Quantitative Finance

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/rquf20

Optimal asset allocation for commodity sovereign wealth funds

Alfonso A. Irarrazabal, Lin Ma & Juan Carlos Parra-Alvarez

To cite this article: Alfonso A. Irarrazabal, Lin Ma & Juan Carlos Parra-Alvarez (2023) Optimal asset allocation for commodity sovereign wealth funds, Quantitative Finance, 23:3, 471-495, DOI: 10.1080/14697688.2022.2158918

To link to this article: https://doi.org/10.1080/14697688.2022.2158918

Published online: 10 Jan 2023.

Quantitative Finance, 2023
Vol. 23, No. 3, 471–495, https://doi.org/10.1080/14697688.2022.2158918

Optimal asset allocation for commodity sovereign wealth funds

ALFONSO A. IRARRAZABAL†, LIN MA ‡ and JUAN CARLOS PARRA-ALVAREZ §¶||*

†BI Norwegian Business School, Oslo, Norway
‡Center for International Climate Research, CICERO, Oslo, Norway
§Aalborg University Business School, Fibigerstræde 2, 9220 Aalborg Ø, Denmark
¶CREATES, Center for Research in Econometric Analysis of Time Series, Fuglesangs Alle 4, 8210 Aarhus V, Denmark
||DFI, Danish Finance Institute, Frederiksberg, 2000, Denmark

(*Received 29 July 2021; accepted 9 December 2022; published online 11 January 2023*)

This paper studies the dynamic asset allocation problem faced by an infinitively lived commodity-based sovereign wealth fund under incomplete markets. Assuming that the fund receives a non-tradable stream of commodity revenues until a predetermined date, the optimal consumption and investment strategies are state and time-dependent. Using data from the Norwegian Petroleum Fund, we find that the optimal demand for equity should decrease gradually from 60% to 40% over the next 60 years. However, the solution is particularly sensitive to the correlation between oil and stock price changes. We also estimate wealth-equivalent welfare losses, relative to the optimal rule, when following alternative suboptimal investment rules.

Keywords: Optimal asset allocation; Sovereign wealth fund; Commodities; Income risk; Suboptimal investments

JEL Classifications: E21, G11, G23, Q32

1. Introduction

Sovereign wealth funds (SWFs hereafter) are institutional investors that engage in long-run investment strategies with the objective to ensure a gradual transfer of wealth across generations. Although these investment funds have existed for decades, there has been a significant increase of SWFs since 2000. The source of income of most SWFs comes from commodity revenues and/or the accumulation of foreign exchange reserves. As of 2019, there were 48 different commodity-based SWFs in the world administering US$4 trillion in assets (SWF Institute 2020), corresponding to US$1,163 billion more than the estimated size of hedge funds worldwide (Statista 2019), and to 5% of the global investment industry (Fages *et al.* 2019)†. Since commodity prices are extremely volatile (cf. Deaton and Laroque 1992), investors face the challenge to design optimal investment strategies that help them manage the associated income risk. To the extent that commodity revenues are correlated with stock prices, investors have the possibility to hedge this volatility away by adjusting their exposure to stocks.

In this paper, we study the optimal consumption-investment decision of oil-based SWFs when the risk from its volatile revenues is only partially hedgeable due to market incompleteness. To do so, we use an otherwise standard strategic asset allocation model with stochastic income similar to those in Bodie *et al.* (1992), Heaton and Lucas (1997), Viceira (2001), Campbell and Viceira (2002), Cocco *et al.* (2005), Munk and Sørensen (2010), and Bosserhoff *et al.* (2022). However, since most SWFs are set up by countries interested in sustaining a standard of living for all future generations, we assume that the fund's planning horizon is infinite, while the commodity revenues are received only for a known fixed number of years that is determined exogenously by institutional factors, e.g. political decisions following pressure from constituents, changes in the production intensity of oil, commodity depletion, etc. In order to distinguish the effects of atemporal risk aversion from those due to intertemporal substitution, we assume that the preferences of the SWF's manager over intermediate consumption are recursive as in Duffie and Epstein (1992a, 1992b). In turn, this allows us to reconcile high risk-taking induced by large risk premiums with a low tolerance for volatile consumption. Moreover, all

*Corresponding author. Email: jcpalva@business.aau.dk
† A description of the investment behavior of SWFs in equity markets and their allocation strategies across industries and sectors is provided by Miceli (2013).

of the uncertainty in the income stream is assumed to come from the stochastic behavior of oil prices.

To solve the dynamic portfolio problem faced by the SWF, we invoke the principle of optimality (Bellman 1957) and break the planning horizon according to the terminal date of the commodity income flow. The resulting subproblems are then solved backwards in time. First, we solve an infinite horizon problem in which the SWF does not receive any revenues. The optimal value function, as well as the levels of investment and consumption for this problem, can be obtained in closed form. Using the resulting indirect utility as a terminal condition, we then solve a finite horizon problem in which the SWF receives a stream of stochastic commodity revenues for a predetermined number of years. Similar to the case of asset allocation over the life cycle with uncertain labor income, we find that the SWF's optimal investment profile is both state- and time-dependent. The optimal demand for the risky asset includes an intertemporal hedging component that depends on the correlation between risky assets and commodity income. Negative (positive) values of this correlation result in a decreasing (increasing) demand for stocks over time that converges to a long-run value equal to the (leveraged) Sharpe ratio of the risky asset.

To assess the degree of market incompleteness, we use monthly data on the S&P500 index and the WTI price of crude oil from 1973 to 2019 to estimate the correlation between oil and stock price changes using a continuous-time vector autoregressive model following the maximum-likelihood framework in Ait-Sahalia (2002, 2008). In line with previous results in the literature, we find statistical evidence of a time-varying correlation. Consistent with conventional wisdom, and similar to the evidence reported by Jones and Kaul (1996), Sadorsky (1999), Gorton and Rouwenhorst (2006), Lee and Chiou (2011) and Bhardwaj *et al.* (2015), we find a negative, but low, average correlation of -7% for the period 1973–2007. However, at the outset of the financial crisis that led to the Great Recession, the correlation becomes positive and high. More specifically, for the period 2008–2019, we estimate a statistically significant correlation coefficient that reaches a value of around 30% by the end of 2018. Similar results have been documented in Filis *et al.* (2011), Buyuksahin and Robe (2014), Bernanke (2016), Lombardi and Ravazzolo (2016) and Datta *et al.* (2021), who have argued that this phenomenon could be the result of a generalized weakening of global aggregate demand, the growth in the commodity-market activity, or the zero lower bound on nominal interest rates. Consequently, we conclude that oil income is not fully spanned by traded assets and hence its associated risk cannot be fully hedged through the financial markets. Although similar limits to diversification have been documented for long-term investors with stochastic labor income (see Campbell 1996, Davis and Willen 2000, Campbell and Viceira 2002), our estimates suggest that oil-based SWFs face a covariance structure between asset returns and income that is diametrically different: low correlation, and large volatility of income that exceeds that of stock prices. Therefore, in contrast to the case of uncertain labor income with no liquidity/investment constraints, it is no longer possible to accurately approximate the optimal investment strategy for oil-based SWF investors using the assumption of complete markets (cf. Bick *et al.* 2009, Munk and Sørensen 2010).

In the presence of stochastic oil revenues and market incompleteness, the finite horizon component of the model does not admit a closed-form solution. Therefore, we resort to numerical methods to approximate the optimal consumption and investment decisions. In particular, we use the state space reduction of Duffie *et al.* (1997), and the corresponding finite difference representation introduced by Munk and Sørensen (2010), which we implement numerically using the method described in Gomez (2019). We calibrate our model to match salient features of the Norwegian Government Pension Fund Global (GPFG), popularly known as the Petroleum Fund. For a low and negative correlation between oil and stock prices, similar to that observed between 1973 and 2007, our quantitative results indicate that the SWF should allocate around 60% of its financial wealth into stocks at the beginning of the planning horizon and decrease it monotonically thereafter to reach a value of 40% after 60 years. This initial overshooting, relative to its long-run value, is the result of two complementary effects: (i) a wealth or leverage effect from the capitalized value of the future stream of commodity revenues, and (ii) a positive hedging demand that accounts for 20% of the total demand for stocks. This additional demand is primarily driven by the high volatility of oil prices, and not by their correlation with the risky asset. If, on the contrary, the correlation is positive, the model implies a large recomposition of the investment portfolio with a large fraction of wealth allocated into the risk-free asset. In particular, with a correlation coefficient of 30%, we find an initial allocation to the risky asset of around 30% that should increase monotonically towards its long-run value of 40%. In either case, the optimal investment strategy goes hand-in-hand with a gradual transfer of wealth into the economy as measured by a relatively constant optimal consumption-to-wealth ratio that fluctuates between 2% and 3% per year. This consumption pattern is consistent with the fiscal rule (*handlingsregelen*) introduced by the Norwegian parliament in 2001 with the objective to spend oil revenues in a gradual and controlled way that helps preventing any undesirable overheating of the economy and/or the occurrence of a Dutch disease. Lindset and Mork (2019) have recently shown that such a smooth path is consistent with the government's desire for smoothness in taxes and public expenditures. As a corollary to our quantitative experiments, we conclude that if the correlation between stock and oil prices remains positive and large in the near future, the Norwegian GPFG should consider lowering its exposure to equity. Our simulations suggest that the current mandate on the stock/bond mix is not compatible with an investment strategy that exploits all the diversification possibilities in an optimal way, and instead exposes the fund to otherwise hedgeable risks.

Our results relate to a number of recent contributions to the study of asset allocation for oil-based SWFs. Scherer (2011) studies the portfolio problem of a SWF fund that must decide how to allocate its oil revenues into different asset classes. Through the lens of a standard mean–variance analysis (Markowitz 1952), and assuming that the value of the oil resources relative to the government's aggregate wealth is constant over time, he finds that the optimal demand for risky assets includes a hedging demand component that is a function of the oil wealth-to-financial wealth ratio, and of the correlation between oil price changes and asset returns. He shows that, when the set of investment opportunities includes assets that correlate negatively with oil prices, the SWF should then decrease its position in the risky asset as the oil reserves decrease. However, his approach abstracts from the optimal consumption-saving decision, and from the implications that the finite nature of the commodity revenues might have on the optimal investment strategy†. Closer to our approach is the work by van den Bremer *et al.* (2016). They extend the work in van den Bremer and der Ploeg (2013) to study the role played by non-tradable commodity assets in the optimal consumption and investment decisions of an infinitively lived SWF that is assumed to receive a stream of oil income to perpetuity. They also conclude that the optimal investment profile of an SWF should take into account the amount of underground wealth through a hedging demand component. Moreover, they show that any undiversifiable risk should be alleviated by an increase in precautionary savings against current consumption. However, to study the implications of their model under incomplete markets, they approximate the optimal allocations by assuming that consumption is a linear function of wealth, a result that only holds if markets are complete.

Next, we evaluate the welfare costs of not following an investment strategy that optimally exploits the intertemporal hedging opportunities available to the SWF. This exercise is motivated by the investment mandate given to the Norwegian GPFG according to which the equity/bond mix in the aggregate portfolio is fixed. Currently, the fund's portfolio manager (The Norges Bank Investment Management, NBIM) is allowed to invest between 60% and 80% of its wealth in equities. What are the consequences of deviating from the optimal investment strategy? Associated with a given suboptimal policy, we answer this question by introducing a measure of wealth-equivalent welfare compensation. The latter is defined as the percentage of additional initial financial wealth that the government would need to transfer to the portfolio administrator in order to achieve the same indirect utility or welfare that could be otherwise obtained by following the optimal investment strategy. In particular, we consider two different suboptimal investment profiles: (i) a constant investment share, and (ii) and ad-hoc deterministic rule that fixes the equity holdings in every period equal to the median optimal investment share. Using our benchmark calibration with a negative correlation between stock and oil prices, we find that following a strategy that fixes the position in equities at 70% (the midpoint of the current mandate of the Norwegian GPFG) would require a wealth compensation equivalent to 12.5% more of the initial endowment. An alternative interpretation of this result is that following a constant investment rule leads to significant welfare losses. We show that these losses can be considerably reduced by implementing instead a time-varying, but ad-hoc, investment rule. For practical purposes, this policy may be considered as a second-best policy in an environment with institutional constraints that prevent the SWF investor to hedge commodity fluctuations periodically.

We are not the first to report large welfare losses from the implementation of suboptimal policies. Campbell and Viceira (1999), find that failing to hedge in the presence of time-varying risk premia leads to large welfare losses relative to the optimal policy, specially for mildly risk-averse investors with positive positions in equity. Similarly, Gomes (2007) and Larsen and Munk (2012) report considerable utility losses from ignoring the intertemporal hedging opportunities for investors facing interest rate risk and stock volatility risk. Finally, Bick *et al.* (2009) study the welfare losses incurred by an investor with stochastic labor income that uses the investment rule that would prevail under complete markets when markets are in fact incomplete. They find that the losses of following this misspecified suboptimal policy are at most 14% of the initial total wealth when the true correlation between income and equities is zero, and drops to 3.2% if the correlation is 60%.

The remainder of the paper proceeds as follows. Section 2 provides a brief introduction and description of the Norwegian GPFG, with a particular focus on the institutional framework it faces and the investment strategy followed since its inception. Section 3 formalizes the optimal allocation problem faced by a commodity-based SWF and provides economic intuition behind the optimal consumption and investment policies when the oil income is both spanned and unspanned by the financial market. In Section 4, we discuss the calibration of the model and discuss the optimal allocations when markets are incomplete. We also study the sensibility of the optimal policies to changes in the correlation between stock and oil prices, the investor's coefficient of relative risk aversion, and to different assumptions on the terminal date of the commodity income inflow. Section 5 studies the welfare costs of following suboptimal policies, and Section 6 concludes.

2. The Norwegian sovereign wealth fund

Norway has one of the world's largest established SWFs, the Government Pension Fund Global (GPFG). In 2019, the market value of the GPFG amounted to US$1,148 billion, nearly 3.5 times the real GDP of mainland Norway,‡ and about 26% larger in market value than the China Investment Corporation which, according to the Sovereign Wealth Fund Institute, is the second biggest SWF. Panel (a) in figure 1 shows the uninterrupted accumulation of wealth for the period 1998–2019.

The GPFG was created by the Norwegian Parliament in 1990 under the Government Pension Fund Act in order to

† Using monthly data for the period 1997–2008, the author does not find any significant correlation between the nominal returns on the MSCI World index and the nominal price of crude oil. Therefore, he concludes that global equities do not provide a hedge against fluctuations in oil prices. In the face of a similar insignificant correlation, Døskeland (2007) proposed to use a cointegration approach to identify the long-term relationship between financial assets and non-tradable assets. When applied to the Norwegian case, he finds a similar result: the government should increase its current (initial) holdings in equity and reduce it over time.

‡ According to Statistics Norway, mainland Norway refers to all the domestic production activity with the exception of exploration of crude oil and natural gas, transport via pipelines and ocean transport.

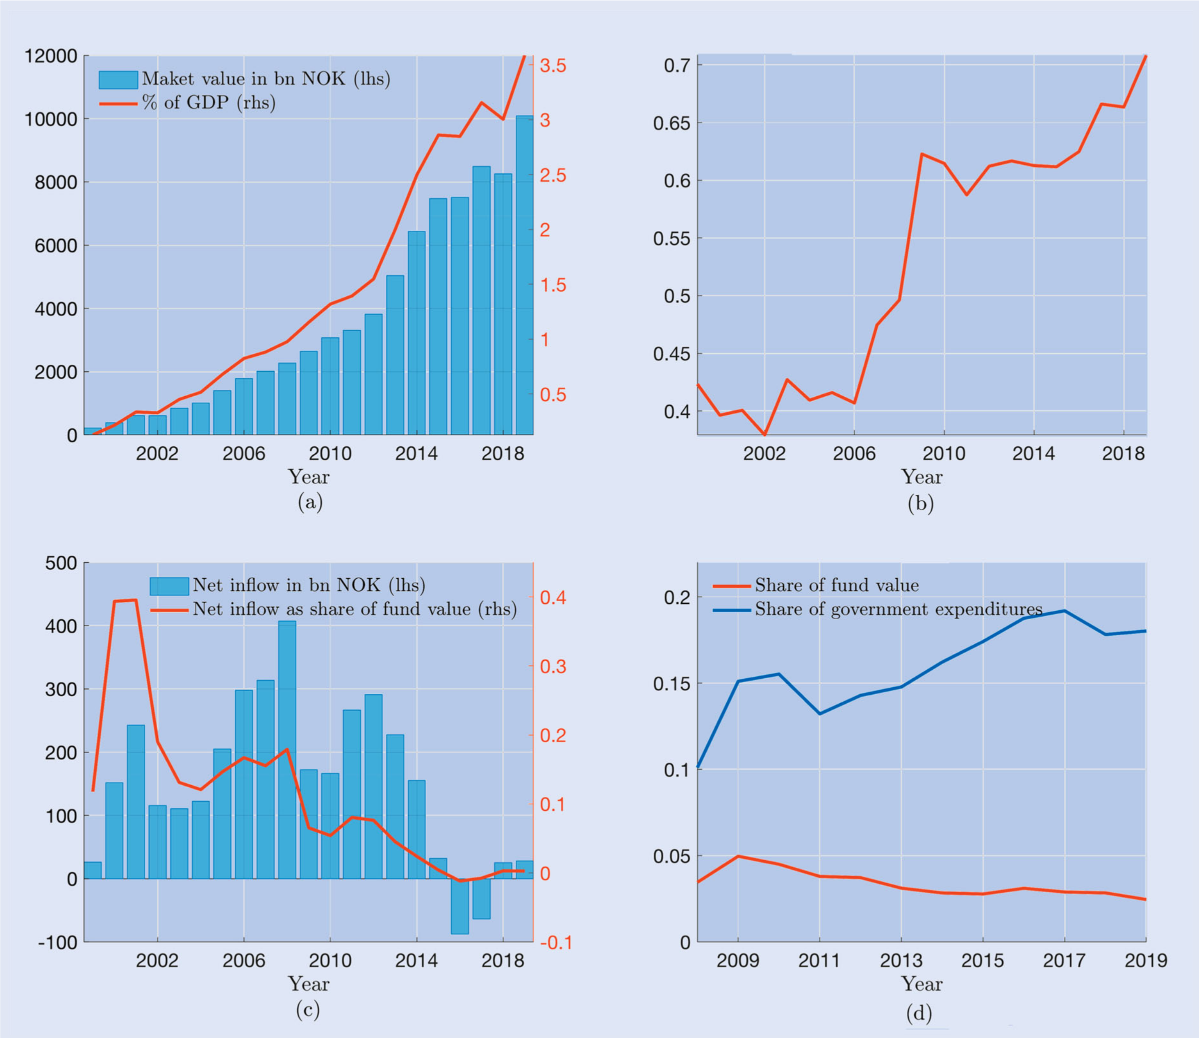

Figure 1. Main facts about the GPFG (1998–2019). Panel (a) plots the market value of the GPFG in billions of NOK and as a fraction of mainland real GDP; Panel (b) plots the fraction of the GPFG's market value invested in equities; Panel (c) plots the total capital inflows (net of managements costs) to the fund associated with oil revenues in billions of NOK and as a fraction of the GPFG's market value of wealth; Panel (d) plots the transfers from the GPFG to the Central Government as a fraction of the GPFC's market value and as a fraction of the total budgeted expenditures for the period 2006–2019.
Norges Bank Investment Management, Statistics Norway and Ministry of Finance.

ensure long-term responsible management of the revenues generated from oil-related activities. Specifically, the objective of the fund is to manage the financial challenges posed by an aging population and to serve as a countercyclical fiscal tool that can be used to neutralize declines in the price of oil and in the economic activity in general. The Ministry of Finance has the overall responsibility for the fund's management. Accordingly, it issues a set of guidelines that are executed by the fund's Executive Board who defines an investment policy that is implemented by the portfolio manager. The Ministry's guidelines delimit the types of risks that the SWF can take, and the Executive Board acts consequently by setting up an asset allocation strategy that distributes the fund's wealth into different asset classes.

