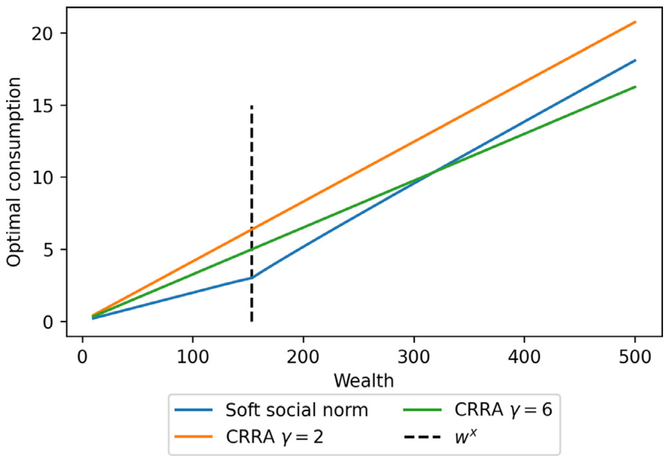

Fig. 3. Policy function for consumption.
The red and green lines show optimal consumption for the CRRA preference orderings with relative risk aversion of 2 and 6, respectively. The dashed vertical line denotes the wealth level w^x at which the optimal consumption equals the social norm. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

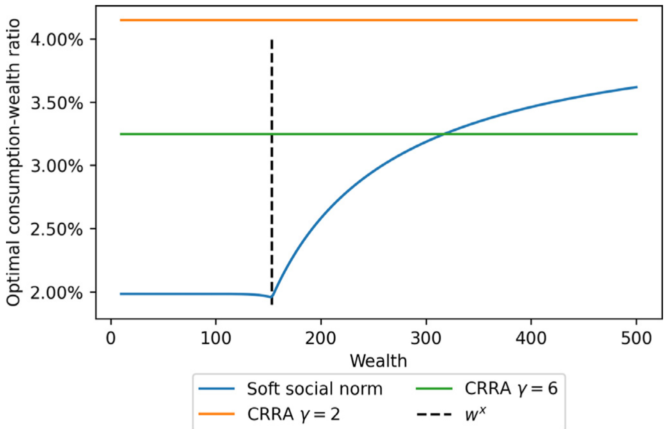

Fig. 4. Policy function for the consumption-wealth ratio.
The red and green lines show optimal consumption-wealth ratios for the CRRA preference orderings with relative risk aversion of 2 and 6, respectively. The dashed vertical line denotes the wealth level w^x. at which the optimal consumption equals the social norm. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Taken together, Figs. 2 and 3 reveal how our agent smooths consumption when wealth falls towards and/or below w^x. The main mechanism is that risk then is taken down, so that wealth fluctuates less. Some additional smoothing is obtained as the marginal propensity to consume (MPC) out of wealth falls abruptly when wealth drops below w^x. However, for wealth just above that level, the MPC is actually at its highest. So, when wealth falls to levels close to, but just above w^x, consumption is cut more than in proportion to the decline in wealth. The resulting rise in saving then helps avoid a further drop below w^x.

Fig. 4 repeats the picture from Fig. 3, but now in the form of the consumption-wealth ratio. Here, the optimal ratios for CRRA$_1$ and CRRA$_2$ become horizontal lines. The latter serves as an upper limit, toward which the policy function converges as wealth grows very large. However, for $w < w^x$, the policy function lies well below the level that is optimal for CRRA$_1$ preferences and falls even a little bit further before optimal consumption reaches the social norm. Then, it rises sharply before levelling out gradually. This is a case of dynamic substitution that seems rather unique to our specification. The motivation for consuming less than with CRRA$_1$ preferences at low wealth levels comes from the prospect of reaching beyond the norm if the agent can succeed in building sufficient wealth.

To make this dynamic substitution clearer, consider the s-period marginal rate of intertemporal substitution

$$M_s = e^{-\rho s}\frac{u_c(c_{t+s}, x_{t+s})}{u_c(c_t, x_t)}$$

for pairs of consumption/norm ratios such that $c_{t+s}/x_{t+s} > c_t/x_t$. Let M_s^1 denote the corresponding factor for an individual with CRRA$_1$ preferences globally. Then, clearly, if $c_{t+s}/x_{t+s} < 1$, $M_s = M_s^1$. For $c_{t+s}/x_{t+s} \geq 1 > c_t/x_t$, it is easily verified that

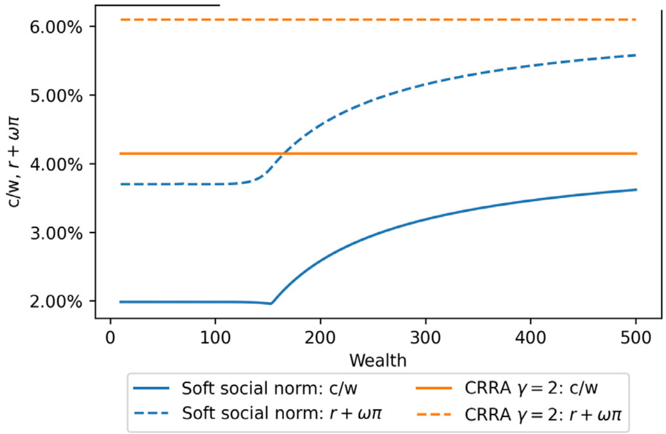

Fig. 5. Closeup of consumption-wealth ratio and expected return around w^x revealing strong saving motive.

$M_s = M_s^1 (c_{t+s}/x_{t+s})^{\gamma_1-\gamma_2} \geq M_s^1$. And for $c_{t+s}/x_{t+s} > c_t/x_t \geq 1$,

$$M_s = M_s^1 \left(\frac{c_{t+s}/x_{t+s}}{c_t/x_t} \right)^{\gamma_1-\gamma_2} > M_s^1.$$

This means that our agent discounts the marginal utility of a brighter future less severely than the CRRA_1 agent and is thus willing to sacrifice more to potentially obtain it. This saving motive is reinforced by the expectation that the norm is expected to grow over time at rate g. The latter point became clear in experiments that we did with lower growth rates, which resulted in significantly higher consumption-wealth ratios at low wealth levels. With a faster growing social norm, it takes more saving to help consumption catch up, so that the incentive implicit in the stochastic discount factor applies to a greater amount of saving.