In 1998 the Norges Bank Investment Management (NBIM) was created to administer the fund's portfolio. The NBIM receives oil revenues in the form of transfers from the government and combines them with the fund's own accumulated wealth to implement the requirements defined by the Executive Board. Panel (b) in figure 1 shows the effective share of wealth is invested in equity markets over the period 1998–2019. This investment profile mimics the management guidelines set by the Ministry of Finance, who in 1998 set a cap on the amount of wealth invested in stocks to 40%. In 2007, this cap was increased such that the equity portfolio constitutes between 60% and 80% of the total portfolio. The remaining fraction of wealth is distributed between fixed income (20–40%) and real estate assets (up to 7%).

The GPFG is set up such that two types of revenues are transferred to NBIM directly: government oil revenues and the fund's return. Panel (c) shows the annual (net) inflows to the NBIM for the period 1998–2019. Throughout the period, the net transfers to the fund have decreased as a fraction of the GPFG's total wealth. As shown in figure 2, this behavior is consistent with three factors: (i) the drop in the average real price of oil observed from 2006, Panel (a); (ii) the decline in the production of crude oil in Norway, Panel (b); and (iii) the fall in the proved reserves of oil, Panel (c). All of these factors have led to a reduced operating surplus from extraction activities.

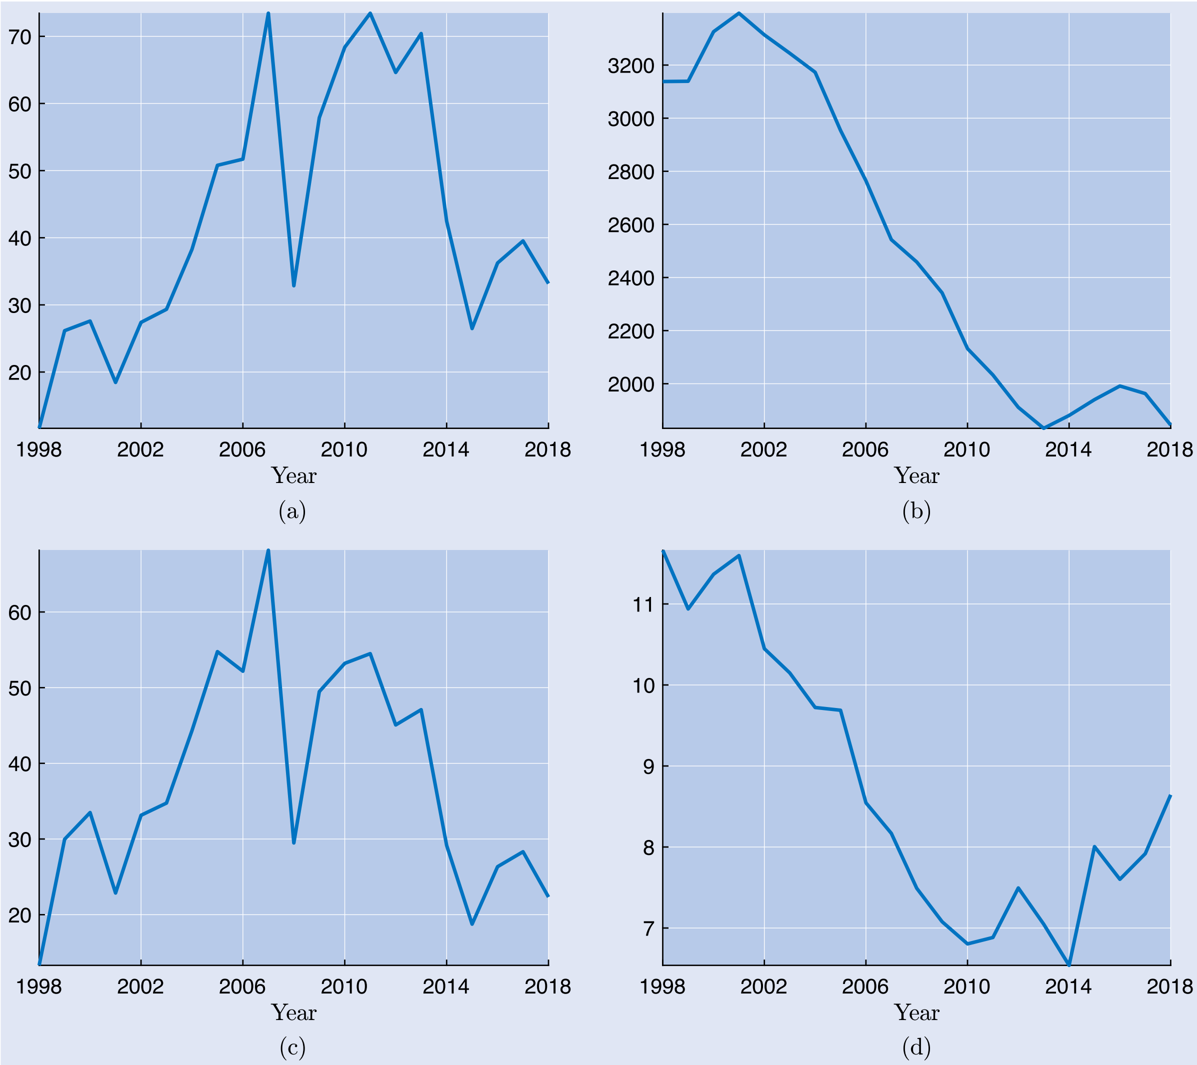

Figure 2. Main components of oil revenues (1998–2018). Panel (a) plots the real price of oil per barrel in U.S. dollars as measured by the WTI price index and deflated by the U.S. CPI base 2000=100 (P_t). Panel (b) plots the production of crude oil in Norway* (Q_t). Panel (c) plots the total oil revenues in U.S. dollars from oil production ($P_t \times Q_t$). Panel (d) plots the total proved reserves of underground oil available at the end of a given year**.
BP Statistical Review of World Energy (BP 2019)
*Includes crude oil, shale oil, oil sands, condensates and NGLs. Excludes liquid fuels from other sources such as biomass and derivatives of coal and natural gas. **Correspond to quantities that geological and engineering information indicates with reasonable certainty can be recovered in the future from known reservoirs under existing economic and operating conditions.

In addition to the NBIM funding sources, the Ministry of Finance established in 2001 a fiscal spending rule that stipulates how to transfer the oil income and its associated returns back to the Norwegian economy in a smooth and controlled way. These transfers from NBIM to the Central Government budget should follow the expected real return on the fund and must be directed to finance non-oil fiscal budget deficits. According to NBIM (2016): '*The spending rule is not a legal requirement, but rather a political economic yardstick which secures the original fund objectives and strengthens the intergenerational perspective.*' The transfer rule was initially set at 4% and in February of 2017, it was reduced to 3%. Panel (d) in figure 1 plots a realized measure of the transfer rule computed as the value of the non-oil deficits budgeted by the Ministry of Finance as a fraction of the GPFG's market value. For the period 2006–2019, the transfers to the Central Government averaged 3.4% of the fund's wealth, and have become an important source of funding of the government total expenditures: in 2018 these transfers represented 18% of the total government spending.

3. The allocation problem of a SWF investor

This section describes the problem faced by a price-taking commodity SWF manager. Time is assumed to evolve continuously. Our framework is a stylized representation of the decision problem faced by the Executive Board of the Norwegian GPFG (the fund's manager) introduced in Section 2. We focus on the optimal asset allocation decisions made by the fund's manager who is also required to ensure a smooth stream of transfers to the government conditional on a decreasing and finite path of commodity revenues. The manager's planning horizon is assumed to be infinite to capture the long-term objective of building financial wealth that ensures sustained transfers for future generations. Moreover, the fund receives a stream of commodity income continuously over a known fixed number of years that are determined exogenously by the fund's owner.

3.1. Description of the model

For a mathematical formulation of the problem faced by the fund's manager, we consider an economy where uncertainty is described by a complete probability space $(\Omega, \mathbb{P}, \mathcal{F})$ endowed with a filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t\geq 0}$, representing the information available at time t.

3.1.1. Investment opportunities. The fund's manager has costless access to two tradable assets. A money market account with a constant, continuously compounded, real return $r > 0$ (risk-free bond), and a risky asset (stock market index) with adapted price process $S = (S_t)_{t\geq 0}$ that evolves over time according to a geometric Brownian motion

$$\mathrm{d}S_t = \mu S_t\,\mathrm{d}t + \sigma_S S_t\,\mathrm{d}Z_{S,t}, \quad S_0 > 0, \tag{1}$$

where μ denotes the asset's constant instantaneous return, $\sigma_S > 0$ is the constant volatility, and $Z_{S,t}$ is a standard Brownian motion with respect to $\mathbb{F}$. Thus, we assume that the manager's investment opportunity set is constant.

3.1.2. Oil income. Assuming zero exploration (and discoveries) of new reserves, the availability of the natural resource Q_t decreases exogenously over time at a constant extraction rate $\kappa_Q > 0$†, i.e.

$$\mathrm{d}Q_t = -\kappa_Q Q_t\,\mathrm{d}t, \quad Q_0 > 0. \tag{2}$$

The price per unit of the commodity is given by the adapted price process $P = (P_t)_{t\geq 0}$ with dynamics described by a geometric Brownian motion with drift rate κ_P, and constant volatility $\sigma_P > 0$,

$$\mathrm{d}P_t = \kappa_P P_t\,\mathrm{d}t + \sigma_P P_t\left(\rho_{PS}\,\mathrm{d}Z_{S,t} + \sqrt{1-\rho_{PS}^2}\,\mathrm{d}Z_{P,t}\right), \quad P_0 > 0, \tag{3}$$

where $Z_{P,t}$ is a standard Brownian motion with respect to $\mathbb{F}$ and independent of $Z_{S,t}$. Moreover, $\langle Z_S, Z_P\rangle_t = \rho_{PS} t$ where $|\rho_{PS}| \leq 1$ denotes the instantaneous correlation between the return to the risky asset and the change in the commodity price. As mentioned earlier, we assume throughout that the SWF takes the price of the commodity as given. The price is determined in the international markets and unaffected by the extraction rate.

† Notice that the commodity converges to zero asymptotically but never gets fully depleted in finite time.

Let $Y_t := P_t Q_t$ denote the fund's continuous stream of non-negative income. In the following, we assume that the fund receives a continuous stream of exogenously given commodity income until the known time $\hat{T} \geq 0$. Itô's Lemma implies that the fund's revenue process $Y = (Y_t)_{t\geq 0}$ evolves according to

$$\mathrm{d}Y_t = \kappa Y_t\,\mathrm{d}t + \sigma_P Y_t\left(\rho_{PS}\,\mathrm{d}Z_{S,t} + \sqrt{1-\rho_{PS}^2}\,\mathrm{d}Z_{P,t}\right), \quad Y_0 > 0, \tag{4}$$

for $t \in (0, \hat{T}]$, and where $\kappa = (\kappa_P - \kappa_Q)$ represents the expected income growth. For $t > \hat{T}$, $Y_t = 0$. Equation (4) assumes that all the uncertainty in the oil income arises from the exogenous variation in the commodity price set in the world markets.

In the particular case $|\rho_{PS}| = 1$, financial markets are complete and all the uncertainty in the oil income process is spanned by the stock price process. In other words, the stream of revenues can be perfectly replicated by some trading strategy in the financial markets and hence, valued as a traded asset. If $|\rho_{PS}| < 1$, the markets are said to be incomplete and the income risk is not fully spanned by financial markets.

3.1.3. Preferences. We assume that the manager has recursive preferences as first proposed by Kreps and Porteus (1978), Epstein and Zin (1989) and Weil (1990), and extended to continuous time by Duffie and Epstein (1992a, 1992b). This allows us to disentangle the effects that risk aversion and intertemporal substitution have separately on the optimal investment and consumption decisions. In particular, the preferences of the fund's manager are given by

$$V_t = \mathbb{E}_t\left[\int_t^\infty f(C_s, V_s)\,\mathrm{d}s\right], \tag{5}$$

where $f(C_s, V_s)$ is a normalized aggregator of the current consumption rate, C_s, and utility, V_s. The process $V = (V_t)_{t\geq 0}$ is referred to as the continuation value process associated with C. In its more general form, the aggregator is defined as

$$f(C, V) = \beta\theta V\left\{\left[\frac{C}{[(1-\gamma)V]^{\frac{1}{1-\gamma}}}\right]^{1-\frac{1}{\psi}} - 1\right\}, \tag{6}$$

where $\beta > 0$ is the rate of time preference, $\gamma > 0$ denotes the coefficient of relative risk aversion (RRA) towards atemporal bets, $\psi > 0$ denotes the elasticity of intertemporal substitution (EIS), so that $1/\psi$ can be understood as aversion towards intertemporal fluctuations, and $\theta = (1-\gamma)/(1-\psi^{-1})$. The normalized aggregator exhibits the property that for $\gamma > 1/\psi$, the investor prefers early over later resolution of uncertainty. The ability to separate the investor's risk aversion from her aversion to intertemporal substitution is important in order to generate a smooth path for the consumption-to-wealth ratio that emulates the intergenerational concerns of the government without affecting the short-term allocation strategies that can be achieved through portfolio diversification. If $\gamma = 1/\psi$ it follows that $\theta = 1$ which makes the recursion in (6) linear, and the preferences in (5) collapse to the usual time-additive utility model. In this case, it is no longer possible to separately investigate how the manager's optimal portfolio is affected by her attitudes towards risk without affecting at the same time intertemporal choices.

3.1.4. The manager's decision problem. Let W_t denote the fund's financial wealth at time t, i.e. the value of the portfolio of financial assets held at time t. Given an initial endowment of financial wealth, $W_0 = w > 0$, and an initial level of oil income, $Y_0 = y > 0$, the fund's manager must choose consumption rates $C = (C_t)_{t\geq 0}$, and investment strategy $\alpha = (\alpha_t)_{t\geq 0}$, with α_t denoting the fraction of financial wealth invested in the risky asset at time t, and $1 - \alpha_t$ the fraction of wealth invested in the risk-free asset. The stochastic variables C_t and α_t are assumed to be $\mathcal{F}_t$-measurable, so the processes C and α are adapted. Moreover, we assume that (i) $C_t \geq 0$, (ii) the process C is such that there exists a uniquely determined semimartingale $V = (V_t)_{t\geq 0}$ satisfying (5); (iii) the fund is able to continuously rebalance its portfolio, and does not face restrictions on borrowing or short sales, (iv) conditional on the instantaneous flow of revenues, the investment strategy α is self-financing and supports the consumption plan C, and (iv) the manager knows the stochastic processes that drive the oil and stock prices. The set $\mathcal{A}$ of all consumption and investment strategies that satisfy the above conditions on the interval $[0, \infty]$ is said to be admissible. For a given consumption and investment strategies $(C, \alpha) \in \mathcal{A}$, the fund's financial wealth evolves according to

$$\mathrm{d}W_t = \left(\mu_{P,t} W_t - C_t + Y_t\right) \mathrm{d}t + \sigma_{P,t} W_t \,\mathrm{d}Z_{S,t}, \qquad W_0 = w > 0 \tag{7}$$

where $\mu_{P,t} := \alpha_t(\mu - r) + r$ and $\sigma_{P,t} := \alpha_t \sigma_S$ represent, respectively, the instantaneous expected return and volatility per unit of financial wealth on the composite portfolio held by the fund, and where Y_t solves (4) with $Y_t = 0$ for all $t > \hat{T}$.

Then, the problem faced by the fund's manager is to find the consumption and investment strategies (C, α) that maximize the present discounted value of her non-expected recursive utility†. In other words, to compute for all $t \in [0, \infty)$

$$J_t = \max_{(C,\alpha)\in\mathcal{A}} V_t, \tag{8}$$

with V_t given in (5), and where $J_t := J(W_t, Y_t)$ is the indirect utility or value function.

† Note that the fund could also choose derivatives instruments, e.g. futures and forward contracts, to hedge unspanned oil income risk. This is an interesting avenue to explore; however, we abstain from doing it for several reasons. First, we do not observe in practice NBIM taking positions in futures contracts on oil on behalf of the Norwegian GPFG. Marking to market of daily profit and losses for a large number of contracts in oil may lead to liquidity constraints, a point made in van den Bremer *et al.* (2016). Second, the institutional mandate of the Norwegian GPFG requires NBIM to exclusively focus on the risk of the portfolio (see Norges Bank Investment Management 2022). Third, it may be that because of the sheer size of the portfolio that NBIM manages, the futures market is not able to provide sufficient liquidity.

3.2. Solution under complete and incomplete markets

To solve the problem faced by the fund's manager we invoke the principle of optimality to break the infinite horizon implicit in problem in (8) into two subproblems according to the predetermined terminal time of income $\hat{T}$, i.e.

$$J_0 = \max_{(C,\alpha)\in\mathcal{A}_t} \mathbb{E}_0 \left[\int_0^{\hat{T}} f(C_t, J_t)\,\mathrm{d}t + J\left(W_{\hat{T}}\right) \right], \tag{9}$$

where $\mathcal{A}_t \subseteq \mathcal{A}$ is the set of admissible strategies for $t \in [0, \hat{T}]$. Thus for $t \leq \hat{T}$, the manager solves a finite horizon asset allocation problem with stochastic income and terminal utility $J(W_{\hat{T}})$, where $W_{\hat{T}}$ represents the fund's financial wealth at time $\hat{T}$‡. For $t > \hat{T}$, the manager no longer receives oil income and faces the following infinite horizon allocation problem (cf. Svensson 1989, Campbell and Viceira 2002, and Hsuku 2007)

$$J\left(W_{\hat{T}}\right) = \max_{(C,\alpha)\in\mathcal{A}} \mathbb{E}_{\hat{T}} \left[\int_{\hat{T}}^{\infty} f(C_s, J_s)\,\mathrm{d}s \right]. \tag{10}$$

Therefore, our solution strategy consists of solving the allocation problem in two stages. In the first stage, we compute the optimal consumption and investment policies that will prevail for $t > \hat{T}$. The second stage uses the value of the optimal program at time $\hat{T}$ found in the first stage as a terminal condition to solve for the optimal allocations for all $0 \leq t \leq \hat{T}$. Notice that since the investment horizon is infinite, the optimal value of $J(W_{\hat{T}})$ (from the first stage) ensures that the wealth at $\hat{T}$ is the optimal transfer of resources to future generations (those for which $t > \hat{T}$).

3.2.1. First stage. For $t > \hat{T}$, it follows that $Y_t = 0$ and the problem faced by the fund's manager becomes a standard infinite horizon dynamic asset allocation problem with constant investment opportunities and complete financial markets. As shown in Appendix 1, the *Hamilton–Jacobi–Bellman* (HJB) equation for this problem is

$$0 = \max_{\{C,\alpha\}\in\mathcal{A}} \left\{ f(C, J) + J_W \left[rW + (\mu - r)\,\alpha W - C \right] + \frac{1}{2}\sigma_S^2 J_{WW} \left(\alpha W\right)^2 \right\}, \tag{11}$$

with $J_W := \partial J(W)/\partial W$ and $J_{WW} := \partial^2 J(W)/\partial W^2$, and where we have suppressed time indexes to reflect the recursive nature of the corresponding dynamic programming problem. The next proposition summarizes the closed-form solution to the first stage problem.

‡ The stochastic optimal control problem in (9) is equivalent to a finite horizon dynamic asset allocation problem with terminal utility $J(W_{\hat{T}}) = A(W_{\hat{T}}^{1-\gamma}/1 - \gamma)$. The constant A, which measures the relative utility weight of wealth at instant $\hat{T}$ and intermediate consumption, is then chosen to be consistent with the assumption that the fund manager optimizes the intertemporal utility of consumption even after the flow of revenues terminates.