Fig. 5 illustrates this dynamic substitution more clearly by displaying the optimal consumption-wealth ratio (solid blue) together with the expected portfolio return (dashed blue) at each wealth level, where the latter, $r + \omega^*\pi$, varies with the equity share. The solid and dashed red lines show the corresponding consumption-wealth ratio and expected portfolio return, respectively, for the CRRA_2 preferences, which we think of as "no social norm." The distance between the expected portfolio return and the consumption-wealth ratio can be interpreted as an expected saving rate out of wealth. Starting from the lowest wealth level, the graph in Fig. 5 then shows this saving rate rising with wealth as the consumption-wealth ratio falls even as the expected return rises with the equity share, reflecting the correspondingly improved chance of climbing out of below-norm consumption levels. After crossing the barrier w^x it continues to rise. In fact, it rises above the corresponding rate for the case of no social norm even though the rate of portfolio return is much lower because of the low equity share. It then levels off and asymptotically approaches the saving rate of case with no social norm.

5. Behavior over time

We simulate the model solution over 50 years to study the dynamic development of wealth, risk taking, and consumption. The step size is specified as 12 per year, corresponding to months, so that the simulation covers 600 months. We draw 50,000 duplications and start with an initial wealth level implying consumption at 10% above the social norm.

Fig. 6 presents a fan chart for the distribution of future wealth relative to the social norm. It is increasingly skewed to the right, consistent with the well-known result that normally distributed instantaneous returns imply log-normal distributions for future wealth. However, the mean is pretty much horizontal, indicating a central tendency for consumption to grow at the same rate as the social norm. This result is not necessarily implied by the model, however. Although the consumption growth rate implied by the agent's choices depends on the growth rate of the social norm, this relationship is not one to one, at least not at all dates. However, as noted in Section 4.1, we have, by trial and error, calibrated the growth rate of the social norm so as to roughly match the average consumption growth.

Fig. 7 presents a similar fan chart for the equity share, using the same draws. It is highly skewed to the right, which is natural given the lower limit given by ω_1, which is about 25%. Even so, the mean is fairly stable across time horizons, which fits with our notion of consumption above the social norm as the normal case. However, because the norm grows over time, the equity share tends to remain well below than the 75% upper limit in our model.

Fig. 8 shows the corresponding pattern for consumption relative to norm. Because the MPC is lower for $w < w^x$ than for $w > w^x$ the skewness of the distributions is even more pronounced than for wealth. The kink at this point is reflected as a kind of corner indicated by an arrow. The lower MPC below this wealth level limits the distribution's downside, which makes the mean rise slightly with the horizon.

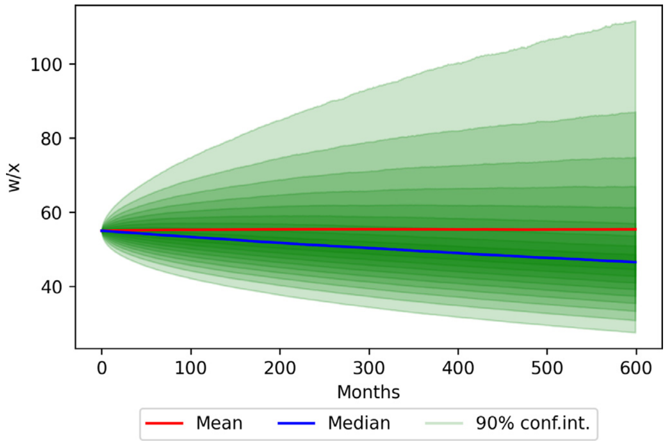

Fig. 6. Simulated development of wealth relative to the social norm.

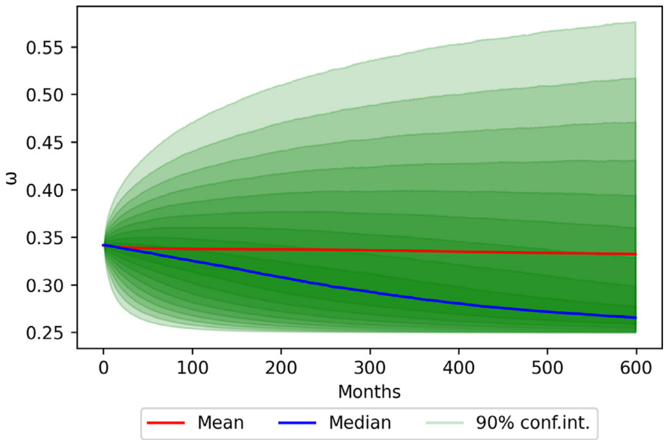

Fig. 7. Simulated development of equity share.

Taken together, Figs. 5, 7, and 8 illustrate the efforts that our agent undertakes to curtail the probability of having to cut consumption below the social norm. Although the initial wealth in our example is high enough to maintain the consumption-wealth ratio that a CRRA_1 agent would choose, our agent chooses to save more, thus making wealth grow faster than the social norm, which—in combination with low risk taking—puts a lid on the probability of financial outcomes that are sufficiently negative to call for below-norm consumption.

The incentive to stay away from the region where $c < x$ is naturally greater the higher the curvature of the utility function in that region. In the limiting case where $\gamma_1 \to \infty$, this incentive becomes absolute in the sense that utility falls to minus infinity in that region. For a solution to exist in this case, initial wealth must be large enough to safely support and sustain consumption at or above the ever-growing social norm, which requires

$$w_0 \geq \frac{x_0}{r-g} \tag{21}$$

When this inequality binds, the agent will be constrained to always just consuming the social norm with all wealth invested in the riskless asset because otherwise the case of $c - x < 0$ will have a positive probability, which cannot be optimal[11]. Moving beyond this point will be impossible because the agent's income cannot increase when equity investment is kept at zero, which is necessary to avoid the risk of consumption falling below the norm. If the inequality is strict, however, the agent will be able to take some risk and raise saving and thus be able to maintain consumption higher than the social norm. The feasibility constraint in Eq. (21) resembles those found in most habit models, as well as in other models with some kind of lower bound on consumption. The important difference that, in our model, it applies only in the limiting case where $\gamma_1 \to \infty$. A similar bound would apply in the model of Choi et al. (2022) if $\beta \to \infty$. However, in contrast to their model, where such a constraint could apply at any time, it can be relevant only for our agent's initial wealth level.