PROPOSITION 3.1 *In the absence of commodity revenues, i.e. $Y_t = 0$, and constant investment opportunities, the manager's optimal value function for any $t \geq \hat{T}$ is given by*

$$J(W) = \frac{\beta^\theta}{1-\gamma} \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W^{1-\gamma}, \tag{12}$$

where the constant $\mathbb{G}_\infty$ is given by

$$\mathbb{G}_\infty = \beta\psi + (1-\psi) r + \frac{(1-\psi)}{2\gamma} \left(\frac{\mu - r}{\sigma_S}\right)^2. \tag{13}$$

The optimal share of financial wealth invested in the risky asset is given by

$$\alpha_t = \frac{1}{\gamma} \frac{\mu - r}{\sigma_S^2}, \tag{14}$$

while the optimal consumption-to-financial wealth ratio is

$$\frac{C_t}{W_t} = \mathbb{G}_\infty = r_P + \psi (\beta - r_P), \tag{15}$$

where $r_P := \mu_P - \frac{\gamma}{2}\sigma_P^2$ is the (risk-adjusted) expected return on the composite portfolio (see van den Bremer *et al.* 2016 *and* Wang *et al.* 2016*).*

Proof See Appendix 1. ∎

Equation (14) suggests that the optimal demand for the risky asset is constant for $t \geq \hat{T}$ and completely characterized by the market price of risk (the ratio of the expected risk premium to the asset's volatility) rescaled by the asset's volatility, and the manager's coefficient of RRA. It is also independent of the EIS and the subjective discount rate. The higher the coefficient of RRA, the lower the investment in stocks and the higher the investment in the risk-free asset (see Merton 1969).

On the other hand, the optimal consumption is given by the modified Keynes–Ramsey rule in (15). It suggests that optimal consumption is a linear function of the fund's financial wealth. The marginal propensity to consume is constant, and its value is determined by the coefficient of RRA, the manager's EIS, and the subjective discount rate.

3.2.2. Second stage. The problem faced by the fund's manager from the perspective of time $t < \hat{T}$ corresponds to a finite horizon allocation problem with stochastic income similar to that in Munk and Sørensen (2010) and Wang *et al.* (2016) with terminal utility given by the power utility of financial wealth in (12) evaluated at $t = \hat{T}$. As shown in Appendix 2, the HJB equation for this problem is

$$\begin{aligned} 0 = \max_{C,\alpha} \Big\{ & f(C,J) + J_t + J_W [rW + \alpha(\mu - r) W + Y - C] \\ & + \frac{1}{2}\sigma_S^2 (\alpha W)^2 J_{WW} + \kappa Y J_Y \\ & + \frac{1}{2}\sigma_P^2 Y^2 J_{YY} + \sigma_S \sigma_P \rho_{PS} (\alpha W) Y J_{WY} \Big\}, \end{aligned} \tag{16}$$

where $J_t := \partial J(t,W,Y)/\partial t$, $J_Y := \partial J(t,W,Y)/\partial Y$, $J_{YY} := \partial^2 J(t,W,Y)/\partial Y^2$, and $J_{WY} := \partial^2 J(t,W,Y)/\partial W \partial Y$. The first-order conditions for an interior solution for any $t < \hat{T}$ are given by

$$\frac{C_t}{W_t} = \left(\frac{\beta}{J_W}\right)^\psi \frac{[(1-\gamma) J]^{\frac{1-\psi\gamma}{1-\gamma}}}{W_t} \tag{17}$$

$$\alpha_t = \underbrace{\frac{1}{-\frac{W_t J_{WW}}{J_W}} \frac{\mu - r}{\sigma_S^2}}_{\text{Myopic demand}} + \underbrace{\frac{Y_t J_{WY}}{J_W} \frac{1}{-\frac{W_t J_{WW}}{J_W}} \frac{\sigma_P \rho_{PS}}{\sigma_S}}_{\text{Hedging demand}}. \tag{18}$$

Equation (17) results from the standard envelope condition $f_C = J_W$. Accordingly, the optimal consumption-to-financial wealth ratio is such that the marginal benefit of an additional unit of consumption is equal to the marginal utility of an additional unit of financial wealth. Equation (18) determines the optimal share of financial wealth allocated to the risky asset as the sum of two components. The first term, usually referred to as the *myopic* or *speculative demand*, corresponds to the investment strategy that a manager with a short investment horizon will follow, i.e. an investor that ignores what happens beyond the immediate next instant. It is defined by the standard mean–variance analysis of Markowitz (1952) that suggest that the demand for risky assets should be proportional to the asset's risk premium over the risk-free asset, $(\mu - r)$, and inversely proportional to the asset's volatility, σ_S, and the investor's risk aversion captured by the curvature of the value function, $-WJ_{WW}/J_W$.

The second term, usually referred to as *intertemporal hedging demand* or *excess risky demand*, represents the additional demand required by an investor with a long investment horizon in order to hedge against the risk of unexpected changes in the commodity income that cannot be fully eliminated (see Merton 1969, 1971, 1973). It is determined by the volatility of income, σ_P, and its correlation with the stock returns, ρ_{PS}, the manager's coefficient of RRA, and his aversion to income risk as measured by YJ_{WY}/J_W. Importantly, this component suggests that the investor should increase her holding of the risky asset for increased levels of aversion to income risk and whenever its returns co-vary negatively with changes in the commodity income.

As seen from (17) and (18), the solutions to the optimal consumption and investment shares depend on the unknown time-dependent value function $J(t,W,Y)$. When financial markets are incomplete, i.e. for $|\rho_{PS}| < 1$, no closed-form solution is available, and we need to resort to numerical methods in order to approximate the optimal allocations. However, under the simplifying assumption of complete markets, i.e. $|\rho_{PS}| = 1$, the oil income can be valued as a stream of dividends which allows us to derive an analytical solution. Although this assumption is challenged by the empirically observed correlation between stock and oil prices, we now make use of the complete market solution to build the economic intuition on the main determinants of the optimal consumption and investment decisions when oil income is not perfectly spanned by the stock market. However, our main results are computed under the assumption of incomplete markets.

3.2.3. Complete market solution. Let us first define $\mathcal{O}_t := \mathcal{O}(Y_t, t; \hat{T})$ to be the time t value of the oil wealth, i.e. the present discounted value of all future oil income transfers to the fund's manager.

LEMMA 3.1 (Oil wealth under complete financial markets) *Assume a complete financial market, i.e.* $|\rho_{PS}| = 1$. *Then, the fund's manager oil wealth at time t is given by*

$$\mathcal{O}\left(Y_t, t; \hat{T}\right) = Y_t \mathcal{M}\left(t; \hat{T}\right), \quad \forall t < \hat{T}. \tag{19}$$

The time-dependent function $\mathcal{M}(t; \hat{T})$ *defines the commodity income multiplier*

$$\mathcal{M}\left(t; \hat{T}\right) = \frac{1}{r - \kappa \pm \sigma_P \lambda}\left[1 - \exp\left\{-\left(r - \kappa \pm \sigma_P \lambda\right)\left(\hat{T} - t\right)\right\}\right], \tag{20}$$

for $(r - \kappa \pm \sigma_P \lambda) \neq 0$, *and where* $\lambda = (\mu - r)/\sigma_S$ *is the market price of risk. For* $t \geq \hat{T}$ *the oil income is zero,* $Y_t = 0$, *and thus* $\mathcal{O}(Y_t, t; \hat{T}) = 0$.

Proof See Appendix A.2. ■

Lemma 3.1 shows that when markets are complete it is possible to decompose the level of oil wealth as the product between the current level of oil income, Y_t, and the time-dependent income multiplier, $\mathcal{M}(t; \hat{T})$. For a given $\hat{T}$, the oil wealth decreases as the fund approaches the income terminal date at a rate that depends on the financial market returns, the expected growth rate of oil income, and the volatility of oil income. Furthermore, the oil wealth is higher for an income stream that is negatively correlated with the stock market, than for a similar income stream, but positively correlated with the stock market. Given a fixed $\hat{T}$, a positive (negative) correlation implies that the future expected income will be discounted at a rate higher (lower) than the return on the risk-free asset.

PROPOSITION 3.2 *Assume a complete financial market, i.e.* $|\rho_{PS}| = 1$. *Let* $Y_t \mathcal{M}(t; \hat{T})$ *denote the market value of the fund's oil wealth at time t. Then, the optimal consumption-to-financial wealth ratio for all* $t < \hat{T}$ *is given by*

$$\frac{C_t}{W_t} = \mathbb{G}_\infty\left(1 + \frac{Y_t}{W_t}\mathcal{M}\left(t; \hat{T}\right)\right), \tag{21}$$

with $\mathbb{G}_\infty$ *is given in* (13). *The optimal share of financial wealth invested in the risky asset for all* $t < \hat{T}$ *is*

$$\alpha_t = \frac{1}{\gamma}\left(\frac{\mu - r}{\sigma_S^2}\right)\left(1 + \frac{Y_t}{W_t}\mathcal{M}\left(t; \hat{T}\right)\right) - \frac{Y_t}{W_t}\mathcal{M}\left(t; \hat{T}\right)\frac{\sigma_P \rho_{PS}}{\sigma_S}. \tag{22}$$

Proof See Appendix A.2. ■

The optimal consumption is given by the modified Keynes–Ramsey rule in (21). In the presence of stochastic oil income, consumption at a given point in time is linear in the fund's total wealth, $(W_t + \mathcal{O}_t)$. The marginal propensity to consume out of total wealth is constant, and its determinants are the same as in the case of no oil wealth. However, using Lemma 3.1 it is straightforward to show that the propensity to consume out of income is increasing in the financial wealth-to-oil income ratio, W_t/Y_t, and the expected income growth rate (oil income multiplier), and decreasing in the current income rate. As opposed to the case $Y_t = 0$, the optimal consumption-to-financial wealth ratio is no longer constant. Instead, the marginal propensity to consume out of the financial wealth will fall over time in line with the decrease in the oil wealth-to-financial wealth ratio. As the fund's oil wealth runs out over time, the consumption-to-financial wealth ratio converges to the constant level given in (15).

The first term on the right-hand side of (22) is the myopic demand for the risky asset, while the second term is the intertemporal hedging demand. Similar to the case without income, the optimal investment share is independent of the EIS. However, the presence of stochastic oil income ($\sigma_P > 0$) will have a magnifying effect on the investment share through the intertemporal hedging component, as long as the commodity income is correlated with the returns of the risky asset, $\rho_{PS} \neq 0$. The direction of this effect will depend on the sign of the instantaneous correlation. A negative (positive) correlation implies a positive (negative) hedging demand. Importantly, the optimal investment share is no longer constant over time. In particular, the trajectory depends on the path of the oil wealth-to-financial wealth ratio, $\mathcal{O}_t/W_t$, and it will converge to the constant level in (14) as time approaches the terminal date, $\hat{T}$. This result also holds for the case of a deterministic stream of income ($\sigma_P = 0$), or a stochastic stream of income that is uncorrelated with asset returns ($\rho_{PS} = 0$). The optimal demand for the risky asset in (22) can be alternatively written as

$$\alpha_t = \frac{1}{\gamma}\left(\frac{\mu - r}{\sigma_S^2}\right) + \left(\frac{1}{\gamma}\left(\frac{\mu - r}{\sigma_S^2}\right) - \frac{\sigma_P \rho_{PS}}{\sigma_S}\right)\frac{Y_t}{W_t}\mathcal{M}\left(t; \hat{T}\right), \tag{23}$$

where the first term is identical to the optimal investment without income, and the second term represents the effect of commodity income on the optimal investment strategy. Consequently, as long as the oil wealth-to-financial wealth ratio is positive, the convergence of the optimal investment strategy to its long-run value will depend on the coefficient of RRA, the expected excess return, and the correlation coefficient. For $(\mu - r)/\sigma_S > \gamma \sigma_P \rho_{PS}$, the convergence will be from above: if the expected excess return exceeds the covariance between the stock and oil prices, the fund's manager should decrease the fraction of financial wealth invested in the stock market as $t \to \hat{T}$. Furthermore, low values of the coefficient of RRA are associated with an intertemporal demand for the risky asset that exhibits larger deviations from the long-run value along its transition.

3.2.4. Incomplete market solution. Whenever† $|\rho_{PS}| < 1$, it is no longer possible to value the oil income as a traded

† To bring down the exposure to income volatility one could follow the ideas of Bick *et al.* (2009, 2013) and introduce an additional traded asset (e.g. derivative) to 'artificially' complete the market (see

asset and thus the portfolio strategy in Proposition 3.2 turns out to be suboptimal. In particular, it will overestimate the hedging demand component by incorrectly assuming that the income risk can be fully replicated in the financial markets. This will lead to higher investment shares and lower consumption rates. For the case of an investor with time-additive preferences ($\theta = 1$) and stochastic labor income, Munk and Sørensen (2010) show that the complete market solution can be used to approximate the optimal allocations in incomplete markets when the investor does not face liquidity nor investment constraints. The reason is that although the allocation policies are suboptimal, they imply small utility losses even for small correlations between income and asset prices. However, their recommendation does not carry to the case of a commodity SWF like the one studied in this paper. Even though the correlation coefficient between the price of oil and the price of equity is low, the high volatility of the oil price will result in a portfolio weight on the risky asset that is unreasonable high relative that suggested by the optimal strategy. This, in turn, can lead to large utility losses from implementing suboptimal allocations due to an excessive exposure to risk.

To characterize the model's optimal policies when oil income is not perfectly hedgeable we use a finite difference approach based on the work in Gomez (2019). In particular, the solution to the HJB equation in (16) is approximated backwards in time starting from the terminal value given by (12) evaluated at $t = \hat{T}$. At each point in time, we maximize the HJB equation over all possible consumption and investment choices along a predefined grid for the state variables. To obtain a stable and more efficient approximation of the unknown value function, we use the state space reduction of Duffie *et al.* (1997) to reduce the number of state variables from two to one by exploiting the homogeneity properties of the HJB equation. A complete description of the state reduction problem as well as of the finite difference solution method can be found in Appendix A.1.

4. Quantitative model under incomplete markets

This section explores the quantitative predictions of the model. We begin our analysis by calibrating the model parameters. Using these values, we compute the optimal value function, as well as the consumption and investment policies, that determine the long-run behavior of the fund in the absence of oil income. We then approximate the solution to the finite horizon problem to obtain the optimal consumption and investment profiles in the presence of a continuous stream of oil income. We end the section by investigating the sensitivity of our results to different values of the correlation coefficient and the coefficient of RRA, as well as to different assumptions on the income terminal date.

footnote 5). The resulting hedging strategy provides additional exposure to oil income risk compared with a strategy where the fund only invests in one risky asset. Extensive research has shown that derivative securities can be used to complete financial markets and improve welfare (see Liu and Pan 2003 and Hsuku 2007 and the references therein.).

Table 1. Parameter values.

Panel (a): Estimated parameters, $\boldsymbol{\theta}_1$			
Description	Parameter	MLE	Standard error
Drift of oil price growth	κ_P	0.0973	0.0606
Drift stock price growth	μ	0.0413	0.0264
Diffusion oil price growth	σ_P	0.3522	0.0908
Diffusion stock price growth	σ_S	0.1592	0.0087
Corr. stock price and oil price growth	ρ_{PS}	−0.0676	0.0403

Panel (b): Calibrated parameters, $\boldsymbol{\theta}_2$		
Description	Parameter	Value
Risk-free rate	r	0.011
Discount rate	β	0.020
Risk aversion coefficient	γ	3.000
Elasticity of intertemporal substitution	ψ	2.000
Extraction rate	κ_Q	0.078
Terminal date of income flow (years)	$\hat{T}$	60

Note: Panel (a) reports the maximum-likelihood estimates and associated standard errors for the parameters that describe the dynamics of the exogenous driving forces in the economy. The estimation uses monthly data on U.S. real stock prices and real WTI oil prices for the period 1973–2007. Effective number of observations is 420. Panel (b) report calibrated parameter values that describe the investor's preferences and some additional parameters that replicate salient features of the investment problem faced by the Norwegian SWF.

4.1. Calibration

We separate the parameters of the model into two groups, $\boldsymbol{\Theta} = \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2\}^\top$. The first group includes all the parameters associated with the exogenous processes that drive the dynamics of the stock and oil prices, $\boldsymbol{\theta}_1 = \{\mu, \sigma_S, \rho_{PS}, \kappa_P, \sigma_P\}^\top$. The second group includes those parameters related to the fund's preferences, the return on the market's risk-free asset, and some additional parameters associated to the generation of income, $\boldsymbol{\theta}_2 = \{\beta, \gamma, \psi, r, \kappa_Q, \hat{T}\}^\top$. Table 1 summarizes our benchmark calibration. Time is measured in years and parameters should be interpreted accordingly.

4.1.1. Estimated parameters ($\boldsymbol{\theta}_1$). In the following, we assume that the risky asset is represented by the U.S. stock market, and the oil price is gauged by the West Texas Intermediate (WTI) price of crude oil. To measure stock prices, we use monthly nominal returns on the value-weighted index excluding distribution from CRSP, while the monthly WTI price of oil is retrieved from the FRED database. All prices in the model are real. We use the monthly value of the Consumer Price Index from the U.S. Bureau of Labor Statistics as the deflator. Panel (a) in figure 3 illustrates the monthly year-over-year (yoy) returns for each of the variables for the period 1974:1-2018:12.