[11] Clearly, $r > g$ is needed for (21) to be valid. This requirement is also needed in a typical habit model as well as in other models with lower bounds on consumption. Whether or not it is needed in our model depends on the values of the other parameters.

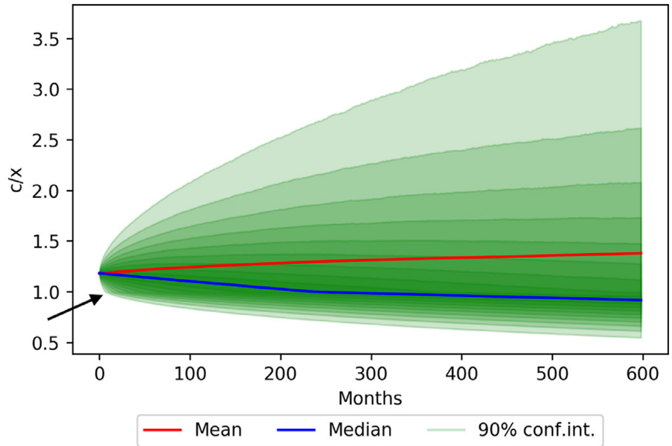

Fig. 8. Simulated development of consumption relative to the social norm.

6. Discussion

This section discusses the implications of our model for the practice of portfolio management and spending policies of endowment funds and sovereign wealth funds. We start by comparing the implications of our model with alternative models in the literature.

6.1. Comparison with alternative models

Of the alternative models in the literature, ours comes close to the intertemporal loss aversion (ILA) models of Shrikhande (1997), Watson and Scott (2015), and Choi et al. (2022) in terms of assuming a finite penalty for consumption below some norm rather than deeming it infeasible in the sense of generating infinitely negative utility. However, the implications in terms of behavior of consumption and risk taking are quite different. For persisting (or permanent) wealth changes, the ILA models predict no lasting changes in risk taking or consumption relative to wealth, whereas our model predicts procyclical risk taking and a consumption-wealth ratio that rises with wealth if the optimal consumption exceeds the social norm. However, if wealth declines far enough for optimal consumption to fall short of the social norm, the consumption-wealth ratio will start rising again as the agent seeks to limit the extent of below-norm consumption. For local wealth changes, the predictions of the ILA model depend on how close the agent is to wanting to change consumption, an issue that does not come up in our model.

In terms of behavioral predictions, our model comes closer to the habit-formation models with instantaneous utility functions like the one in Eq. (7). The policy functions for such an agent have been derived by Constantinides (1990) and were further explored by Lindset and Mork (2019). Dybvig's (1995) model of non-decreasing consumption can be interpreted as a habit-formation model where the habit is updated instantaneously to match current consumption. Like ours, habit models imply procyclical risk taking; however, they are more extreme by requiring the agent to forgo all risk taking if that is needed for the riskless part of the portfolio to be able to fund the habit level of consumption. Our specification is softer by implying a positive lower bound on risk taking. In terms of consumption, both models predict some smoothing relative to the CRRA benchmark. In terms of welfare, however, the habit models predict that any gain in happiness caused by an increase in wealth eventually will dissipate as people get used to the more luxurious living. Our model implies a similar prediction as the social norm is assumed to rise steadily over time; however, our model views this movement as exogenous and not a result of the agent's actions.

Habit models have an advantage over ours in that the decline in curvature as wealth grows there is continuous rather than the result of a jump. However, because the curvature grows without limit as consumption falls towards the habit level, habit models function best if the habit can be understood as a fairly low fraction of actual consumption. Furthermore, and importantly, they imply an absolute requirement that consumption always be kept above the habit level. This requirement is satisfied endogenously because their agent in good times sets aside a sufficient part of wealth in the riskless asset to ensure sufficient funds to be available to defray the cost of maintaining the habit in bad times.

However, an agent wanting to initiate behavior based on such preferences needs to have sufficient initial wealth to allow a large enough set-aside in safe investments to satisfy the initial habit. In our model, this requirement arises only in the limiting case of $\gamma_1 \to \infty$, as discussed above. Accordingly, the owners of an endowment fund or a sovereign wealth fund that consider switching to a strategy compatible with habit preferences will need to set aside sufficient funds in safe assets to maintain what they consider habitual, i.e. minimal spending. As mentioned in the introduction, this requirement may conflict, sometimes seriously, with preconceived ideas about minimum spending. We furthermore find it intuitively appealing to allow consumption to fall below the social norm in bad times because the restraint on consumption in good times may seem unduly harsh. Thus, our specification allows for more flexibility both initially and over time.

Habit-formation models have been advanced as explanations of the equity premium puzzle (e.g. Constantinides 1990, Campbell and Cochrane 1999) and the observed smoothness of aggregate consumption. Because our model is partial and not necessarily intended to describe the behavior of all agents in a general-equilibrium model, we make no claims regarding the equity premium puzzle.

Nor can we claim that our model offers a good explanation for the smoothness of aggregate consumption. It is true that consumption is smoothed slightly relative to wealth as illustrated by the decreasing consumption-wealth ratio for $w < w^x$ in Figs. 3 and 4. However, the sharply rising ratio for $w > w^x$ pulls very much in the opposite direction. The standard deviation (in annual terms) of the log change in consumption in our model[12] is 9.1% for the parameter values in Table 1. That is lower than the 13.4% implied by the low-curvature specification $CRRA_2$, mainly because wealth fluctuates less for our more risk-averse agent. However, it is significantly higher than the 4.5%[13] implied by the high-curvature specification $CRRA_1$.

6.2. Implications for endowment and sovereign wealth funds

The analysis of our specification is intended as a contribution to the ongoing discussion about portfolio and withdrawal rules for endowment funds and sovereign wealth funds such as the Norwegian GPFG. Because one of our concerns is that portfolio and withdrawal rules must be consistent with each other, we emphasize the implications of our analysis along both dimensions. Although all such funds typically have rules about withdrawals, these rules may be bypassed in practice when it comes to actual spending as short-term borrowing may be used to handle fluctuating spending needs. We consider it worth pointing out that such deviations in fact have portfolio implications in that short-term borrowing is equivalent to a lowering of the riskless share of the portfolio. Moreover, the GPFG is, by statute, barred from making such bypasses in that the Government Pension Fund Act of 1990 prohibits loan-financed spending as long as there is money left in the fund[14]. The motivation for this rule is to prevent policy makers from using borrowing to circumvent the withdrawal rules in general. This legal detail makes the practical implementation of withdrawal rules so much more important for the GPFG.