Let $\mathbf{X}_t = (S_t, P_t)^\top$. According to (1) and (3), the price dynamics can be represented by the following system of Markovian stochastic differential equations (SDE)

$$\mathrm{d}\vec{X}_t = \mu\left(\vec{X}_t; \boldsymbol{\theta}_1\right) \mathrm{d}t + \sigma\left(\vec{X}_t; \boldsymbol{\theta}_1\right) \mathrm{d}\vec{Z}_t, \tag{24}$$

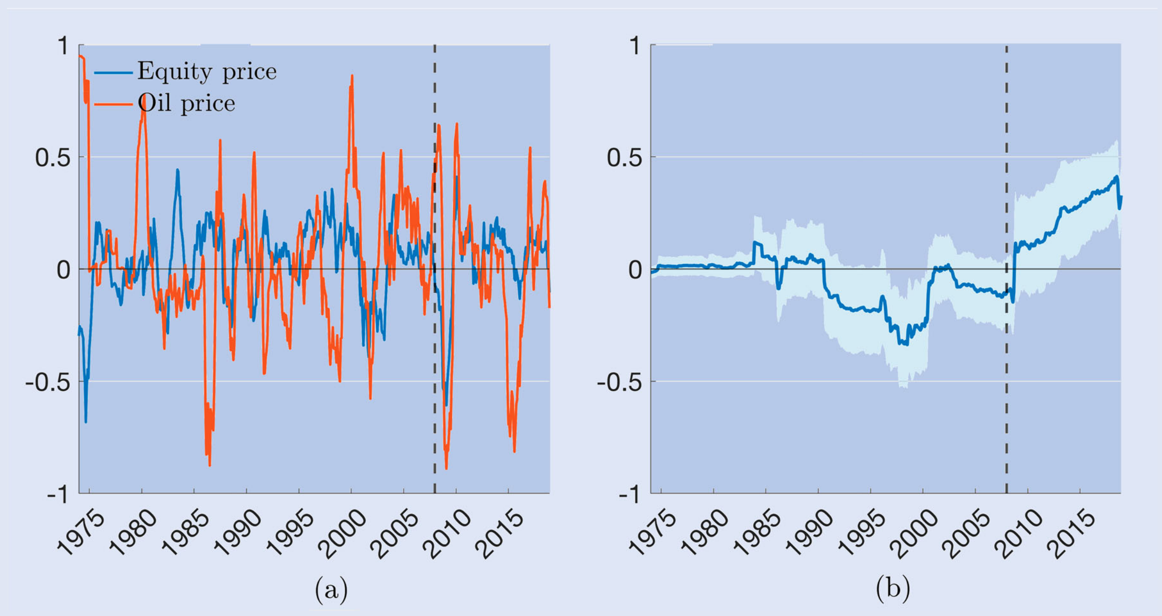

Figure 3. Crude oil and equity prices (1974:1–2018:12). Panel (a) plots the monthly year-over-year (yoy) change in the real price of stocks (VWI CRSP) and crude oil (WTI). Panel (b) plots the estimated instantaneous correlation between the real price of stocks and crude oil using a 10 years estimation rolling window. Dashed areas represent 95% confidence intervals. The dashed vertical line delimits the end of the sample used in the maximum-likelihood estimation (December 2017).

where $\mathbf{Z}_t = (Z_{S,t}, Z_{P,t})^\top$ is a vector of independent standard Brownian motions, and where

$$\mu\left(\vec{X}_t;\boldsymbol{\theta}_1\right) = \begin{bmatrix} \mu S_t \\ \kappa_P P_t \end{bmatrix},$$

$$\sigma\left(\vec{X}_t;\boldsymbol{\theta}_1\right) = \begin{bmatrix} \sigma_S S_t & 0 \\ \sigma_P \rho_{PS} P_t & \sigma_P\sqrt{1-\rho_{PS}^2}P_t \end{bmatrix}.$$

Let $\mathbb{P}(\vec{X}_0, \vec{X}_\Delta, \ldots, \vec{X}_{n\Delta})$ denote the joint density of a sample of n discrete measurements $\vec{X} = \{\vec{X}_{i\Delta}\}_{i=0}^n$, where $\Delta = 1/12$ denotes the fixed (monthly) observation frequency. Using the properties of joint densities, and the Markovian nature of the process in (24), it is possible to decompose $\mathbb{P}(\vec{X}_0, \vec{X}_\Delta, \ldots, \vec{X}_{n\Delta})$ as the product of a conditional and a marginal density

$$\mathbb{P}\left(\vec{X}_0, \vec{X}_\Delta, \ldots, \vec{X}_{n\Delta};\boldsymbol{\theta}_1\right) = \mathbb{P}\left(\vec{X}_0;\boldsymbol{\theta}_1\right)\prod_{i=1}^{n}\mathbb{P}\left(\vec{X}_{i\Delta}\,|\,\vec{X}_{(i-1)\Delta}\right). \tag{25}$$

Ignoring the dependence on the initial observation, $\mathbb{P}(\vec{X}_0;\boldsymbol{\theta}_1)$, and taking logarithms, the log-likelihood of the data reads

$$\mathcal{L}_n\left(\vec{X};\boldsymbol{\theta}_1\right) = \sum_{i=1}^{n}\log\mathbb{P}\left(\vec{X}_{i\Delta}\,|\,\vec{X}_{(i-1)\Delta};\boldsymbol{\theta}_1\right), \tag{26}$$

whilst the maximum-likelihood estimator (MLE) of $\boldsymbol{\theta}_1$ is defined as

$$\hat{\boldsymbol{\theta}}_1 = \underset{\boldsymbol{\theta}_1}{\arg\max}\ \mathcal{L}_n\left(\boldsymbol{\theta}_1;\vec{X}\right). \tag{27}$$

In general, the conditional probability density $\mathbb{P}(\vec{X}_{i\Delta}|\vec{X}_{(i-1)\Delta};\boldsymbol{\theta}_1)$, and hence the log-likelihood function, is not available in closed form. We follow Ait-Sahalia (2002, 2008) and approximate the log-likelihood function in (26) by

$$\mathcal{L}_n\left(\vec{X};\boldsymbol{\theta}_1\right) \approx -\frac{1}{2}\log|\boldsymbol{\Sigma}(\vec{X}_t;\boldsymbol{\theta}_1)| + \mathcal{L}_n\left(\vec{Y};\boldsymbol{\theta}_1\right), \tag{28}$$

where $\boldsymbol{\Sigma} = \sigma\sigma^\top$ is the infinitesimal variance–covariance of the stochastic process $\vec{X}$, and $\mathcal{L}_n(\vec{Y};\boldsymbol{\theta}_1)$ is the first-order closed-form approximation to the log-likelihood function of the transformed unitary diffusion process $\mathbf{Y} = \{\vec{Y}_{i\Delta}\}_{i=0}^n$ given by†

$$\mathcal{L}_n\left(\vec{Y};\boldsymbol{\theta}_1\right) = -\log(2\pi\Delta) + \frac{C_{\vec{Y}}^{(-1)}(\vec{Y}\,|\,\vec{Y}_0;\boldsymbol{\theta}_1)}{\Delta} + C_{\vec{Y}}^{(0)}(\vec{Y}\,|\,\vec{Y}_0;\boldsymbol{\theta}_1) + C_{\vec{Y}}^{(1)}(\vec{Y}\,|\,\vec{Y}_0;\boldsymbol{\theta}_1)\Delta.$$

The approximation constants $C_{\vec{Y}}^{(-1)}$, $C_{\vec{Y}}^{(0)}$ and $C_{\vec{Y}}^{(1)}$ are provided in Ait-Sahalia (2008, Theorem 1). The approximated log-likelihood function in (28) converges to the true log-likelihood function of the data as $\Delta \to 0$, and thus all the standard statistical properties of the (quasi-) MLE, including classical inference, carry over.

Panel (a) in table 1 reports the estimation results for the model in (24) using monthly data that span the period from 1973:1–2007:12. It also presents standard errors robust to autocorrelation and heteroskedasticity. The estimated (annual) volatilities of the stock and oil price changes are 15.92% and 35.22%, respectively, while the instantaneous drift parameters, although not statistically significant at conventional confidence levels, imply an annual expected stock return

† The multivariate diffusion in (24) is reducible in the sense that it is possible to transform the diffusion process $\vec{X}$ into a diffusion process $\vec{Y}$ with diffusion matrix equal to the identity matrix (see Ait-Sahalia 2008).

of 4.13%, and an annual expected growth of oil price of 9.73%. The estimated instantaneous correlation between stock and oil prices indicates a negative and statistically significant association equal to -6.76% (per annum) during this period. Our estimate is consistent with the evidence provided in Jones and Kaul (1996), Sadorsky (1999), Gorton and Rouwenhorst (2006), Lee and Chiou (2011) and Bhardwaj *et al.* (2015) for the period prior to the Great Recession.

Although the negative correlation between oil prices and equity is in line with conventional wisdom, recent evidence suggests that this correlation has increased considerably during the last decade. Buyuksahin and Robe (2014) attribute this increase to the observed growth in commodity-market activity led by hedge funds but also to macroeconomic fundamentals and the TED spread. Lombardi and Ravazzolo (2016) provide further evidence of the higher correlation observed at the onset of the financial crisis of 2008. Additionally, Datta *et al.* (2021) argue that the increase in the oil-stock price correlation that started in 2008 can be explained by the nominal interest rates being constrained by the zero lower bound. To verify this claim, we extend our sample period until the end of 2018 and, using the maximum-likelihood procedure described above, produce rolling estimates using a fixed window of 10 years. Panel (b) in figure 3 reports the rolling estimates together with 95% confidence bands. The results suggest that the long-term co-movement between oil price changes and stock price changes is not constant. Instead, the estimates suggest three different phases over the last 45 years: (i) a period of zero correlation between 1973 and 1989; (ii) a period of negative correlation between 1990 and 2007; and (iii) a period of positive correlation from 2008 until today. We use the variation in these correlations to perform a sensitivity analysis.

4.1.2. Fixed parameters ($\boldsymbol{\theta}_2$). The calibration of the second group of parameters is summarized in Panel (b) of table 1. We set the subjective discount rate to $\beta = 2\%$ per year, the coefficient of RRA to $\gamma = 3.0$ and the EIS to $\psi = 2.0$. These values are standard in the asset pricing and asset allocation literature and, as it is shown below, imply a relatively smooth path for the consumption-to-financial wealth ratio that is consistent with the spending rule mandate of the Norwegian SWF.

The exogenously given income terminal date is calibrated to $\hat{T} = 60$ years to resemble the number of years that it would take to exhaust 99% of the oil reserves available in 2018 under the following two assumptions: (i) zero exploration and discoveries of new reserves, and (ii) a constant extraction rate of $\kappa_Q = 7.8\%$ per year, consistent with the 2018 production-to-reserves ratio reported by BP (2019) which, together with the estimated average growth rate of oil price, implies an expected oil income growth rate of $\kappa = 1.9\%$ per year.

Finally, the return on the risk-free asset is calibrated to the sample average of the annualized real return on the 90-day U.S. Treasury bill. For the postwar period, this corresponds to $r = 1.11\%$ per year. Together with the estimated values for the equity's expected return, μ, and volatility, σ_S, our calibration implies that starting from year $\hat{T} = 60$, the fund's manager should invest 40% of her financial wealth on the risky asset and consume 2.3% of her financial wealth each period in perpetuity.

4.2. *Optimal consumption and investment policies*

In the following, we use the benchmark calibration in table 1 to illustrate the solution of the model when the oil income is not fully spanned by the financial markets. Due to market incompleteness, we approximate the solution numerically for all $t \in [0, \hat{T}]$. We then simulate each of the variables in the model and report the median value over 10,000 samples. In the simulations, we assume an initial financial wealth-to-oil income ratio equal to $W_0/Y_0 = 9.4$, a number that is consistent with that reported by the Norwegian GPFG in 2018.

Figure 4 plots the optimal path for selected variables for all $t \in [0, \hat{T}]$, together with intervals around the median that represent the 15th and 85th percentiles of the simulated distributions. Panel (a) illustrates the optimal portfolio share of financial wealth invested in equity. The results indicate that the fund's manager should initially invest 60% of her financial wealth into stocks, a figure that is 20% above the optimal share that would prevail in the long run. In fact, as $t \to \hat{T}$ the fund's manager should decrease gradually its position in the risky asset until she reaches a long-run allocation of 40% in 60 years.

The excess demand for the risky asset, relative to the long-run position, is the result of two complementary effects. First, a wealth or leverage effect from the capitalized value of all future oil income transferred to the fund, i.e. the non-tradable 'underground' oil wealth available at the beginning of the investment horizon. Second, a positive hedging demand that, given our benchmark calibration, is primarily driven by the high volatility of oil prices, and not by their correlation with the risky asset. Panel (b) plots this hedging component using the decomposition in (18). We see that the intertemporal hedging demand accounts for nearly 12% of the additional financial wealth invested in equity markets at the beginning of the investment horizon. This additional demand creates a hedge against fluctuations in oil prices and thus in the fund's revenues. Similar to the overall investment share in equity, the hedging demand decreases over time hand-in-hand with the oil wealth-to-financial wealth ratio. In the long run, when the fund stops receiving oil revenues, the hedging demand becomes zero.

Our results suggest that the optimal portfolio profile is in sharp contrast with the effective investment strategy executed by the Norwegian GPFG between 1998 and 2007 shown in figure 1. In a period characterized by a low and negative correlation between stock and oil prices, the fund's investment strategy was rather conservative. Although in line with the mandate given by the Executive Board to the NBIM, the share of financial wealth invested in equity remained relatively constant at around 40%.

The median consumption-to-financial wealth ratio is illustrated in Panel (c). Our benchmark calibration produces a relatively stable optimal spending rule over time as a function of the fund's financial wealth. At the beginning of the investment horizon, when the total wealth is large, the optimal consumption is 2.7% of the fund's financial wealth. After 10

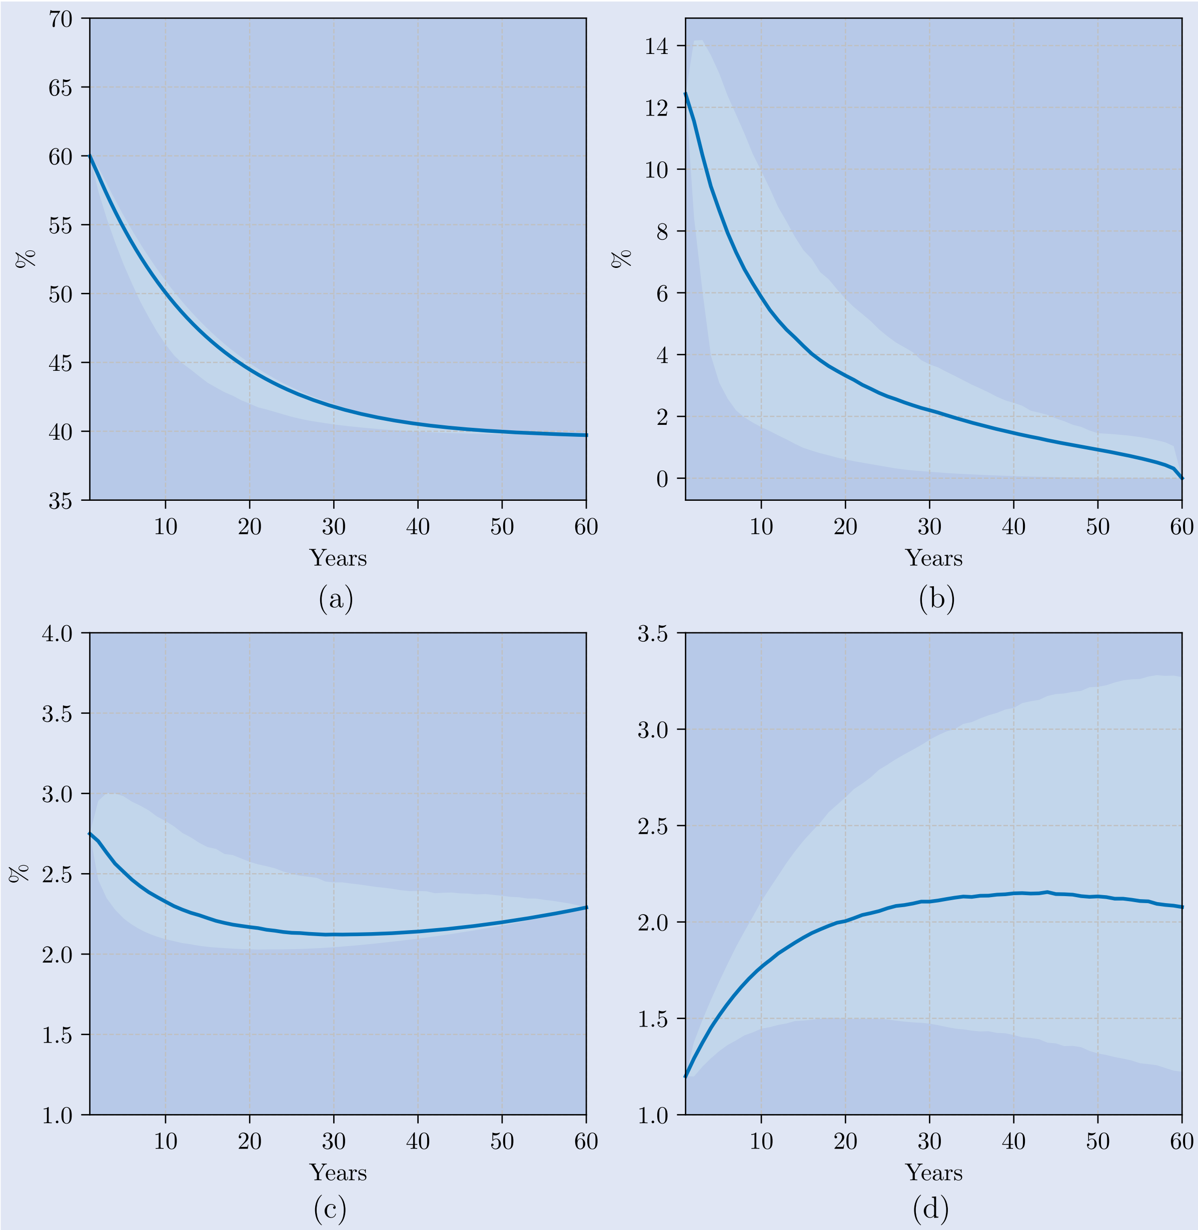

Figure 4. Optimal strategies under incomplete markets. Panels (a)–(d) plot, respectively, the optimal share of financial wealth invested in equity, the hedging demand as a fraction of financial wealth, the optimal consumption-to-financial wealth ratio, and the evolution of financial wealth-to-mainland GDP ratio using the parameters in table 1. The solid lines represent the median value over M =10,000 simulated paths, each of them of $T = 60$ sample points. The shaded areas represent the 15 and 85 percentiles from the sampling distribution of the simulated series.

years, the spending path stabilizes at around 2.2% of the financial wealth for the remaining time horizon, until it reaches its long-run value of 2.3% after 60 years. Our results, although somewhat lower, are consistent with the constant transfer rule described in the investment mandate of the Norwegian GPFG.

Finally, Panel (d) shows the path of the financial wealth-to-mainland GDP ratio over time. In the simulations, we use an initial value of $W_0/GDP_0 = 3.32$, a value that is consistent with the real GDP of mainland Norway for the year 2019. Assuming a constant growth rate for the real GDP of mainland Norway of 1.07% per year (the average growth rate observed over the period 2000-2019), our results suggest that following the optimal consumption-investment policies allows the fund to reach, over the course of 60 years, a level of financial wealth-to-GDP ratio that doubles its initial endowment. This, in turn, implies that the fund's financial wealth will grow at an average annual rate of 2.55% from 9.4 in year zero to 42.36 in year 60. The terminal wealth in period 60 also reflects the optimal transfer of wealth to future generations.

Our results can also be used to argue against the popular advice of using the complete market solution to approximate the optimal consumption and investment strategies of an unconstrained investor with stochastic income when markets are in fact incomplete (see for example Bick *et al.* 2009, Munk and Sørensen 2010, van den Bremer *et al.* 2016). Suppose that the oil income risk is perfectly spanned by the risky asset and use Lemma 3.1 to compute the value of the underground oil wealth at each point in time. Then, using Proposition 3.2, and the observed correlation coefficient between oil and stock price changes, we can approximate the investment and consumption policies†. In general, we find that the use of the complete market solution would command the SWF to invest 535% of its financial wealth at the beginning of the planning horizon, a number that far exceeds the true portfolio share of 60% that prevails under incomplete markets. The associated intertemporal hedging demand component amounts to nearly 140%, when the true excess demand is 12%. Under this scenario, the SWF needs to borrow an unreasonable large amount of funds in order to achieve the optimal portfolio. The reason behind this large demand for the risky asset is that, given the covariance structure between oil and stock price changes, the SWF will overestimate the value of its underground oil wealth when erroneously assuming that markets are complete. This will, in turn, imply that the investor will rely on a wrong measure of its true leverage possibilities, over invest in the risky asset, and expose the fund to unnecessary large amounts of risk. Therefore, our results suggest not to approximate the optimal investment strategy for oil-based SWFs using the assumption of complete markets, since the resulting portfolio weights on the risky asset will be unreasonable high in the absence of liquidity or borrowing constraints. However, as shown in Munk and Sørensen (2010), the accuracy of the approximation remains valid in this case if the volatility of equity returns exceeds that of income‡.

4.3. Parameter sensitivity

In this subsection, we examine the sensitivity of the optimal investment strategy to changes in the correlation between the stock price and the oil price, and the SWF's coefficient of RRA.§

4.3.1. Correlation between asset returns and oil price changes. Motivated by the time-varying estimates reported in figure 3, we ask what the consequences of different values of the correlation between the shocks to the oil price and the stock price changes for the optimal portfolio allocation are. The results are illustrated in figure 5 where we consider different values of ρ_{PS}, all of them consistent with the three different episodes observed between 1974 and 2018. These include periods of high negative and positive correlation, as well as periods of zero correlation.

Panel (a) plots the median share of financial wealth invested in equity across 10,000 simulated paths. It shows that the optimal investment strategy is sensitive to changes in the correlation coefficient. In particular, the portfolio weight on the risky asset is a nonlinear and decreasing function of ρ_{PS}. As discussed previously, a negative correlation between the price of oil and the price of equity commands an initial position in stocks that exceeds its long-run value. Using (22) under complete markets as an approximation, we observe that the larger is the negative association between stock and oil prices, the larger are the gains from hedging, and thus the higher is the optimal demand for equity. In particular, while a correlation of nearly -7% implies an initial allocation of 60%, a correlation of -30% increases this position to over 80%. As shown in Panel (b), the hedging demand in these two examples account, respectively, for 12% and 65% of the initial financial wealth allocated to stocks.

If the price of oil and the price of equity are uncorrelated, then the optimal portfolio share is below 60%. This allocation is completely determined by the speculative demand component since there is no room for hedging. The decreasing path in this scenario is exclusively explained by a falling oil wealth-to-financial wealth ratio.

Finally, a positive correlation of 30% results in a portfolio that invests a large fraction of wealth into the risk-free asset. In particular, the initial fraction of wealth invested in equity is reduced to 30%, a value that represents half of the share under our benchmark calibration. This low but positive fraction invested in the risky asset is the result of a long position that is simultaneously counterbalanced by a short position that aims to minimize the oil income risk via the hedging component. Moreover, the optimal investment rule should increase monotonically over time to reach its long-run value of 40% at terminal time $\hat{T}$. Hence, if the correlation stays positive, as suggested by the recent empirical evidence, the current stock/bond mix in the portfolio of the Norwegian GPFG implies an unnecessary large exposure to risk.

4.3.2. Risk aversion. As discussed previously, the coefficient of RRA, γ, plays an important role in shaping the intertemporal demand for the risky asset. Consequently, panel (a) in figure 6 illustrates the median value of the optimal demand for equity for different values of this parameter across 10,000 simulated paths. In general, larger risk aversion leads the investor to take less risk. While our benchmark calibration with $\gamma = 3$ commands the fund to invest nearly 60% of its financial wealth into the risky asset at the beginning of the planning horizon, this fraction drops to 30% and 21% for coefficients of RRA equal to $\gamma = 6$ and $\gamma = 9$, respectively. Moreover, the uncertainty around the optimal investment strategy, as indicated by the shaded areas, also becomes smaller the lower is the coefficient of RRA. Similar conclusions are reported by Campbell and Viceira (1999) in a dynamic asset allocation model of an infinitely lived

† The corresponding median trajectories from 10,000 simulations of the model together with their 15th and 85th percentiles are reported in Appendix 3.

‡ Using quarterly data on U.S. aggregate income from the National Income and Product Accounts (NIPA) for the period 1951–2003, Munk and Sørensen (2010) estimate a volatility of disposable labor income of 2.08%, and a correlation coefficient between income and equity price changes of 16.73%. Their estimate of the volatility of the S&P500 index is 16.13%. When using aggregate income data from the Panel Study of Income Dynamics (PSID) survey, their estimate of the labor income volatility is 1.64%. On the other hand, and consistent with the evidence reported in Carroll and Samwick (1997) and Chamberlain and Hirano (1999) using PSID data, Viceira (2001) uses a volatility of labor income of 10% in his benchmark calibration. A similar value is estimated in Cocco *et al.* (2005), who additionally estimate a correlation coefficient with equity returns between 0% and -1.75%.

§ We also perform a sensitivity check to explore the role of different initial values of ratio W_0/Y_0 on the optimal allocation. We find that initial values have moderate effects on the optimal equity holding, whereas the impact on the consumption-wealth ratio is more sizable. However, the effects are not quantitatively important. The results are available upon request.

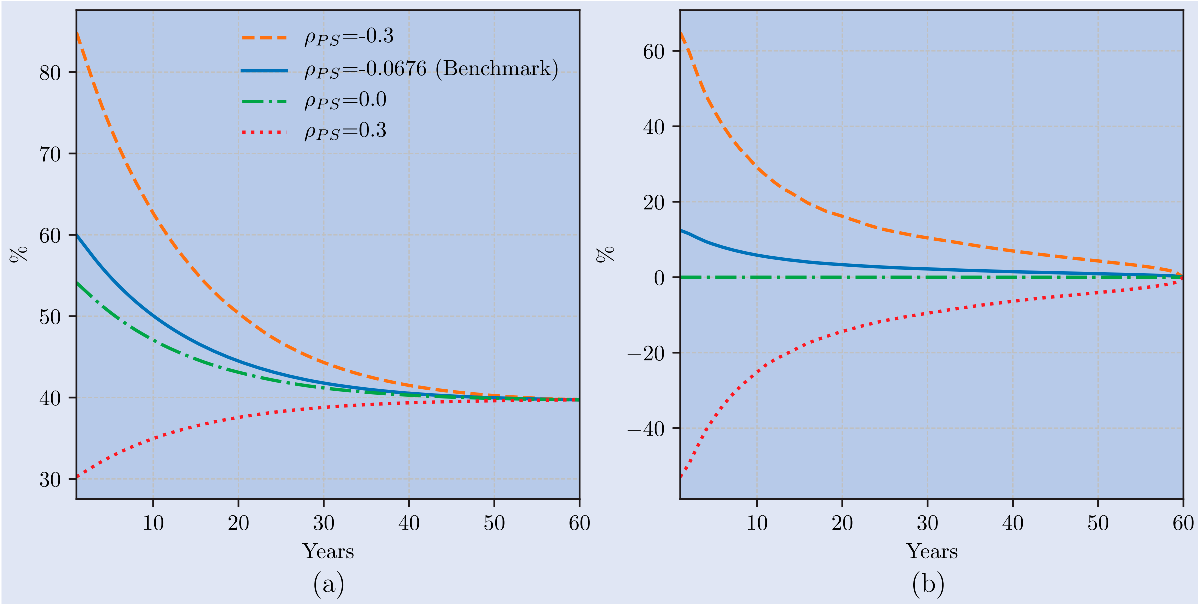

Figure 5. Sensitivity of the optimal investment share to the correlation coefficient, ρ_{PS}. Panel (a) plots the median value of the fraction of financial wealth invested into equity. The median is computed from M =10,000 simulated paths of the model for different values of the instantaneous correlation between real asset returns and changes in the real price of oil, ρ_{PS}. Panel (b) plots the associated median hedging demand.

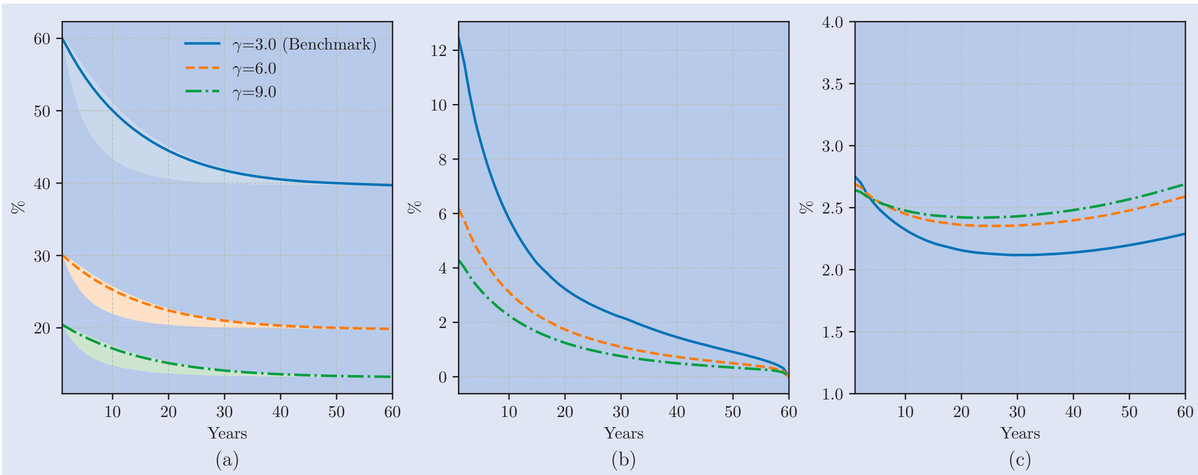

Figure 6. Sensitivity of optimal allocations to the risk aversion coefficient. Panels (a)–(c) plot, respectively, the median value of the investment share on equity, of the hedging demand, and of the consumption-to-financial wealth ratio over M =10,000 simulated paths of the model for different values in γ.

investor without stochastic income but with a time-varying equity premium.

As $t \to \hat{T}$, the allocation on the risky asset converges to an investment share that remains constant through time. As shown in (14), this stationary level decreases with the value of the coefficient of RRA. In particular, the long-run optimal allocation on the risky asset is nearly 20% for $\gamma = 6$, and around 13% for $\gamma = 9$. These values represent a sizable correction relative to the 40% share implied by our benchmark scenario.

Panel (b) plots the median value of the hedging demand component associated with the total demand in Panel (a). In general, we find that the larger is the investor's risk aversion, the more conservative she is to hedge against the oil income risk, thus the less is the hedging component. This is in line with the intuition obtained in the complete market case. As shown in equation (22), when the oil income and stock price are negatively correlated, then the larger is the coefficient of RRA, the smaller is the hedging component.

Finally, in Panel (c), we report the effects of different values of the coefficient of RRA on the dynamics of the median optimal consumption-to-financial wealth ratio. In general, we find that the impact of different levels of risk aversion on the consumption ratio is less pronounced than for equity holdings. In other words, changes in the investor's risk aversion have a larger effect on the manager's asset allocation than on the consumption choice over time. For $\gamma \in (3.0, 6.0, 9.0)$, the optimal ratio is very smooth over time, and it fluctuates between 2.0% and 3.0%. Given the value of the EIS in our benchmark calibration ($\psi = 2.0$), we find that the optimal consumption ratio declines with γ. As an example, consider year $\hat{T} = 60$. The median optimal consumption is 2.3% of the financial wealth for an investor with $\gamma = 3.0$, whereas it increases to 2.6% for an investor with $\gamma = 9.0$. As shown in Campbell and Viceira (1999) this is not necessarily always the case as one should expect to see an opposite relation between the optimal consumption ratio and the coefficient of RRA for investors that are extremely reluctant to substitute consumption across periods and hence have low values of the EIS.

4.4. The effect of alternative assumptions on $\hat{T}$

The results described so far assume that the flow of income ends with certainty after 60 years. Here, we investigate how the optimal consumption and investment strategies change with different assumptions on the terminal income date $\hat{T}$. First, we study how sensitive the optimal allocations are to alternative, but still fixed, values of $\hat{T}$. Second, we assume that the terminal date is unknown to the fund's manager and study its implication for optimal equity holding. The results under the assumption of incomplete markets are summarized in figure 7 where we plot the optimal share of financial wealth invested in the risky asset and the optimal consumption-wealth ratio, both at the beginning of the planning horizon, for different terminal dates. The solid lines report the optimal strategies for distinct terminal dates $\hat{T}$, each known with certainty by the manager, while the dashed lines represent the optimal policies for different *expected* terminal dates. All the remaining parameters are fixed to those in table 1.

4.4.1. Predetermined terminal date. Under the maintained assumption that the fund's manager knows with certainty the date in which it will stop receiving income, we assess the effects of different terminal dates on the optimal strategies. In particular, we solve the second stage for $\hat{T} \in [10, 30, 60, 100, 200, 300]$ years.

Panel (a) in figure 7 shows a positive relationship between the initial optimal investment share and the terminal date. Interestingly, we find that as we increase $\hat{T}$, the fraction of financial wealth invested in equity increases monotonically to a value of 60%. On the other hand, panel (b) shows that the initial consumption rate decreases monotonically to a level of 2.5% of the initial financial wealth. The positive relation between α_0 and $\hat{T}$ is partly explained by the larger oil wealth that results from accumulating oil income over a longer time horizon. This effect is reinforced by the negative correlation between oil price and stock price changes, ρ_{PS}. On the other hand, the negative relation between C_0/W_0 and $\hat{T}$ is the result of two opposing effects: a positive income effect from the larger accumulation of oil wealth mentioned before, and a negative effect from a marginal propensity to consume out of financial wealth that is decreasing in the terminal date. Under our calibration, the latter effect dominates. Finally, notice that as we increase $\hat{T}$, the relative importance of the terminal utility $J(W_{\hat{T}})$ decreases due to the rate of discount.

4.4.2. Random terminal date. The assumption of a certain terminal date might seem unrealistic. Therefore, we now consider the case where $\hat{T} \in [0, \infty)$ is instead a random variable that denotes the time at which the fund will stop receiving oil revenues. Let λ denote the conditional arrival rate of an event that forces the fund's owner to interrupt the flow of income to the portfolio manager at time $t > \hat{T}$. Hence the expected terminal date is given by $\mathbb{E}[\hat{T}] = 1/\lambda$. Appendix 4 shows that, under the assumption that the variation in the terminal date is uncorrelated with the shocks to oil income and financial returns, the manager's allocation problem can be written as

$$J_0 = \max_{\{C,\alpha\}\in\mathcal{A}} \mathbb{E}_0 \left[\int_0^\infty e^{-\lambda t} \left(f(C_t, J_t) + \lambda A \frac{W_t^{1-\gamma}}{1-\gamma} \right) \mathrm{d}t \right] \tag{29}$$

where $A := \beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}}$, with $\mathbb{G}_\infty$ given in (13). Note that as $\lambda \to 0$, the fund is more likely to receive commodity income for a longer time horizon on average. The HJB equation for this problem is given by

$$\begin{aligned} 0 = {} & -\lambda J + u(C_t, W_t, J_t) + J_W\,(rW + \alpha\,(\mu - r)\,W + Y - C) \\ & + \frac{1}{2} J_{WW} \sigma_S^2\,(\alpha W)^2 + \kappa Y J_Y + \frac{1}{2}\sigma_P^2 Y^2 J_{YY} \\ & + \sigma_S \sigma_P \rho_{PS}\,(\alpha W)\, Y J_{WY}, \end{aligned} \tag{30}$$

with $u(C_t, W_t, J_t) := f(C_t, J_t) + \lambda A \frac{W_t^{1-\gamma}}{1-\gamma}$. We approximate the solution to (30) using an infinite horizon version of the finite difference scheme described earlier.

Panel (a) in figure 7 plots the initial investment share for different expected terminal dates $\mathbb{E}[\hat{T}] \in [10, 30, 60, 100, 200, 300]$ years. As for the case of a known terminal date, the investment share becomes larger the higher is the expected terminal date. In fact, α_0 converges monotonically to a value of around 50% when the expected terminal date is 300 years. More interestingly, the figure shows the impact of uncertainty on the investment share at the beginning of the planning horizon. For example, consider the case where $\mathbb{E}[\hat{T}] = 60$ years. The optimal initial investment strategy is around 48%, below the optimal strategy that would prevail if the fund's manager knew with certainty that the terminal date was $\hat{T} = 60$ years. The lower share is due to the effects of uncertainty on the terminal date which, by assumption, cannot be hedged by the fund. Panel (b) plots the effects of different expected terminal dates on the initial consumption-wealth ratio. Contrary to the case of the certain terminal date, we obtain a monotonically increasing initial consumption per unit of financial wealth that is the result of two complementary factors. First, an increase in the precautionary savings behavior that reduces the initial optimal consumption as a response to the undiversifiable risk associated with an unknown terminal date of the income flow that is potentially close to the beginning of the planning horizon. Second, a discounting effect that weakens as the expected terminal date increases meaning that consumption decisions associated with lower expected terminal dates will be also more affected by the uncertainty around the terminal date.

5. Welfare costs of suboptimal investment rules

So far, we have shown how the portfolio manager of an SWF should optimally allocate its wealth among securities, and

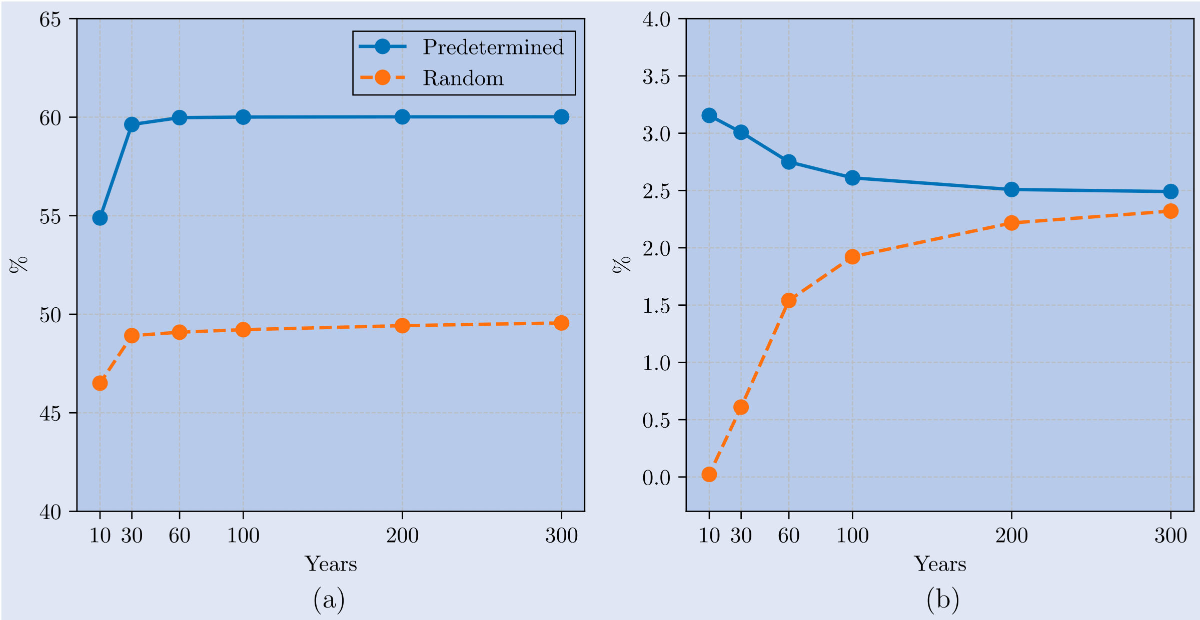

Figure 7. Optimal initial strategies for different assumptions on $\hat{T}$. Panels (a) and (b) plot, respectively, the optimal initial share of financial wealth invested in equity and the optimal initial consumption-wealth ratio for different predetermined and certain values of the terminal date $\hat{T}$ (solid lines), and different expected terminal dates (dashed lines) determined by the hazard rate λ.

how to use the equity markets to hedge against fluctuations in the commodity income. In particular, the optimal investment policy implies a portfolio mix that is time-dependent. Its evolution over time is determined by the path followed by the oil wealth-to-financial wealth ratio, the coefficient of RRA, and the correlation between the commodity income and the asset returns. However, the investment strategy followed by most commodity-based SWFs is exogenously given by the fund's owner. More specifically, the fund's exposure to risk is determined by a long-term investment mandate that usually recommends to hold a relatively constant position in equity without necessarily timing the market, as opposed to the strategy that would otherwise maximize welfare. To shed light about the potential costs of following policies that are suboptimal, this section calculates the welfare cost of following alternative policy rules. To this end, we ask what the amount of additional initial financial wealth that needs to be transferred to the fund at the beginning of the planning horizon is so that the implementation of a suboptimal policy provides the same level of utility that could be achieved otherwise with the optimal rule (see Cochrane 1989). In other words, we compute the wealth compensating variation τ_0 that yields

$$J(t_0, W_0, Y_0) = \widetilde{J}(t_0, (1+\tau_0)W_0, Y_0), \tag{31}$$

where J is the indirect utility that solves (8), and $\widetilde{J}$ is the value function that results from following an investment strategy different from that implied by (18), but where consumption is allowed to adjust optimally. Note that τ_0 can alternatively be interpreted as a wealth-equivalent *loss* for the owner of the SWF since in the absence of compensation the fund's (indirect) utility from following a suboptimal policy will be lower. If the wealth compensation τ_0 is small, the welfare gains of implementing the maximizing investment rule can be probably outweighed by some of the features our model extracts from, e.g. transaction costs related to portfolio rebalancing, stochastic investment opportunities, etc. Thus, our stylized framework provides a lower bound on the welfare costs of following suboptimal investment policies.