The main GPFG withdrawal rule, referred to as the Fiscal Rule, passed by the Storting (Parliament) in 2001[15], is to allow payouts in support of the government's annual budget corresponding to the expected real return on the fund's balance at the beginning of the year, currently stipulated as 3%. If strictly adhered to, this rule would require spending to follow the ups and downs of financial markets, independently of perceived spending needs. However, the Fiscal Rule also includes provisions to allow for smoothing, which may be approximately represented as the MIT-Tobin rule[16]

$$c_t = \lambda \bar{\eta} w_t + (1 - \lambda) c_{t-1}, \tag{22}$$

where, $\bar{\eta}$ is the target withdrawal rate (3% for the GPFG).

Our analysis implies a number of recommendations for changes in this rule. First, the target ratio $\bar{\eta}$, the expected real financial return, is much too high. It is not sustainable, as shown by Dybvig and Qin (2021). Although our model does not recommend a fixed target ratio, the range of recommended ratios displayed in Fig. 4 is consistently lower than the one implied by the low-curvature preference ordering $CRRA_2$, which in turn is significantly lower than the expected rate of portfolio return.

Second, although our model implies some smoothing of payouts at wealth levels below the social norm, the pattern above this wealth level is very much the opposite. In particular, financial losses should call for spending cuts that are more than proportional to the financial loss. This result follows from a desire to save as a safeguard against the risk of future spending below the social norm. The message from our analysis is that the current rule does not take this risk sufficiently seriously. The widespread use of withdrawal rules based on smoothing rules such as the one in Eq. (22) may perhaps be explained by beliefs in mean-reverting equity prices. Cochrane's (2022) discussion may give some support to this belief. However, mean reversion also means that long-term rates of return will be lower than the short-term ones, as analyzed by Mork and Trønnes (2023).

Third, managing the risk of payouts falling below the social norm implies significant limitations on the share of risky assets in the portfolio. However, the equity share need not, and should not be constant. By being procyclical it should allow

[12] The simulated results are averages of the time-series estimated standard deviations for each of the duplications.

[13] This number is almost exactly like the one obtained from quarterly U.S. data on per-capita consumption since 1947. That does not mean that our $CRRA_1$ specification explains U.S. aggregate data, however, because our model is very much partial whereas a test of the equity premium needs a general-equilibrium model. In particular, with a representative agent, it needs a model with zero net holdings of the risk-free asset, so that the equity share is constrained to unity. With CRRA preferences, that constraint would make the volatility of consumption equal to that of equity wealth, which we have calibrated at 17.9%.

[14] https://www.regjeringen.no/contentassets/9d68c55c272c41e99f0bf45d24397d8c/government-pension-fund-act-01.01.2020.pdf, accessed October 5, 2022. All borrowing done by the Norwegian government after the 1990 passing of this Act is for the purpose of funding policy financial institutions, the largest of which is the *Lånekassen* student loan agency.

[15] https://www.regjeringen.no/no/dokumenter/stmeld-nr-29-2000-2001-/id194346/, accessed November 2, 2022.

[16] The Norwegian Fical Rule is actually more complicated, on the one hand because the rule applies to the structural rather than the actual non-oil deficit and on the other hand because it allows for discretionary fiscal policy as temporary deviations from the rule. However, all such deviations must be funded directly from the GPFG. These details are further discussed and analyzed in Mork et al. (2022).

Table 2
Comparison with rules of thumb Loss of expected utility expressed in units of pct loss of w_0.

Constrained optimal constant $\omega = 28.9\%$ and $c/w = 2.1\%$	0.3%
Constrained optimal values updated every 10 years	0.1%
Optimal initial values held constant, then updated every 10 years	1.8%
CRRA with constrained optimal ω	21.2%
CRRA with constrained optimal ω, smoothed, $\lambda = 0.5$ (annual)	21.2%

higher risk taking when wealth is high enough for the probability of falling below the social norm is modest, and vice versa. High financial returns should be followed by increases in the equity share and vice versa. That means that our agent sometimes may want to "buy at the top" and "sell at the bottom." Again, the motivation "at the bottom" is to limit the risk of falling and/or staying below the social norm. But similarly, because this risk is modest "at the top," it will then be prudent to harvest more of the equity premium.

6.3. Mandate and implementation

In principle, our investor could delegate all investment decisions to portfolio managers by communicating the policy functions of the optimal solution. In practice, simpler rules are preferred because they are easier to communicate, and they make monitoring of manager performance easier. This subsection offers an evaluation of the expected-utility cost of formulating mandates as simpler rules of thumb rather than the optimal solution.

We determine this cost by simulating each of the alternatives over 100 years[17] and comparing their implied expected utility, computed as the mean subjectively discounted utility, based on the soft-social-norm preferences. We use the same 50,000 duplications in all alternatives so as to avoid differences among alternatives being caused by sampling errors. We start with consumption at 10% above the social norm, so that consumption falling below the norm is not an immediate threat, but an event close enough to guard against. For each alternative, we then divide the difference in expected utility from the optimal case by marginal utility[18] so as to express the difference as a loss measured in consumption units. Lastly, by dividing this loss by the marginal value of the initial level of wealth, we express it as the equivalent to a percentage loss of initial wealth.

As a first alternative, we consider a mandate where the equity share and the consumption-wealth ratio are held constant over time, but where the constant values are derived as optimal subject to this constraint. We determine the constrained-optimal constants as the pair of values that provide the highest simulated expected discounted utility over 100 years.

The result is presented as the first line in Table 2. The estimated implied loss is small, in fact, so small that many would call it trivial[19]. It becomes even smaller if we assume that the values for equity share and the consumption-wealth ratio are updated by recomputing the same procedure for each simulated scenario every 10 years[20]. It remains less than two percent of the initial wealth if instead we assume that the fund owner requires managers to use values that are optimal on the initial date and, for reasons of convenience, keeps it constant for 10 years at a time.

These results, which are reported in the first three lines of Table 2, may seem astonishing considering the considerable range of optimal values displayed in Figs. 2 and 4. The explanation is that the optimal solution chooses a combination of equity share and consumption-wealth ratio that makes consumption remain fairly stable relative to the social norm as the latter changes over time.