In what follows, we consider two different suboptimal investment strategies using the benchmark model where the terminal date is known a priori by the fund's manager. The first policy fixes the fraction of wealth invested in the risky asset to $\widetilde{\alpha}_t = 70\%$ for all $t \leq \hat{T}$. This fixed rule echoes the investment mandate given by the Norwegian GPFG to the NBIM according to which the allocation on equity should amount to 60–80% of the total portfolio. The second rule assumes instead that every period the fund invests a fraction of its financial wealth into equity that is equal to the median investment share recommended by the optimal asset allocation model across M =10,000 simulations, i.e. $\widetilde{\alpha}_t = \text{Median}(\alpha_t^1, \ldots, \alpha_t^M)$. In this case, the investment rule is no longer constant but time-varying. We regard this rule as a near-optimal policy in the sense that it corresponds to a perturbed version of the optimal strategy and therefore can be used to study the welfare costs of small misspecifications in the dynamic asset allocation model or the use of inaccurate parameter values.

Table 2 summarizes our findings, where we report the wealth-equivalent compensation required at the beginning of the planning horizon for different values of coefficient of RRA, γ, and different values of the correlation coefficient between stock and oil price changes, ρ_{PS}. All the values are measured as a percentage of the fund's initial endowment of financial wealth.

Panel (a) shows the results when the SWF follows the constant investment rule of 70%, regardless of the value of the coefficient of RRA and of the correlation coefficient. This strategy implies that the SWF does not time the market, and instead assumes that holding a constant position in equity

Table 2. Wealth-equivalent compensation under suboptimal portfolio rules.

	(a) Constant rule				(b) Median rule			
	ρ_{PS}				ρ_{PS}			
γ	−0.3	−0.07	0	0.3	−0.3	−0.07	0	0.3
2	23.82	6.42	4.14	1.41	23.38	5.82	3.51	0.04
3	17.85	12.49	12.33	16.29	9.77	2.01	1.02	0.44
4	26.17	26.64	27.76	36.33	5.72	1.00	0.43	0.64
5	38.22	43.32	45.58	58.47	3.94	0.60	0.22	0.72
6	51.83	61.32	64.71	81.98	2.98	0.41	0.13	0.74

Panel (a) reports the wealth compensation required (% of initial financial wealth) when following a suboptimal portfolio rule that is constant and equal to 70% over time, while consumption adjusts optimally. Panel (b) reports the wealth compensation required (% of initial financial wealth) when following an ad-hoc suboptimal portfolio rule that in each period fixes the investment share equal to the median value of the optimal portfolio rule, while consumption adjusts optimally.

will imply a reduced level of risk for the overall portfolio over the long run (see Siegel 2014). However, our results suggest that such a strategy can lead to large wealth losses. As an example, for our benchmark calibration with $\gamma = 3$ and $\rho_{PS} = -0.07$, the wealth-equivalent loss is equivalent to 12.5% of the initial financial wealth. In other words, for the suboptimal policy to deliver the same level of welfare that can be obtained using the optimal investment policy, the fund will require an injection of capital equivalent to 12.5% of the initial endowment. To understand why the magnitude of the loss, recall that under the optimal strategy, the fund should initially invest 60% of its financial wealth and thereafter decrease this fraction over time to reach a long-run value of around 40%. On the contrary, when following the suboptimal policy, the fund invests instead 70% period-by-period, a value that exceeds the optimal allocation at every point in time. Hence, the wealth-equivalent loss in the benchmark scenario reflects the excessive exposure to risk implied by the suboptimal policy, an exposition that becomes larger the closer the fund is to stop receiving revenues from commodity-related activities.

Our results also suggest that the welfare losses tend to be substantial for moderate to highly risk-averse investors, particularly when the absolute value of the correlation between oil and stock price changes is large. For example, for a correlation of 30% which resembles that observed in during the last decade, the losses that result from implementing a constant investment rule of 70% in equity over time fluctuate between 16% and 82% of the initial endowment of the fund. Although smaller in magnitude, large losses are also incurred for large and negative correlation coefficients.

In Panel (b), we report the wealth-equivalent losses incurred by the fund when following a suboptimal, but time-dependent, investment rule corresponding to the median of the optimal share obtained from repeated simulations of the model. In contrast to the previous rule, this alternative fixed rule is now a function of the investor's coefficient of RRA and the correlation coefficient between oil and stock price changes. Our results indicate that using a time-dependent rule, whose path is allowed to adjust to the investor preferences and the market interdependencies, leads to considerably smaller losses. For our benchmark calibration, the loss from not implementing the optimal strategy is equivalent to 2% of the fund's initial endowment. Interestingly, if $\rho_{PS} \leq 0$ the loss becomes smaller the more risk averse is the investor, a pattern that is in contrast to that documented in Panel (a). This inverse relation is partly explained by the reduced variability in the optimal investment share (see figure 6) that accompanies the lower median equity holdings of highly risk-averse agents. However, in the case of a positive correlation, we find that the losses, although small, increase with the coefficient of RRA. For $\rho_{PS} = 30\%$ the losses never exceed 1%, which makes the use of a time-dependent fixed rule an attractive alternative in case that the optimal rule is not readily available to the fund. A similar conclusion can be drawn for small negative correlation coefficients and relatively high coefficients of RRA.

6. Conclusions

In this paper, we have extended the standard dynamic asset allocation problem for long-term investors with stochastic income to accommodate the portfolio problem faced by commodity-based SWFs. In particular, we study the optimal consumption-investment decision of a SWF whose primary source of income comes from oil-related activities. Fluctuations in the income stream are assumed to be primarily driven by variations in the exogenous and volatile price of oil. The model features Epstein-Zin-Weil recursive preferences that conveniently separate risk aversion from the elasticity of intertemporal substitution. Since most SWFs are set up by countries interested in sustaining a standard of living for all future generations, we assume that the fund's planning horizon is infinite. However, the oil revenues are received only for a fixed number of periods.

Using data on the S&P500 index and the WTI price for crude oil for the period 1973–2019, we find statistical evidence of a time-varying, but imperfect, correlation between the risky asset and the commodity price. This suggests that the income risk faced by the SWF cannot be perfectly replicated by a trading strategy in the financial markets. More specifically, we find that the average correlation for the period prior to 2007 was −7%. Interestingly, the direction and magnitude of this correlation changed considerably with the start of the Great Recession, and by the end of 2018, it had reached a positive value of 30%. We solve the SWF's manager allocation problem under the assumption of incomplete markets. The remaining parameters are chosen to replicate some salient features of the Government Pension Fund Global (GPFG) in Norway. We find that the fund should initially allocate around 60% of its financial wealth into risky assets and thereafter decrease its position gradually until it reaches a long-run share of 40% over a period of 60 years. The implied intertemporal hedging demand component is found to be large at the beginning of the planning horizon accounting for 20% of the total demand for risky assets.

In contrast with our findings, the fraction of financial wealth invested by the GFPG in equity has exhibited an upward trend since its inception: 40% in 1998 to 70% in 2019. Given the imperfect, but negative, correlation between stock prices and oil income, our results suggest that the GPFG has followed a suboptimal investment strategy that has not exploited all the available hedging opportunities, and thus has taken larger and unnecessary amounts of risk. If instead we consider a positive correlation, similar to that observed after the Great Recession, we find an initial allocation of financial wealth into the risky asset of around 30% that should gradually increase towards its long-run share of 40%. Although in this case, the model implies an increasing investment profile, the optimal portfolio share in equity is substantially lower than any of the allocations reported by the Norwegian SWF during the last 20 years, and thus, also suggests a suboptimal use of the hedging possibilities. In a world where the positive correlation is becoming stronger, SWFs should therefore reduce their exposure to risk by moving away from stocks in the near future.

We also studied the welfare costs of not following an investment strategy that optimally exploits the intertemporal hedging opportunities available to the SWF. To do so, we use a measure of wealth-equivalent welfare compensation. We find significant welfare costs for a SWF that follows a constant investment policy with a fixed equity/bond mix equal to 70/30%. We also find that these losses are considerably reduced if instead the SWF implements a time-dependent, but ad-hoc, investment policy that although suboptimal resembles the decreasing path of the oil wealth-to-financial wealth ratio. For practical purposes, this alternative policy may be considered as a second-best policy in an environment with institutional constraints that prevent the SWF's manager to hedge oil price fluctuations periodically.

Our study offers a unified framework that can be used by commodity-based SWFs to design optimal investment policies that are consistent with the long-term objective of ensuring a smooth intergenerational consumption, and the short-run objective of shielding the economy from fluctuations in the volatile oil revenues. Given the importance of the covariance structure between oil prices and asset returns, future work should further investigate the time-varying nature of the correlation between tradable and non-tradable assets, and the inclusion of more asset classes with time-varying risk premiums.

Acknowledgements

We appreciate the comments and suggestions from two anonymous referees, as well as from Juan Pablo Nicolini, Knut Einar Rosendahl, Olvar Bergland, Camilo Castillo, Olesia Verchenco, and seminar participants at Norges Bank, Norwegian University of Life Sciences, Central Bank of Chile and the Ministry of Finance of Norway. We also thank Yacine Ait-Sahalia for providing us with Matlab code to perform the maximum-likelihood estimation.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Lin Ma http://orcid.org/0000-0001-8414-7744
Juan Carlos Parra-Alvarez http://orcid.org/0000-0001-9076-9021

References

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L. and Moll, B., Income and wealth distribution in macroeconomics: A continuous-time approach. *Rev. Econ. Stud.*, 2022, **89**, 45–86.
Ait-Sahalia, Y., Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica*, 2002, **70**, 223–262.
Ait-Sahalia, Y., Closed-form likelihood expansions for multivariate diffusions. *Ann. Statist.*, 2008, **36**, 906–937.
Bellman, R., *Dynamic Programming*, 1957 (Princeton University Press: Princeton, NJ).
Bernanke, B., The relationship between stocks and oil prices. The Brookings Institution, 2016. Retrieved from https://www.brookings.edu/blog/ben-bernanke.
Bhardwaj, G., Gorton, G. and Rouwenhorst, G., Facts and fantasies about commodity futures ten years later. NBER Working Papers, National Bureau of Economic Research, Inc., 2015.
Bick, B., Kraft, H. and Munk, C., Investment, income, and incompleteness. Working Paper Series: Finance and Accounting Nr. 202., 2009.
Bick, B., Kraft, H. and Munk, C., Solving constrained consumption-investment problems by simulation of artificial market strategies. *Manage. Sci.*, 2013, **59**, 485–503.
Bodie, Z., Merton, R.C. and Samuelson, W.F., Labor supply flexibility and portfolio choice in a life cycle model. *J. Econ. Dyn. Control*, 1992, **16**, 427–449.
Bosserhoff, F., Chen, A., Sørensen, N. and Stadje, M., On the investment strategies in occupational pension plans. *Quant. Finance*, 2022, **22**, 889–905.
BP, BP Statistical Review of World Energy 2019. Retrieved from https://www.bp.com/en/global/corporate/energy-economics/statistical-review-of-world-energy.html.
Buyuksahin, B. and Robe, M., Speculators, commodities and cross-market linkages. *J. Int. Money. Finance.*, 2014, **42**, 38–70.
Campbell, J.Y., Understanding risk and return. *J. Polit. Econ.*, 1996, **104**, 298–345.
Campbell, J.Y. and Viceira, L.M., Consumption and portfolio decisions when expected returns are time varying. *Q. J. Econ.*, 1999, **114**, 433–495.
Campbell, J.Y. and Viceira, L.M., *Strategic Asset Allocation: Portfolio Choice for Long-term Investors*, 2002 (Oxford University Press: Oxford).
Candler, G., Finite difference methods for continuous-time dynamic programming. In *Computational Methods for the Study of*

Dynamic Economies, edited by R. Marimon and A. Scott, 1999 (Oxford University Press: Oxford).
Carroll, C. and Samwick, A., The nature of precautionary wealth. *J. Monet. Econ.*, 1997, **40**, 41–71.
Chamberlain, G. and Hirano, K., Predictive distributions based on longitudinal earnings data. *Ann. Econ. Stat.*, 1999, 211–242.
Cocco, J.F., Gomes, F.J. and Maenhout, P.J., Consumption and portfolio choice over the life cycle. *Rev. Financ. Stud.*, 2005, **18**, 491–533.
Cochrane, J.H., The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives. *American Econ. Rev.*, 1989, **79**, 319–337.
Datta, D.D., Johannsen, B.K., Kwon, H. and Vigfusson, R.J., Oil, equities, and the zero lower bound. *American Econ. J.: Macroeconom.*, 2021, **13**, 214–253.
Davis, S. and Willen, P., Using financial assets to hedge labor income risks: Estimating the benefits. Working paper, 2000.
Deaton, A. and Laroque, G., On the behaviour of commodity prices. *Rev. Econ. Stud.*, 1992, **59**, 1–23.
Døskeland, T., Strategic asset allocation for a country: The Norwegian case. *Financ. Markets Portfolio Manag.*, 2007, **21**, 167–201.
Duffie, D. and Epstein, L.G., Asset pricing with stochastic differential utility. *Rev. Financ. Stud.*, 1992a, **5**, 411–436.
Duffie, D. and Epstein, L.G., Stochastic differential utility. *Econometrica*, 1992b, **60**, 353–394.
Duffie, D., Fleming, W., Soner, H.M. and Zariphopoulou, T., Hedging in incomplete markets with HARA utility. *J. Econ. Dyn. Control*, 1997, **21**, 753–782.
Epstein, L. and Zin, S., Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 1989, **57**, 937–969.
Fages, R., Heredia, L., Carrubba, J., Eyal, O., Frankle, D., Kurihara, K., Mace, B., Palmisani, E., Pardasani, N., Schulte, T., Sheridan, B. and Xu, Q., Global asset management 2019: Will These '20s Roar?' 2019. Retrieved from https://www.bcg.com/publications/2019.
Filis, G., Degiannakis, S. and Floros, C., Dynamic correlation between stock market and oil prices: The case of oil-importing and oil-exporting countries. *Int. Rev. Financ. Anal.*, 2011, **20**, 152–164.
Gomes, F.J., Exploiting short-run predictability. *J. Bank. Finance*, 2007, **31**, 1427–1440.
Gomez, M., Solving PDEs associated with economic models. Technical Report, 2019. Available at https://github.com/matthieugomez/EconPDEs.jl.
Gorton, G. and Rouwenhorst, K.G., Facts and fantasies about commodity futures. *Financ. Anal. J.*, 2006, **62**, 47–68.
Heaton, J. and Lucas, D., Market frictions, savings behavior, and portfolio choice. *Macroecon. Dyn.*, 1997, **1**, 76–101.
Hsuku, Y.-H., Dynamic consumption and asset allocation with derivative securities. *Quant. Finance*, 2007, **7**, 137–149.
Jones, C.M. and Kaul, G., Oil and the stock markets. *J. Finance*, 1996, **51**, 463–491.
Kreps, D. and Porteus, E.L., Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 1978, **46**, 185–200.
Larsen, L.S. and Munk, C., The costs of suboptimal dynamic asset allocation: General results and applications to interest rate risk, stock volatility risk, and growth/value tilts. *J. Econ. Dyn. Control*, 2012, **36**, 266–293.
Lee, Y.-H. and Chiou, J.-S., Oil sensitivity and its asymmetric impact on the stock market. *Energy*, 2011, **36**, 168–174.
Lindset, S. and Mork, K.A., Risk taking and fiscal smoothing with sovereign wealth funds in advanced economies. *Int. J. Financ. Stud.*, 2019, **7**, 1–24.
Liu, J. and Pan, J., Dynamic derivative strategies. *J. Financ. Econ.*, 2003, **69**, 401–430.
Lombardi, M. and Ravazzolo, F., On the correlation between commodity and equity returns: Implications for portfolio allocation. *J. Commodity Markets*, 2016, **2**, 45–57.
Markowitz, H., Portfolio selection. *J. Finance*, 1952, **7**, 77–91.
Merton, R., Lifetime portfolio selection under uncertainty: The continuous-time case. *Rev. Econ. Stat.*, 1969, **51**, 247–257.
Merton, R., An intertemporal capital asset pricing model. *Econometrica*, 1973, **41**, 867–887.
Merton, R.C., Optimum consumption and portfolio rules in a continuous-time model. *J. Econ. Theory.*, 1971, **3**, 373–413.
Miceli, V., Do sovereign wealth funds herd in equity markets? *Quant. Finance*, 2013, **13**, 1503–1518.
Munk, C. and Sørensen, C., Dynamic asset allocation with stochastic income and interest rates. *J. Financ. Econ.*, 2010, **96**, 433–462.
NBIM, Discussion Note 04/16: Asset allocation with government revenues and spending commitments, 2016. Retrieved from https://www.nbim.no.
Norges Bank Investment Management, Norges bank investment management performance results: GIPS Report 31 December 1997 through 31 December 2021, 2022. https://www.nbim.no/en/publications/reports/2021/gips-report/.
Sadorsky, P., Oil price shocks and stock market activity. *Energy Econom.*, 1999, **21**, 449–469.
Scherer, B., Portfolio choice for oil-based sovereign wealth funds. *J. Alternative Investments*, 2011, **13**, 231–246.
Siegel, J., *Stocks for the Long Run*, 5th ed., 2014 (McGraw-Hill Education: New York).
Statista, Statista dossier on hedge funds, 2019. Retrieved from https://www.statista.com/topics/5064/hedge-funds/.
Svensson, L., Portfolio choice with non-expected utility in continuous time. *Econ. Lett.*, 1989, **30**, 313–317.
SWF Institute, Sovereign Wealth Fund Rankings, 2020. Retrieved from https://www.swfinstitute.org.
van den Bremer, T., Frederick, V.D.P. and Wills, S., The elephant in the ground: Managing oil and sovereign wealth. *Eur. Econ. Rev.*, 2016, **82**, 113–131.
van den Bremer, T.S. and der Ploeg, F. van, Managing and harnessing volatile oil windfalls. *IMF Econom. Rev.*, 2013, **61**, 130–167.
Viceira, L.M., Optimal portfolio choice for long-horizon investors with nontradable labor income. *J. Finance.*, 2001, **56**, 433–470.
Wang, C., Wang, N. and Yang, J., Optimal consumption and savings with stochastic income and recursive utility. *J. Econ. Theory.*, 2016, **165**, 292–331.
Weil, P., Nonexpected utility in macroeconomics. *Q. J. Econ.*, 1990, **105**, 29–42.