The constrained-optimal choice of equity share and consumption-wealth ratio thus depends on the starting point. If, instead of starting with consumption 10% above the norm, the agent starts with a wealth level implying consumption 10% *below* that norm, the constrained-optimal equity share and the consumption-wealth ratio are both lower than with the higher initial wealth in Table 2. For much higher wealth levels, these values will naturally converge toward the optimal values for the unconstrained specification $CRRA_2$, which are 75.0% and 4.1%, respectively.

The near optimality of constant policy functions may give the impression that our preference ordering can be closely approximated by a CRRA function or Epstein–Zin preferences with constant risk parameter and elasticity of intertemporal substitution. That is not the case, however, as can be readily seen from the fact that the constrained-optimal combination of equity share $\omega = 28.9$ % and consumption-wealth ratio of $c/w = 2.1$ % could not have been implied by any of the two alternatives. For expected CRRA preferences and our values for equity variability and equity premium, an equity share of 28.9% would have required a curvature parameter of $\gamma = 5.2$, which in turn would have implied a consumption-wealth ratio as high as $c/w = 3.3$ %. A higher curvature parameter would have implied a lower consumption-wealth ratio, but also a

[17] Our reason for now extending the simulation to 100 years is that the discount factor for this horizon becomes very close to zero. With $\rho = 0.025$, $e^{-100\rho} = 0.018$.

[18] For alternatives whose results are reasonably close to the unconstrained optimal solution, we use the derivative of the value function at the starting wealth level. For alternatives with significantly inferior results (in practice, the last two lines of Table 2), we use the average of the same derivative at the starting wealth level and the one yielding the same value of the value function as the expected discounted utility of that alternative.

[19] For a fund worth USD 1.3 trillion, 0.3% is still USD 3.9 billion. We leave it to the reader to judge whether that is trivial.

[20] Updates with approximate 10-year intervals have been the practice so far of the GPFG.

lower equity share. In fact, for expected CRRA, the consumption-wealth ratio cannot fall below the risk-free rate, which in our case equals 2.5%. With Epstein–Zin preferences, the corresponding lower bound is defined as the smaller of the riskless rate and the subjective discount rate. With the latter calibrated at 4%, the lower bound is again 2.5%.

Thus, with our parameter values, the constrained-optimal values for the consumption-wealth ratio and the equity share lie outside the possible combinations for expected CRRA utility and Epstein–Zin preferences. Although this may not be the case for other sets of parameter values or other starting points for wealth, our example clearly shows that the constrained-optimal policy functions for soft social norms can be quite different from these frequently used preference specifications. This is another example of the essential role that dynamic substitution plays in our specification. High saving helps build wealth despite low equity income and thus serves as a relative safeguard against suffering consumption below the social norm.

We considered including an alternative based on the rule of equating consumption to the expected financial return, with or without smoothing, and with varying degrees of risk taking. Unfortunately, such rules fit poorly into our framework because we have ignored non-financial income[21] that would typically grow over time to match the trend growth of the social norm. Because consumption, when set to equal the expected financial return, would then not grow in expectation, there would be a steady tendency for consumption to fall behind the exogenously determined social norm over time. As a result, we would exaggerate the cost of following such rules because they look worse in our partial framework than they would do in a more realistic, but also more complicated framework with non-financial income.

As a substitute, we consider an alternative based on the behavior of a person with CRRA expected-utility preferences with risk aversion implying the same equity share as in the constrained-optimal case of the first line in Table 2. With our parameter values, that requires a curvature parameter of $\gamma = 5.2$ which, as noted above, would imply a consumption-wealth ratio of 3.3%, about 50% higher than in the constrained optimum. The welfare loss from following that behavioral pattern, shown in the fourth line of Table 2, corresponds to a more than 20% loss of initial wealth, two orders of magnitude greater than the constrained optimum.

As a final alternative, we consider a spending rule like the one in Eq. (22), but where $\bar{\eta}$ is the consumption-wealth ratio of the CRRA preferences with $\gamma = 5.2$, and $\lambda = 0.056$ per month, implying an annual rate of adjustment[22] of about 0.5. The result is shown in the last line of Table 2.

Clearly, the smoothing does not help. In fact, it makes the result slightly worse, albeit only in the second decimal of the percentage, which is not shown. The benefits of smoothing are small as long as consumption stays above the social norm; and smoothing has the negative effect of widening the distribution of future wealth levels. The latter effect is almost negligible in our case, however, making the net effect of smoothing essentially nil.

7. Summary and conclusions

Inspired by debates about investment and withdrawal strategies for endowment funds and sovereign wealth funds, this paper has explored the implications of a CRRA preference ordering modified by a soft social norm as a possible representation of investor preferences. Although we find this representation appealing, our purpose has not necessarily been to promote it as normatively better than alternatives, but to explore how the behavior implied by this model differs from that of other models. Even so, we believe this exploration could be useful in terms of guiding the concrete decisions about risk taking and withdrawal policy that will have to be made by investors whose preferences are influenced by factors similar to those we consider.

Our model shares one important element with the loss-aversion literature, namely that the penalty from letting consumption fall below the norm is finite rather than infinite, like it is in the typical habit-formation models. In other respects, however, our model behaves quite differently. In particular, it shares the two leading features of habit models and other models with lower bounds on consumption, namely, smoother consumption than with CRRA preferences and high and countercyclical risk aversion, albeit lower risk aversion than the typical habit model. The main difference is that our agent has a much higher propensity to save in bad times, being quite willing to substitute consumption dynamically by saving a lot when wealth is low in the hope of building sufficient wealth to allow consumption above the social norm in the future. Models of habits and other lower bounds on consumption do not capture this effect because there, consumption above the bound is secured by sufficient holding of safe assets, which in turn requires a severe constraint on consumption in good times. By allowing higher consumption in good times as well as somewhat higher risk taking in bad times, we believe our model may more closely resemble the preferences revealed by the actions observed for large endowment funds and sovereign wealth funds.

That said, our analysis suggests a number of changes that the owners of such funds may want to implement if they share our belief in the importance of social spending norms. Annual spending should be significantly lower than expected financial returns. Spending as a share of wealth should furthermore be procyclical; in particular, financial losses should, as a rule, be followed by larger than proportional spending cuts so as to build enough new wealth to limit the probability of

[21] For individual investors, this would typically be labor income. For endowment funds it would be current revenue such as tuition payments, and for sovereign wealth funds current tax revenues.

[22] 0.056 is the approximate solution to the equation $(1 - \lambda_m)^{12} = 1 - \lambda_a$, where $\lambda_a = 0.5$.

consumption falling below the norm. The only exception would be to keep spending from falling too far below the social norm.

The desire to keep consumption above the social norm carries important implications for financial risk taking. It should be modest, although not as modest as in the typical habit-formation models. It should be procyclical, again not to the same degree as with habit formation, but enough so that the fund's managers sometimes may want to "buy at the top" and "sell at the bottom."

Because the policy functions implied by our model are highly non-linear, delegating them to professional managers may not be simple. However, at least with our parameter values and other assumptions, we find that the welfare loss would be quite modest if mandates for equity share and withdrawals are updated only every 10 years, even if the mandates are set myopically at the values that would have been currently optimal if updates had been made continuously. However, this mandate would depend on the wealth owner's initial wealth and look very differently from one based on expected CRRA utility or Epstein–Zin preferences.

The exogeneity of the social norm may be considered a limitation of our model. A promising modification might be to replace this soft norm by an equally soft habit, where the habit is built up from past consumption rather than exogenously. The resulting model would still differ quite a bit from the habit-formation literature by treating the habit as a soft rather than a hard bound on consumption. We intend to turn to this task in a follow-up paper.

Appendix 1. Analytical solutions to the Merton problem through the dual value formulation

In this Appendix we assume that the soft social norm x is a fixed number, and not dynamic. The case of a dynamic social norm is left for Appendix 2. Whereas the proposed utility in Eq. (6) behaves very differently than the standard CRRA utility function, we are able to derive an analytic solution by using the dual value function approach to solve the HJB equation associated to the value function.

To this end, we will follow the procedure from Rogers (2013) Section 1.3. First, it is completely clear from the time-independence of the utility function in Eq. (6) that it satisfies a standard transversality condition, i.e. that $\lim_{t\to\infty} e^{-\rho t}u_c(c,x) = 0$. The dual value function approach is based on a generalized Legendre transform, and for this it is convenient with a change of variables. Define the new coordinates

$$(t,z) = (t, V_w(t,w))$$

where, $(t,z) \in A := (0,\infty) \times (V_w(t,\infty), V_w(t,0))$. Define a new function $J : A \to \mathbb{R}$ by

$$J(t,z) = V(t,w) - wz.$$

Note that the variable z now implicitly describes the optimal control, since by first order conditions the marginal utility of the optimal consumption c^* is equal to $V_w(t,w)$. We define $\tilde{u}(z,x)$ to be the convex dual of $u(c,x)$, that is

$$\tilde{u}(z,x) = \sup_c \{u(c,x) - zc\}.$$

This is a decreasing function in $z \mapsto \tilde{u}(z,x)$, and it is readily seen that $\tilde{u}(z,x) = u((u_c)^{-1}(z,x),x) - z(u_c)^{-1}(z,x)$. It is readily checked that the inverse of the $(u_c)^{-1}(z,x)$ is given by

$$(u_c)^{-1}(z,x) = \mathbb{1}_{z^{-1/\gamma_1}x^{1-1/\gamma_1}<x} z^{-1/\gamma_1}x^{1-1/\gamma_1} + \mathbb{1}_{z^{-1/\gamma_2}x^{1-1/\gamma_2}\geq x} z^{-1/\gamma_2}x^{1-1/\gamma_2}.$$

Inserting the inverse function into u we can compute that

$$u\big((u_c)^{-1}(z,x),x\big) = \mathbb{1}_{z^{-1/\gamma_1}x^{1-1/\gamma_1}<x}\frac{1}{1-\gamma_1}\Big[(zx)^{1-1/\gamma_1}-1\Big] + \mathbb{1}_{z^{-1/\gamma_2}x^{1-1/\gamma_2}\geq x}\frac{1}{1-\gamma_2}\Big[(zx)^{1-1/\gamma_2}-1\Big].$$

And so using this representation and subtracting $z(u_c)^{-1}(z,x)$ we find that

$$\tilde{u}(z,x) = \mathbb{1}_{z^{-1}<x}(1-\gamma_1)^{-1}\Big[\gamma_1(zx)^{1-1/\gamma_1}-1\Big] + \mathbb{1}_{z^{-1}\geq x}(1-\gamma_2)^{-1}\Big[\gamma_2(zx)^{1-1/\gamma_2}-1\Big].$$

By time homogeneity of the payoff function we know that $V(t,w) = e^{-\rho t}v(w)$, and therefore, the convex dual satisfies $J(t,z) = e^{-\rho t}j(ze^{\rho t})$. From Section 1.3 in Rogers (2013) it then follows that the convex dual is given by the following conditional expectation

$$e^{\rho t}J(t,z) = E\left[\int_0^\infty e^{-\rho s}\tilde{u}(Y_s,x)\,ds | Y_0 = e^{\rho t}z\right].$$

Here, Y is a geometric Brownian motion satisfying

$$dY_t = Y_t\left(\left|\frac{\mu-r}{\sigma}\right|dW_t + (\rho - r)dt\right) = Y_t((\pi/\sigma)dW_t + (\rho-r)dt),$$

where, the last equality follows from the definition of the equity premium π and the fact that π and σ both must be positive for our problem to be well posed.

Observe that by the Fubini theorem, we can take expectations inside the integral, and we are left to compute

$$E\Big[\tilde{u}(Y_s, x)\,|Y_0 = e^{\rho t}z\Big] = E\Big[\mathbb{1}_{Y_s^{-1}<x}(1-\gamma_1)^{-1}\Big[(Y_s x)^{1-1/\gamma_1} - 1\Big]|Y_0 = e^{\rho t}z\Big]$$
$$+ E\Big[\mathbb{1}_{Y_s^{-1}\geq x}(1-\gamma_2)^{-1}\Big[(Y_s x)^{1-1/\gamma_2} - 1\Big]|Y_0 = e^{\rho t}z\Big] = J^1(s,t,z,x) + J^2(s,t,z,x)$$