Appendices

Appendix 1. Optimal allocation: first stage solution

Proof of Proposition 3.1 Following Campbell and Viceira (2002), the optimal allocation problem faced by the fund's manager for $t > \hat{T}$ when $Y_t = 0$ is given by

$$J\left(W_{\hat{T}}\right) = \max_{(C,\alpha)\in\mathcal{A}} \mathbb{E}_{\hat{T}}\left[\int_{\hat{T}}^{\infty} f\left(C_t, J\left(W_t\right)\right) \mathrm{d}t\right]$$

subject to (7), and where the aggregator $f(C, J)$ is given by (6). A necessary condition for optimality for any $t > \hat{T}$ is given by the Hamilton–Jacobi–Bellman (HJB) equation

$$0 = \max_{\{C,\alpha\}\in\mathcal{A}} \left\{f\left(C, J\left(W\right)\right) + \frac{1}{\mathrm{d}t}\mathbb{E}_t\left[\mathrm{d}J\left(W\right)\right]\right\}.$$

An application of Itô's lemma implies

$$\mathrm{d}J\left(W\right) = J_W\,\mathrm{d}W + \frac{1}{2}J_{WW}\left(\sigma_S\alpha W\right)^2\,\mathrm{d}t,$$

where $J_W := \partial J(W)/\partial W$ and $J_{WW} := \partial^2 J(W)/\partial W^2$. Using the martingale difference properties of stochastic integrals, we arrive at

$$0 = \max_{\{C,\alpha\}\in\mathcal{A}} \Big\{ f(C, J) + J_W [rW + (\mu - r)\alpha W - C] + \frac{1}{2}\sigma_S^2 J_{WW} (\alpha W)^2 \Big\}. \tag{A1}$$

The first-order conditions for an interior solution read

$$C^\star = \left(\frac{\beta}{J_W}\right)^{\psi} [(1-\gamma) J]^{\frac{1-\psi\gamma}{1-\gamma}} \tag{A2}$$

$$\alpha^\star = \frac{1}{-\frac{W_t J_{WW}}{J_W}} \frac{\mu - r}{\sigma_S^2}. \tag{A3}$$

By substituting (A2) and (A3) in (A1) we arrive to the maximized HJB equation

$$0 = f\left(C^\star, J\right) + J_W \left[rW + (\mu - r)\alpha^\star W - C^\star\right] + \frac{1}{2}\sigma_S^2 J_{WW} \left(\alpha^\star W\right)^2, \tag{A4}$$

which corresponds to a nonlinear partial difference equation in $J(W)$. We conjecture that a solution to (A4) is given by

$$J(W) = \frac{\beta^\theta}{1-\gamma} \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W^{1-\gamma}, \tag{A5}$$

where $\mathbb{G}_\infty$ is an unknown constant to be determined. Our conjecture implies that

$$J_W = \beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W^{-\gamma}, \quad \text{and} \quad J_{WW} = -\gamma \beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W^{-\gamma-1}. \tag{A6}$$

Substituting (A5) and (A6) into (A4) yields

$$0 = \frac{\beta\psi}{\psi - 1} \beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W^{1-\gamma} \left\{\beta^{-1}\mathbb{G}_\infty - 1\right\} + \beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W^{1-\gamma} \left[r + \alpha^\star(\mu - r) - \frac{C^\star}{W}\right] - \frac{1}{2}\gamma \beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W^{1-\gamma} \left(\alpha^\star\right)^2,$$

where

$$\frac{C^\star}{W} = \mathbb{G}_\infty, \tag{A7}$$

$$\alpha^\star = \frac{1}{\gamma}\frac{\mu - r}{\sigma_S^2}. \tag{A8}$$

Dividing both sides by $\beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W^{1-\gamma}$ and solving for $\mathbb{G}_\infty$ yields (13) which confirms our conjecture. Equations (A7) and (A8) show that both the consumption-to-financial wealth ratio and the share of financial wealth invested in the risky asset are constant for all $t > \hat{T}$. ■

Appendix 2. Optimal allocation: second stage solution

The optimal allocation problem faced by the fund's manager for all $t \le \hat{T}$ is given by

$$J(0, W_0, Y_0) = \max_{(C,\alpha)\in\mathcal{A}_t} \mathbb{E}_0 \left[\int_0^{\hat{T}} f(C_t, J(t, W_t, Y_t))\, \mathrm{d}t + A \frac{W_{\hat{T}}^{1-\gamma}}{1-\gamma}\right] \tag{A9}$$

subject to (4) and (7), and where $A := \beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}}$. The terminal condition fixes the value of the indirect utility function at time $\hat{T}$ to that in (A5), i.e. $J(\hat{T}, W, 0) = J(W)$. A necessary condition for optimality for any $t \in [0, \hat{T}]$ is given by the Hamilton–Jacobi–Bellman (HJB) equation

$$0 = \max_{\{C,\alpha\}\in\mathcal{A}} \left\{ f(C, J(W)) + \frac{1}{\mathrm{d}t}\mathbb{E}_t[\mathrm{d}J(W)] \right\}.$$

An application of Itô's lemma implies

$$\begin{aligned}\mathrm{d}J(t, W, Y) &= \Big[J_t + [rW + \alpha(\mu - r)W + Y - C] J_W + \frac{1}{2}\sigma_S^2 (\alpha W)^2 J_{WW} + \kappa Y J_Y \\ &\quad + \frac{1}{2}\sigma_P^2 Y^2 J_{YY} + \sigma_S \sigma_P \rho_{PS} (\alpha W) Y J_{WY} \Big] \mathrm{d}t \\ &\quad + (\sigma_S \alpha_t W_t J_W + \sigma_P Y_t \rho_{PS} J_Y)\, \mathrm{d}Z_{S,t} + \sigma_P Y_t \sqrt{1 - \rho_{PS}^2} J_Y\, \mathrm{d}Z_{P,t},\end{aligned}$$

where $J_t := \partial J(t, W, Y)/\partial t$, $J_W := \partial J(t, W, Y)/\partial W$, $J_Y := \partial J(t, W, Y)/\partial Y$, $J_{WW} := \partial^2 J(t, W, Y)/\partial W^2$, $J_{YY} := \partial^2 J(t, W, Y)/\partial Y^2$, and $J_{WY} := \partial^2 J(t, W, Y)/\partial W \partial Y$. Using the martingale difference properties of stochastic integrals, we arrive at

$$0 = \max_{\{C,\alpha\}\in\mathcal{A}_t} \Big\{ f(C, J) + J_t + [rW + \alpha(\mu - r)W + Y - C] J_W + \frac{1}{2}\sigma_S^2 (\alpha W)^2 J_{WW} + \kappa Y J_Y + \frac{1}{2}\sigma_P^2 Y^2 J_{YY} + \sigma_S \sigma_P \rho_{PS} (\alpha W) Y J_{WY} \Big\}.$$

Under the assumptions of incomplete markets, the allocation problem for $t \le \hat{T}$, does not admit a closed-form solution. Section A.1 shows how to numerically approximate $J(W, Y, t)$ for all $t < \hat{T}$ such that the maximized HJB equation in (16)

$$\begin{aligned}0 &= \beta\theta J \left\{ \left[\frac{C}{[(1-\gamma)J]^{\frac{1}{1-\gamma}}}\right]^{1-\frac{1}{\psi}} - 1 \right\} \\ &\quad + J_t + [rW + \alpha(\mu - r)W + Y - C] J_W + \frac{1}{2}\sigma_S^2 (\alpha W)^2 J_{WW} \\ &\quad + \kappa Y J_Y + \frac{1}{2}\sigma_P^2 Y^2 J_{YY} + \sigma_S \sigma_P \rho_{PS} (\alpha W) Y J_{WY},\end{aligned} \tag{A10}$$

holds, where C and α are given by the first-order conditions in (17) and (18). The solution must satisfy the terminal condition $J(\hat{T}, W, 0)$. On the contrary, if financial markets are complete an analytical solution to the allocation problem is derived in Section A.2.

A.2.1. Incomplete markets solution

A.2.1.1. The transformed problem. The problem in (A10) consists of solving a nonlinear partial differential equation (PDE) in three state variables: time, financial wealth and income. To simplify the implementation of the numerical approximation, we exploit the homogeneity of the value function with respect to financial wealth and income to reduce the number of state variables from three to two.

As discussed in Wang *et al.* (2016), the value function $J(t, W, Y)$ is homogeneous of degree $(1-\gamma)$ in W and Y. Hence, for any given function $k(t)$ it holds that

$$J(t, k(t)W, k(t)Y) = k(t)^{1-\gamma} J(t, W, Y).$$

In line with Munk and Sørensen (2010), we set $k(t) = e^{-\delta t}/Y$, implying that

$$J(t, W, Y) = Y^{1-\gamma} e^{-\delta(\gamma-1)t} J\left(t, x, e^{-\delta t}\right) = Y^{1-\gamma} \frac{F(t, x)^{1-\gamma}}{1-\gamma}, \tag{A11}$$

where we have defined $x := e^{-\delta t} W/Y$ to be the scaled-adjusted financial wealth-to-income ratio, with $\partial x/\partial Y = -x/Y$, $\partial x/\partial t = -\delta x$ and $\partial x/\partial W = e^{-\delta t}/Y$. The parameter $\delta \geq 0$ prevents the financial wealth-to-income ratio from taking very large values as $t \to \hat{T}$ which could prevent the numerical algorithm to converge on a fixed grid for x. The value for δ is found by trial-and-error conditional on the calibration of the structural parameters.

The introduction of new state variable x allows us to simplify the original problem. In fact, substituting (A11) into (A10) yields

$$
\begin{aligned}
0 = {} & \frac{\beta\psi}{\psi-1}(FY)^{1-\gamma}\left\{\left(\frac{C}{YF}\right)^{1-\frac{1}{\psi}}-1\right\} + \underbrace{F^{-\gamma}Y^{1-\gamma}(F_t-\delta F_x x)}_{\equiv J_t} \\
& + \left[r\frac{W}{Y}+\alpha(\mu-r)\frac{W}{Y}+1-\frac{C}{Y}\right]Y\underbrace{e^{-\delta t}F^{-\gamma}Y^{-\gamma}F_x}_{\equiv J_W} \\
& + \frac{1}{2}\sigma_S^2\alpha^2\frac{W^2}{Y^2}Y^2\underbrace{e^{-2\delta t}F^{-\gamma-1}Y^{-\gamma-1}\left[FF_{xx}-\gamma F_x^2\right]}_{\equiv J_{WW}} \\
& + \kappa Y\underbrace{(FY)^{-\gamma}(F-xF_x)}_{\equiv J_Y} \\
& + \frac{1}{2}Y^2\sigma_Y^2\underbrace{(FY)^{-1-\gamma}\left(x^2FF_{xx}-\gamma(F-xF_x)^2\right)}_{\equiv J_{YY}} \\
& + \sigma_S\sigma_Y\rho_{YS}\alpha\frac{W}{Y}Y^2 \\
& \times\underbrace{e^{-\delta t}(FY)^{-\gamma-1}\left[-xFF_{xx}-(\gamma F_x)(F-xF_x)\right]}_{\equiv J_{WY}}.
\end{aligned}
$$

After some algebra, the transformed PDE reads

$$
\begin{aligned}
0 = {} & \frac{\beta\psi}{\psi-1}\hat{c}^{1-\frac{1}{\psi}}F^{\frac{1}{\psi}} + \left(\kappa-\frac{\beta\psi}{\psi-1}-\frac{\gamma\sigma_P^2}{2}\right)F+F_t \\
& + \Big[\left(r-\delta-\kappa+\sigma_P^2\gamma\right)x \\
& + (\mu-r-\gamma\sigma_S\sigma_P\rho_{PS})\alpha x + e^{-\delta t}-e^{-\delta t}\hat{c}\Big]F_x \\
& + \frac{1}{2}x^2\left(\sigma_S^2\alpha^2+\sigma_P^2-2\sigma_S\sigma_P\rho_{PS}\alpha\right)\left(F_{xx}-\gamma F^{-1}F_x^2\right),
\end{aligned}
\tag{A12}
$$

where the optimal consumption-to-income ratio, $\hat{c} := C/Y$, and the optimal investment share, α, are given by

$$
\hat{c} = e^{\psi\delta t}\beta^{\psi}F_x^{-\psi}F, \tag{A13}
$$

$$
\alpha = \frac{FF_x}{x\left[\gamma F_x^2-FF_{xx}\right]}\left(\frac{\mu-r}{\sigma_S^2}-\frac{\gamma\sigma_P\rho_{PS}}{\sigma_S}\right)+\frac{\sigma_P\rho_{PS}}{\sigma_S}. \tag{A14}
$$

The transformed problem is now that of finding the two state variable function $F(t,x)$ for all $t < \hat{T}$ that solves the PDE in (A12). The terminal condition for the transformed problem, $F(\hat{T},x)$, is related to the original terminal condition through (A11). In particular, note that at $t = \hat{T}$, the optimal value function can be written as

$$
\begin{aligned}
J\left(\hat{T},W,Y\right) &= \frac{\beta^{\theta}}{1-\gamma}\mathbb{G}_\infty^{-\frac{\theta}{\psi}}W^{1-\gamma} \\
&= \frac{\beta^{\theta}}{1-\gamma}\mathbb{G}_\infty^{-\frac{\theta}{\psi}}\left[Y\left(e^{-\delta\hat{T}}\frac{W}{Y}e^{\delta\hat{T}}\right)\right]^{1-\gamma} \\
&= \frac{\beta^{\theta}}{1-\gamma}\mathbb{G}_\infty^{-\frac{\theta}{\psi}}\left(e^{\delta\hat{T}}Yx\right)^{1-\gamma},
\end{aligned}
$$

where we have used the fact that $x(\hat{T}) = e^{-\delta\hat{T}}\frac{W}{Y}$. Then, by the homogeneity property of the value function, it follows that

$$
Y^{1-\gamma}\frac{F\left(\hat{T},x\right)^{1-\gamma}}{1-\gamma} = \frac{\beta^{\theta}}{1-\gamma}\mathbb{G}_\infty^{-\frac{\theta}{\psi}}\left(e^{\delta\hat{T}}Yx\right)^{1-\gamma},
$$

which implies the following value for the terminal condition

$$
F\left(\hat{T},x\right) = \left(\beta^{\theta}\mathbb{G}_\infty^{-\frac{\theta}{\psi}}\right)^{\frac{1}{1-\gamma}}e^{\delta\hat{T}}x, \tag{A15}
$$

where $\mathbb{G}_\infty$ is given in (13).

A.2.1.2. Finite difference approximation. We approximate the solution to the PDE in (A12) for $t \leq \hat{T}$ using the finite difference algorithm for nonlinear PDEs introduced in Gomez (2019). In particular, the finite difference method approximates $F(t,x)$ on a $(N+1)\times(J+1)$ rectangular grid of equally spaced points on the (t,x)−space with values $\{(t_n,x_j)\,|\,n=0,1,\ldots,N, j=0,1,\ldots,J\}$, where $x_j = x_0 + j\Delta x$ and $t_n = n\Delta t$ for some fixed spacing parameters Δx and Δt.

Let $F_{j,n} := F(t_n,x_j)$ denote the approximated value function at grid point (t_n,x_j). For $\hat{T} = t_N = N\Delta t$, the approximated value function is set equal to

$$
F_{j,\hat{T}} = \left(\beta^{\theta}\mathbb{G}_\infty^{-\frac{\theta}{\psi}}\right)^{\frac{1}{1-\gamma}}e^{\delta\hat{T}}x_j, \tag{A16}
$$

for all $j = 0,1\ldots,J$. Given (A16), the optimal investment share and consumption-to-income ratio at time $\hat{T}$, $\hat{c}_{j,\hat{T}}$ and $\alpha_{j,\hat{T}}$, are computed from (14) and (15), respectively.

Now, for each $j = 0,1\ldots,J$ and $t = 0,1,\ldots,N-1$ in the interior of the grid, we compute the time derivative of the value function using the forward difference approximation

$$
F_t^+ \approx D_t^+F_{j,n} = \frac{F_{j,n+1}-F_{j,n}}{\Delta t},
$$

whereas the first-order derivative with respect to the scale-adjusted wealth-to-income ratio is computed with either a forward or a backward difference operator

$$
F_x^+ \approx D_x^+F_{j,n} = \frac{F_{j+1,n}-F_{j,n}}{\Delta x},
$$

$$
F_x^- \approx D_x^-F_{j,n} = \frac{F_{j,n}-F_{j-1,n}}{\Delta x}.
$$

Finally, the second-order derivatives are approximated using the central difference operator

$$
F_{xx} \approx D_x^2F_{j,n} = \frac{F_{j+1,n}-2F_{j,n}+F_{j-1,n}}{(\Delta x)^2}.
$$

Following Candler (1999) and Achdou *et al.* (2022), the choice of difference operator for F_x is based on an upwind differentiation scheme according to which the correct approximation, $D_xF_{j,n}$, is determined by the direction of state variable. In what follows, the direction will be determined by the sign of

$$
\begin{aligned}
z_{j,n} := {} & \left(r-\delta-\kappa+\sigma_P^2\gamma\right)x_{j,n} + (\mu-r-\gamma\sigma_S\sigma_P\rho_{PS})\alpha_{j,n}x_{j,n} \\
& + e^{-\delta n\Delta t} - e^{-\delta n\Delta t}\hat{c}_{j,n},
\end{aligned}
$$

where

$$
\hat{c}_{j,n} = e^{\psi\delta n\Delta t}\beta^{\psi}\left(D_xF_{j,n}\right)^{-\psi}F_{j,n},
$$

$$
\begin{aligned}
\alpha_{j,n} = {} & \frac{F_{j,n}\left(D_xF_{j,n}\right)}{x\left[\gamma\left(D_xF_{j,n}\right)^2-F_{j,n}\left(D_x^2F_{j,n}\right)\right]} \\
& \times\left(\frac{\mu-r}{\sigma_S^2}-\frac{\gamma\sigma_P\rho_{PS}}{\sigma_S}\right)+\frac{\sigma_P\rho_{PS}}{\sigma_S},
\end{aligned}
$$

are the optimal consumption-to-income ratio and investment rate at grid point (j,n). Thus, if the 'drift' variable $z_{j,n}$ is positive we use the forward operator and if it is negative, we use the backward operator. This gives rise to the following upwind operator

$$D_xF_{j,n} = \left(D_x^+F_{j,n}\right)\mathbb{1}_{\{z^+\geq 0\}} + \left(D_x^-F_{j,n}\right)\mathbb{1}_{\{z^-<0\}},$$

where $\mathbb{1}$ denotes the indicator function, and z^+ and z^- the 'drift' variables computed with the forward and backward operators, respectively. Then, the finite difference approximation to the HJB equation at grid point (n,j) is given by

$$\begin{aligned}
&-\frac{F_{j,n+1}-F_{j,n}}{\Delta t}\\
&= \frac{\beta\psi}{\psi-1}\left(\hat{c}_{j,n}\right)^{1-\frac{1}{\psi}}\left(F_{j,n}\right)^{\frac{1}{\psi}} + \left(\kappa - \frac{\beta\psi}{\psi-1} - \frac{\gamma\sigma_P^2}{2}\right)F_{j,n}\\
&\quad + z_{j,n}^+\left(D_x^+F_{j,n}\right)\mathbb{1}_{\{z^+\geq 0\}} + z_{j,n}^-\left(D_x^-F_{j,n}\right)\mathbb{1}_{\{z^-<0\}}\\
&\quad + \frac{1}{2}x_{j,n}^2\left(\sigma_S^2\alpha_{j,n}^2 + \sigma_P^2 - 2\sigma_S\sigma_P\rho_{PS}\alpha_{j,n}\right)\Big[\left(D_x^2F_{j,n}\right)\\
&\quad - \gamma\left(F_{j,n}\right)^{-1}\left(\left(D_x^+F_{j,n}\right)\mathbb{1}_{\{z^+\geq 0\}} + \left(D_x^-F_{j,n}\right)\mathbb{1}_{\{z^-<0\}}\right)^2\Big].
\end{aligned} \quad \text{(A17)}$$

Given a value $F_{j,n+1}$ for all j, the approximation in (A17) can be compactly written as a system of $(J+1)$ nonlinear equations: one for each $n = 0, 1, \ldots, N-1$. An approximation to the value function at time t_n is therefore given by the vector $\vec{F}_n = [F_{0,n}, F_{1,n}, \ldots, F_{J,n}]^\top$ that solves $\vec{G}(\vec{F}_n) = \vec{0}$, where $\vec{F}_n$ denotes the unknown value function at all the grid points in the $x-$lattice at time step n. To compute the approximation $\vec{F}_n$ for all $n = 0, 1, \ldots, N-1$, we iterate backwards on time starting from the terminal condition $\vec{F}_N$ in (A16). This recursion can be written as

$$\vec{0} = \mathbf{G}\left(\vec{F}_n\right) + \frac{1}{\Delta t}\left(\vec{F}_{n+1} - \vec{F}_n\right). \quad \text{(A18)}$$