Manipulating further, using the explicit solution to the geometric Brownian motion, we see that

$$J^1(s,t,z,x) = e^{(1-1/\gamma_1)\rho t}(1-\gamma_1)^{-1}(zx)^{1-1/\gamma_1}E\Big[\mathbb{1}_{Y_s^{-1}<x}(Y_s)^{1-1/\gamma_1}|Y_0 = 1\Big] - (1-\gamma_1)^{-1}E\Big[\mathbb{1}_{Y_s^{-1}<x}|Y_0 = e^{\rho t}z\Big]$$
$$= e^{(1-1/\gamma_1)\rho t}(zx)^{1-1/\gamma_1}N_1(s,x) - (1-\gamma_1)^{-1}P\Big(Y_s^{-1} < x|Y_0 = e^{\rho t}z\Big).$$

Here, $N_1(s,x)$ represents the conditional expectation of a power of a geometric Brownian motion. The conditional probability $P(Y_s^{-1} < x|Y_0 = e^{\rho t}z)$ can further be simplified by a simple change of variable, using the fact that Y_s is log-normally distributed, and we get that

$$P\Big(Y_s^{-1} \leq x|Y_0 = e^{\rho t}z\Big) = e^{\rho t}zP\Big(Y_s^{-1} \leq x|Y_0 = 1\Big).$$

We therefore find that,

$$J^1(s,t,z,x) = e^{(1-1/\gamma_1)\rho t}\ (zx)^{1-1/\gamma_1}N_1(s,x) - e^{\rho t}z\ (1-\gamma_1)^{-1}P\Big(Y_s^{-1} \leq x|Y_0 = 1\Big).$$

A similar representation can be made for J^2. Note that when adding J^1 and J^2 in order to find J, we get

$$J(s,t,z,x) = e^{(1-1/\gamma_1)\rho t}(zx)^{1-1/\gamma_1}N_1(s,x) + e^{(1-1/\gamma_2)\rho t}(zx)^{1-1/\gamma_2}N_2(s,x) \ - \ e^{\rho t}z(1-\gamma_1)^{-1}P\Big(Y_s^{-1} < x|Y_0 = 1\Big)$$
$$- e^{\rho t}z(1-\gamma_2)^{-1}P\Big(Y_s^{-1} \geq x|Y_0 = 1\Big).$$

Our goal is now to explicitly compute the conditional expectations N_i for $i = 1, 2$. By Itô's formula it follows that for $i = 1, 2$, $Z_t^i = (Y_t)^{1-1/\gamma_i}$ is a geometric Brownian motion with drift $(1-1/\gamma_i)(\rho - \ r \ - (1/\gamma_i)(\pi/\sigma)^2) = \rho - \eta_i$, where η_i was defined in Eq. (5), and standard deviation $(1-1/\gamma_i)\pi/\sigma$. We therefore need to evaluate two conditional expectations of two log-normally distributed random variables

$$\sim Log\mathcal{N}\Big((\rho - \eta_i)t,\ \ (1-1/\gamma_i)^2(\pi/\sigma)^2 t\Big).$$

This expectation can be computed explicitly using the following formula: suppose $Y \sim Log\mathcal{N}(\bar{\mu}, \overline{\sigma^2})$, and let $\Theta_{\bar{\mu},\overline{\sigma^2}}$ denote the cumulative probability distribution associated to a normally distributed variable $\sim \mathcal{N}(\bar{\mu}, \overline{\sigma^2})$. Then for any $h > 0$ we have that

$$E[Y|Y > h] = e^{\bar{\mu}+\frac{\overline{\sigma^2}}{2}}\frac{1-\Theta_{\bar{\mu}+\overline{\sigma^2},\overline{\sigma^2}}(ln(h))}{1-\Theta_{\bar{\mu},\overline{\sigma^2}}(ln(h))}$$

and

$$E[Y|Y \leq \ h] = e^{\bar{\mu}+\frac{\overline{\sigma^2}}{2}}\frac{\Theta_{\bar{\mu}+\overline{\sigma^2},\overline{\sigma^2}}(ln(h)) - \Theta_{\bar{\mu},\overline{\sigma^2}}(ln(h))}{1-\Theta_{\bar{\mu},\overline{\sigma^2}}(ln(h))}.$$

This explicit formula can now be used directly to compute the terms N_1 and N_2 appearing in $J(s,t,z,x)$.

Summarizing our findings, computing the integral $J(t,z|x) = \int_0^\infty J(s,t,z,x)ds$, we can write this as the following formula

$$J(t,z|x) = F_1(t,x)(xz)^{1-1/\gamma_1} + F_2(t,x)(xz)^{1-1/\gamma_2} + zG(t,x),$$

where F_1 and F_2 are two functions capturing the weighted conditional expectation of the fractional moment of a geometric Brownian motion, and the function G captures the probabilities of the inverse of a geometric Brownian motion being above or below the log of the soft social norm x. This representation shows explicitly how $J(t,z|x)$ explicitly depends on z as a sum of powers of z.

To now find the value function, one may follow the procedure outlined in Section 3.

Appendix 2. Going from the analytic solution of the fixed social norm problem, to the dynamic social norm

Assume for now that we have found the value function $V(t,w,x)$ for a given (fixed) social norm level x. We are interested in finding the value function $V^g(t,w,x_0)$ describing the case where the social norm level is increasing at rate g, i.e. $x_t = x_0e^{gt}$. Due to the homogeneity of our proposed utility function, described in Eq. (6), it turns out that we can transfer between the two solutions by the simple transform $e^{-gt}c_t^g(w,x_0) = c_t(w,x_0)$, where c is the optimal consumption with fixed social norm, and c^g is the optimal consumption with increasing social norm at rate g. Indeed, since the utility function u is homogeneous of degree zero, we observe that

$$u(c_t, x_t) = u\Big(e^{-gt}c_t, x_0\Big).$$

And therefore,

$$V^g(t,w,x_0) = \max_{c^g \in C(R_+)} \mathbb{E}\left[\int_t^\infty e^{-\rho(s-t)} u\left(c_s^g, x_s\right) ds | W_t = w\right] = \max_{c^g \in C(R_+)} \mathbb{E}\left[\int_t^\infty e^{-\rho(s-t)} u\left(e^{-gs} c_s^g, x_0\right) ds | W_t = w\right]$$

$$= \max_{c \in C(R_+)} \mathbb{E}\left[\int_t^\infty e^{-\rho(s-t)} u(c_s, x_0) ds | W_t = w\right] = V(t,w,x_0).$$

References

Abel, A.B., 1990. Asset prices under habit formation and catching up with the joneses. Am. Econ. Rev. 80 (2, Papers and Proceedings), 38–42.
Alhashel, B., 2015. Sovereign wealth funds: a literature review. J. Econ. Bus. 78, 1–13.
Arouri, M., Boubaker, S., Grais, W., Grira, J., 2018. Rationality or politics? The color of black gold money. Q. Rev. Econ. Financ. 70, 62–76.
Arrow, K., Dasgupta, P., Goulder, L., Daily, G., Ehrlich, P., Heal, G., Levin, S., Mäler, K.-G., Schneider, S., Starret, D., Walker, B., 2004. Are we consuming too much? J. Econ. Perspect. 18 (3), 147–172.
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.M., Wang, G., 2013. Do (some) university endowments earn alpha? Financ. Anal. J. 69 (5), 26–44.
Bernstein, S., Lerner, J., Schoar, A., 2013. The investment strategies of sovereign wealth funds. J. Econ. Perspect. 27, 219–237.
Braunstein, J., 2022. Capital Choices: Sectoral Politics and the Variation of Sovereign Wealth. University of Michigan Press, Ann Arbor.
Brown, J.R., Dimmock, S.G., Kang, J.-K., Weisbrenner, S.J., 2014. How university endowments respond to financial market shocks: evidence and implications. Am. Econ. Rev. 104, 931–962.
Campbell, J., Deaton, A., 1989. Why is consumption so smooth? Rev. Econ. Stud. 56, 357–374.
Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. J. Polit. Econ. 107 (2), 205–251.
Campbell, J.Y., Martin, I.W.R., 2022. Sustainability in a Risky World. Harvard University and London School of Economics.
Campbell, JY., Sigalov, R., 2022. Portfolio choice with sustainable spending: a model of reaching for yield. J. Financ. Econ. 143, 188–206.
Choi, K.J., Junkee, J., Koo, H.K., 2022. Intertemporal preference with loss aversion: consumption and risk-attitude. J. Econ. Theory 200, 10538 2022.
Cochrane, J.H., 2022. Porfolios for long-term investors. Rev. Financ. 26 (1), 1–42.
Constantinides, G.M., 1990. Habit formation: a resolution of the equity premium puzzle. J. Polit. Econ. 98 (3), 519–543.
Dahiya, S., Yermack, D., 2018. Investment returns and distribution policies of non-profit endowment funds. 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. Global Financ. J. 34, 43–53.
Duesenberry, JS., 1949. Income, Saving, and the Theory of Consumer Behavior. Harvard University Press, Cambridge, Mass.
Dybvig, PH., 1995. Duesenberry's Racheting of consumption: optimal dynamic consumption and investment given intolerance for any decline in standard of living. Rev. Econ. Stud. 62 (2), 287–313.
Dybvig, PH., 1999. Using asset allocation to protect spending. Financ. Anal. J. 55 (1), 49–62.
Dybvig, PH. and Qin Z., 2021. "How to squander your endowment: pitfalls and remedies." Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3939984 or 10.2139/ssrn.3939984.
Gabaix, X., Koijen, RS.J., 2022. In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis. Swiss Finance Institute Research Paper No. 20-91, available at.
Galí, J., 1990. Finite horizons, life-cycle savings, and time-series evidence on consumption. J. Monetarh Econ. 26, 433–452.
Grossman, S.J., Laroque, G., 1990. Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods. Econometrica 58 (1), 25–51.
Irarrazabal, A.A., Ma, L., Parra-Alvarez, J.C., 2023. Optimal asset allocation for commodity sovereign wealth funds. Quant. Financ. 23 (3), 471–495.
Jeon, J., Park, K., 2021. Portfolio selection with drawdown constraint on consumption: a generalization model. Math. Methods Oper. Res. 93, 243–289.
Johan, S., Knill, A., Mauck, N., 2013. Determinants of sovereign wealth fund investment in private equity vs. public equity. J. Int. Bus. Stud. 44, 155–172.
Lindset, S., Mork, K.A., 2019. Risk taking and fiscal smoothing with sovereign wealth funds in advanced economies. Int. J. Financ. Stud. 7 (4), 1–24.
Mankiw, N.G, 1985. Small menu costs and large business cycles: a macroeconomic model of monopoly. Q. J. Econ. 100 (2), 529–537.
Mehra, R., Prescott, E.C., 1985. The equity premium: A puzzle. J. Monet. Econ. 15, 145–161.
Merton, R.C., 1971. Optimum consumption and portfolio rules in a contnuous-time model. J. Econ. Theory 3 (4), 373–413.
Mork, K.A., Trønnes, H.A., 2023. Expected Long-Term Rates of Return When Short-Term Returns Are Serially Correlated. Norwegian University of Science and Technology (NTNU).
Mork, K.A., Trønnes, H.A., Bjerketvedt, V.S., 2022. Capital preservation and current spending with sovereign wealth funds and endowment funds: a simulation study. Int. J. Financ. Stud. 10, 1–24.
Paltrinieri, A., Pichler, F., Falzon, J., 2013. Asset management issues in sovereign wealth funds: an empirical analysis. Bank Performance, Risk and Securitization. Palgrave Macmillan, London.
Rackauckas, C., Nie, Q., 2017. DifferentialEquations.jl-2017 – a performant and feature-rich ecosystem for solving differential equations in Julia. J. Open Res. Softw. 5, 1–10.
Rogers, L.C.G., 2013. Optimal Investment. Springer, London.
Rotemberg, JJ., 1982. Sticky prices in the United States. J.Polit. Econ. 90 (6), 1187–1211.
Shin, Y.H., Lim, B.H., Jin Choi, U., 2007. Optimal consumption and portfolio selection problem with downside consumption constraints. Appl. Math. Comput. 188, 1801–1811.
Shrikhande, M.M., 1997. Non-addictive habit formation and the equity remium puzzle. Eur. Financ. Manag. 3 (3), 293–319.
Tobin, J., 1974. What is permanent endowment income? Am. Econ. Rev. 64 (2), 427–432.
van den Bremer, T., van der Ploeg, F., Willis, S., 2016. The elephant in the ground: managing oil and sovereign wealth. Eur. Econ. Rev. 82, 113–131.
Watson, J.G. and Scott J.S., 2015. "First order loss aversion and optimal consumption." Available at SSRN: https://ssrn.com/abstract=2664820 or 10.2139/ssrn.2664820
World Commission on Environment and Development, 1987. Our Common Future. United Nations.