A.2.2. *Complete markets solution*

Proof of Lemma 3.1 Under the assumption of complete markets it follows that $|\rho_{PS}| = 1$. Hence, the dynamics of the fund's manager income under the physical probability measure $\mathbb{P}$ is given by the Geometric Brownian motion

$$\frac{\mathrm{d}Y_t}{Y_t} = \kappa\,\mathrm{d}t + \xi\,\mathrm{d}Z_{S,t}, \quad \text{(A19)}$$

where $\xi := \sigma_P \times \rho_{PS}$. Let the market price of risk be given by $\lambda := (\mu - r)/\sigma_S$. Then, by Girsanov's theorem, the fund's income has the following equivalent Geometric Brownian motion representation under the risk-neutral probability measure $\mathbb{Q}$

$$\frac{\mathrm{d}Y_t}{Y_t} = (\kappa - \xi\lambda)\,\mathrm{d}t + \xi\,\mathrm{d}Z_{S,t}^{\mathbb{Q}}, \quad \text{(A20)}$$

with solution

$$Y_u = Y_t e^{\left((\kappa-\xi\lambda-\frac{1}{2}\xi^2)(u-t)+\xi\left(Z_{S,u}^{\mathbb{Q}}-Z_{S,t}^{\mathbb{Q}}\right)\right)}, \quad \text{for } u \geq t. \quad \text{(A21)}$$

Let $\mathcal{O}_t := \mathcal{O}(Y_t, t; \hat{T})$ denote the expected present discounted value at time t of all future oil income, $\{Y_u\}_{u=t}^{\hat{T}}$, i.e.

$$\mathcal{O}_t = \mathbb{E}_t^{\mathbb{Q}}\left[\int_t^{\hat{T}} e^{-r(u-t)}Y_u\,\mathrm{d}u\right]. \quad \text{(A22)}$$

Multiplying both sides of (A21) by $e^{-r(u-t)}$ and integrating from t to $\hat{T}$ we arrive at

$$\int_t^{\hat{T}} e^{(-r(u-t))}Y_u\,\mathrm{d}u = Y_t\int_t^{\hat{T}} e^{\left((-r+\kappa-\xi\lambda-\frac{1}{2}\xi^2)(u-t)+\xi\left(Z_{S,u}^{\mathbb{Q}}-Z_{S,t}^{\mathbb{Q}}\right)\right)}\,\mathrm{d}u. \quad \text{(A23)}$$

Since $Z_{S,u}^{\mathbb{Q}} - Z_{S,t}^{\mathbb{Q}}$ is normally distributed, the term inside the integral on the right-hand side of (A23) is log-normally distributed. Hence, the expected value conditional on the information at time t is

$$\mathbb{E}_t^{\mathbb{Q}}\left[\int_t^{\hat{T}} e^{-r(u-t)}Y_u\,\mathrm{d}u\right] = Y_t\int_t^{\hat{T}} e^{-(r-\kappa+\xi\lambda)(u-t)}\,\mathrm{d}u. \quad \text{(A24)}$$

It follows that

$$\mathcal{O}_t = Y_t\mathbb{1}_{\{t<\hat{T}\}}\underbrace{\begin{cases}\dfrac{1}{r-\kappa+\xi\lambda}\left(1-e^{-(r-\kappa+\xi\lambda)\left(\hat{T}-t\right)}\right) & \text{if } r-\kappa+\xi\lambda \neq 0\\ (T-t) & \text{if } r-\kappa+\xi\lambda = 0.\end{cases}}_{\equiv\mathcal{M}(t)\text{:Income multiplier}} \quad \text{(A25)}$$

Note that under the risk-neutral probability measure $\mathbb{Q}$, the Feynman–Kac Theorem suggests that $\mathcal{O}(Y_t, t; \hat{T})$ satisfies the following partial differential equation (PDE)

$$\frac{\partial\mathcal{O}_t}{\partial t} + (\kappa - \xi\lambda)\,Y_t\frac{\partial\mathcal{O}_t}{\partial Y_t} + \frac{1}{2}\xi^2Y_t^2\frac{\partial^2\mathcal{O}_t}{\partial Y_t^2} - r\mathcal{O}_t + Y_t = 0, \quad \text{(A26)}$$

with terminal condition $\mathcal{O}(Y_{\hat{T}}, \hat{T}; \hat{T}) = 0$. ■

Proof of Proposition 3.2 Conjecture that the value function takes the form

$$J\left(t, W, Y\right) = \frac{\beta^\theta}{1-\gamma}\mathbb{G}\left(t\right)^{\frac{\theta}{\psi}}\left(W + \mathcal{O}\right)^{1-\gamma} \quad \text{(A27)}$$

where $W + \mathcal{O}$ is the fund's total wealth at a given point in time, and $\mathbb{G}(t)$ is an unknown deterministic function to be determined. Our conjecture uses the ideas in Bodie *et al.* (1992) according to which it is possible to think of the fund's manager as having an initial wealth equal to $W + \mathcal{O}$ and no oil income, instead of having an initial level of financial wealth and an flow of income. Our conjecture implies that

$$\begin{aligned}
J_t &= \frac{\beta^\theta}{(\psi-1)}\mathbb{G}\left(t\right)^{\frac{\theta}{\psi}-1}\frac{\partial\mathbb{G}\left(t\right)}{\partial t}\left(W+\mathcal{O}\right)^{1-\gamma}\\
&\quad + \beta^\theta\mathbb{G}\left(t\right)^{\frac{\theta}{\psi}}\left(W+\mathcal{O}\right)^{-\gamma}\frac{\partial\mathcal{O}}{\partial t},\\
J_W &= \beta^\theta\mathbb{G}\left(t\right)^{\frac{\theta}{\psi}}\left(W+\mathcal{O}\right)^{-\gamma},\\
J_{WW} &= -\gamma\beta^\theta\mathbb{G}\left(t\right)^{\frac{\theta}{\psi}}\left(W+\mathcal{O}\right)^{-\gamma-1},\\
J_Y &= \beta^\theta\mathbb{G}\left(t\right)^{\frac{\theta}{\psi}}\left(W+\mathcal{O}\right)^{-\gamma}\mathcal{O}_Y,\\
J_{YY} &= -\gamma\beta^\theta\mathbb{G}\left(t\right)^{\frac{\theta}{\psi}}\left(W+\mathcal{O}\right)^{-\gamma-1}\left(\mathcal{O}_Y\right)^2,\\
J_{WY} &= -\gamma\beta^\theta\mathbb{G}\left(t\right)^{\frac{\theta}{\psi}}\left(W+\mathcal{O}\right)^{-\gamma-1}\mathcal{O}_Y,
\end{aligned}$$

where we have used the fact that, according to Lemma 3.1, $\mathcal{O}_{YY} = 0$. Subscripts on the value function J denote partial derivatives with respect to the respective state variables. Substituting in (17) and (18) yields

$$C_t = \mathbb{G}\left(t\right)^{-1}\left(W_t + \mathcal{O}_t\right), \quad \text{(A28)}$$

$$\alpha_t = \frac{1}{\gamma}\left(\frac{\mu-r}{\sigma_S^2}\right)\left(1+\frac{\mathcal{O}_t}{W_t}\right) - \frac{\mathcal{O}_t}{W_t}\frac{\sigma_P\rho_{PS}}{\sigma_S}, \quad \text{(A29)}$$

where we have used (A25) to conclude that $\mathcal{O}_t = Y_t\mathcal{O}_Y$.

Substituting into the maximized HJB Equation (A10) and using (A26) together with the fact $rW = r(W+\mathcal{O}) - r\mathcal{O}$, we arrive to the linear ODE

$$\frac{\partial \mathbb{G}(t)}{\partial t} - \left[\beta\psi + (1-\psi)r + (1-\psi)\frac{1}{2\gamma}\left(\frac{\mu - r}{\sigma_S}\right)^2\right] \times \mathbb{G}(t) + 1 = 0. \quad \text{(A30)}$$

Using the terminal condition for the HJB equation

$$J\left(\hat{T}, W, Y\right) = \frac{1}{1-\gamma}\beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W_{\hat{T}}^{1-\gamma},$$

together with our conjecture, we obtain the terminal condition that the ODE in (A30) has to satisfy. In particular

$$\frac{\beta^\theta}{1-\gamma}\mathbb{G}\left(\hat{T}\right)^{\frac{\theta}{\psi}} W_{\hat{T}}^{1-\gamma} = \frac{1}{1-\gamma}\beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}} W_{\hat{T}}^{1-\gamma}$$
$$\Downarrow$$
$$\mathbb{G}\left(\hat{T}\right) = \mathbb{G}_\infty^{-1},$$

which implies that

$$\mathbb{G}(t) = \mathbb{G}_\infty^{-1} \quad \forall\, t < \hat{T}, \quad \text{(A31)}$$

and our conjecture has been verified. ■

Appendix 3. Optimal trajectories under complete markets

Figure A.1 illustrates the optimal path for selected variables for all $t \in [0, \hat{T}]$, together with intervals around the median that represent the 15th and 85th percentiles of their distribution generated from 10,000 simulations of the model. We assume that markets are complete and therefore use Lemma 3.1 to compute the value of the underground oil wealth at each point in time. The optimal consumption-to-financial wealth ratio, and the optimal demand for the risky asset follow from Proposition 3.2 when replacing ρ_{PS} by the estimated correlation coefficient reported in table 1.

Appendix 4. A model with random $\hat{T}$

A.4.1. The random terminal date

Let $\hat{T} \in [0, \infty)$ denote the random time at which the fund's manager stops receiving commodity income, with probability density function $\pi(t)$. Hence, the probability of receiving income between 0 and t (the survival probability) is

$$p_t := \mathbb{P}(\hat{T} > t \mid \hat{T} > 0) = \int_t^\infty \pi(s)\,\mathrm{d}s = 1 - \int_0^t \pi(s)\,\mathrm{d}s, \quad \text{(A32)}$$

from which it follows that

$$\frac{\mathrm{d}p_t}{p_t} = -\frac{\pi_t}{1 - \int_0^t \pi(s)\,\mathrm{d}s}\,\mathrm{d}t, \quad p_0 = 1. \quad \text{(A33)}$$

The hazard rate, i.e. the probability that the fund will stop receiving income in the next instant of time, given that it has received a positive flow of income up to time t, is

$$\lambda_t = \frac{\pi_t}{1 - \int_0^t \pi(s)\,\mathrm{d}s}. \quad \text{(A34)}$$

Hence, λ_t can be interpreted as the conditional arrival rate of an event that will stop the flow of income to the fund at time $t > \hat{T}$. Notice that the unique solution to (A33) is

$$p_t := \mathbb{E}_0\left[\mathbb{1}_{t<\hat{T}}\right] = \exp\left(-\int_0^t \lambda_s\,\mathrm{d}s\right), \quad \text{(A35)}$$

where $\mathbb{1}_x$ is the indicator function. For $\lambda_t = \lambda\ \forall t$, the random time $\hat{T}$ represents the date of the first jump of a standard Poisson process, and (A35) becomes

$$p_t := \mathbb{E}_0\left[\mathbb{1}_{t<\hat{T}}\right] = \exp(-\lambda t). \quad \text{(A36)}$$

Let $d_t := 1 - p_t = \int_0^t \pi(s)\,\mathrm{d}s$ denote the complement of the survival function. It follows that

$$\frac{\mathrm{d}d_t}{\mathrm{d}t} = \pi_t = \lambda_t \exp\left(-\int_0^t \lambda_s\,\mathrm{d}s\right). \quad \text{(A37)}$$

A.4.2. The manager's problem

The problem faced by the fund's manager is given by

$$J_0 = \max_{\{C_t,\alpha_t\}_{t=0}^{\hat{T}}} \mathbb{E}_0^{\hat{T}}\left[\int_0^{\hat{T}} f(C_t, J_t)\,\mathrm{d}t + A\frac{W_{\hat{T}}^{1-\gamma}}{1-\gamma}\right], \quad \text{(A38)}$$

subject to (7) and (4), where $\mathbb{E}_0^{\hat{T}}[\cdot]$ is the expected value conditional on the information available at time $t = 0$, $f(C, J)$ is given in (6), and the constant $A = \beta^\theta \mathbb{G}_\infty^{-\frac{\theta}{\psi}}$, with $\mathbb{G}_\infty$ given in (13). The value of A is consistent with the assumption that the fund's manager optimizes even after the oil income is over, $Y_t = 0\ \forall t > \hat{T}$.

In the following, let us assume that the Brownian motions $Z_{S,t}$ and $Z_{P,t}$, and the random terminal date, $\hat{T}$, are mutually independent. Then, (A38) can be written as

$$\begin{aligned}
J_0 &= \max_{\{C_t,\alpha_t\}_{t=0}^{\hat{T}}} \mathbb{E}_0^{\hat{T}}\left[\int_0^{\hat{T}} f(C_t, J_t)\,\mathrm{d}t + A\frac{W_{\hat{T}}^{1-\gamma}}{1-\gamma}\right] \\
&= \max_{\{C_t,\alpha_t\}_{t=0}^{\hat{T}}} \mathbb{E}_0^{\hat{T}}\left[\int_0^\infty \mathbb{1}_{t<\hat{T}} f(C_t, J_t)\,\mathrm{d}t + \mathbb{1}_{t\geq\hat{T}} A\frac{W_{\hat{T}}^{1-\gamma}}{1-\gamma}\right] \\
&= \max_{\{C_t,\alpha_t\}_{t=0}^{\hat{T}}} \int_0^\infty \mathbb{E}_0^{\hat{T}}\left[\mathbb{1}_{t<\hat{T}}\right] \mathbb{E}_0^{\hat{T}}\left[f(C_t, J_t)\right]\,\mathrm{d}t \\
&\quad + \mathbb{E}_0^{\hat{T}}\left[\mathbb{1}_{t\geq\hat{T}} A\frac{W_{\hat{T}}^{1-\gamma}}{1-\gamma}\right] \\
&= \max_{\{C_t,\alpha_t\}_{t=0}^{\hat{T}}} \mathbb{E}_0\left[\int_0^\infty p_t f(C_t, J_t)\,\mathrm{d}t + \int_0^\infty A\frac{W_t^{1-\gamma}}{1-\gamma}\,\mathrm{d}d_t\right] \\
&= \max_{\{C_t,\alpha_t\}_{t=0}^{\hat{T}}} \mathbb{E}_0\left[\int_0^\infty e^{-\int_0^t \lambda_s\,\mathrm{d}s} f(C_t, J_t)\,\mathrm{d}t\right. \\
&\quad \left. + \int_0^\infty \lambda_t e^{-\int_0^t \lambda_s\,\mathrm{d}s} A\frac{W_t^{1-\gamma}}{1-\gamma}\,\mathrm{d}t\right] \quad \text{(A39)}
\end{aligned}$$

by resolving the uncertainty with respect to the random time $\hat{T}$. For $\lambda_t = \lambda$ for all $t > 0$, the above simplifies to (29) in the main text.

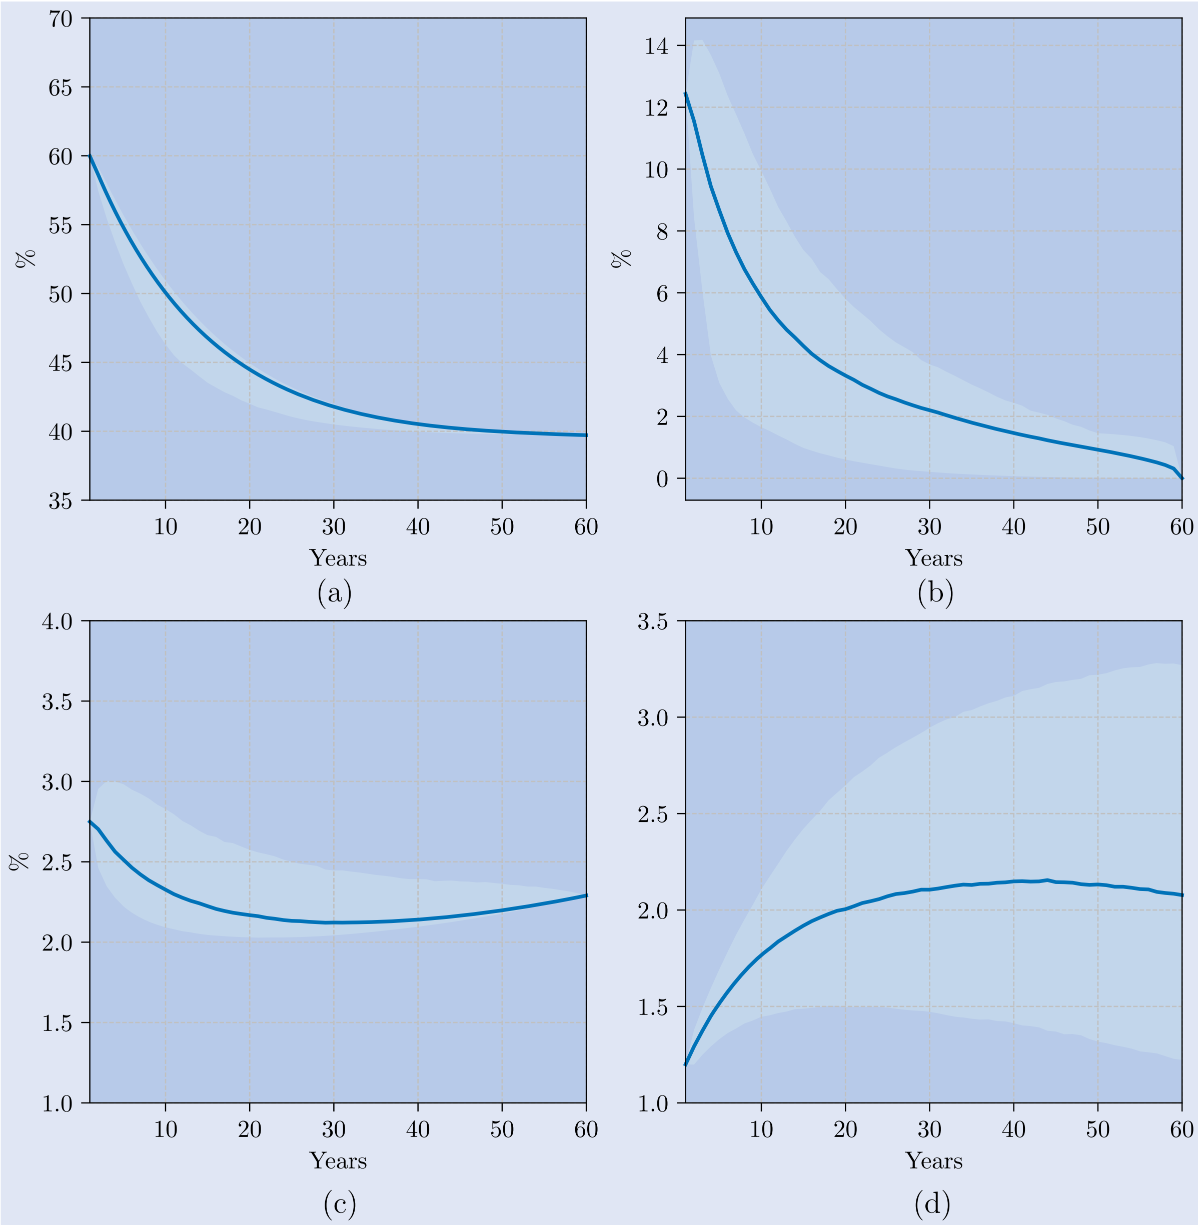

Figure A1. Optimal strategies under complete markets. Panels (a)–(d) plot, respectively, the optimal share of financial wealth invested in equity, the hedging demand as a fraction of the financial wealth, the optimal consumption-to-financial wealth ratio, and the evolution of financial wealth-to-mainland GDP ratio. The optimal consumption-to-financial wealth ratio and the optimal demand for equity are given by (21) and (22) in Proposition 3.2, respectively. The solid lines represent the median value over M =10,000 simulated paths using the parameters in table 1, each of them of $T = 60$ sample points. The shaded areas represent the 15 and 85 percentiles from the sampling distribution of the simulated series.