NORGES BANK
INVESTMENT MANAGEMENT

04|2016

ASSET ALLOCATION WITH GOVERNMENT REVENUES AND SPENDING COMMITMENTS

DISCUSSION NOTE

In this note, we explore the implications of expanding the discussion about the equity share in the Norwegian Government Pension Fund Global to include revenues and spending commitments, using a stylised simulation framework.

Date 01/12/2016
ISSN 1893-966X

dn@nbim.no
www.nbim.no

SUMMARY

- It is common to consider an equity and fixed-income allocation problem in terms of expected returns and risk. These are important inputs, but there are additional considerations that should be taken into account, in particular other revenues and commitments, which can change the problem significantly.

- Revenues and spending commitments are particularly important considerations in the equity share decision for the Norwegian Government Pension Fund Global, since they are significant in size and variability, and co-vary with returns on the assets held in the fund. Using a stylised, reduced-form simulation framework, we explore the implications of including revenues and spending commitments. Our objective is to illustrate some of the trade-offs that ought to be considered by the government, rather than to advise on a specific equity share.

- We expand the analysis from considering the financial portfolio in isolation, to including inflows from oil and a stylised spending rule, to finally looking at the broader government problem with tax revenues and stable government spending.

- When we include income and spending considerations, the risk being borne for any given equity share increases substantially, driven by correlations between the various components in the problem. This implies that analyses that do not account for the behaviour of revenues and spending commitments are likely to understate the risk in the risk-return trade-off actually being faced by the government.

- Of particular importance is the level and flexibility of government expenditure: lower spending levels are less likely to be underfunded, and the longevity of the fund is likely to be significantly extended as a result.

- While an equity share may be chosen on the basis of a given set of assumptions regarding revenues, spending, market variables, and their interaction, a significant change in any of these inputs would imply that the equity share decision should be revisited.

1. Introduction

This note uses a simulation approach to highlight some of the broader considerations faced by the Norwegian government when choosing the equity share in the Government Pension Fund Global (GPFG). Early academic work emphasised the role of expected returns, risk and co-movement of financial assets as key inputs into portfolio choice problems.[1] This work was later extended to explore the effects of adding income from outside the financial portfolio. These more recent studies show that, when taking into account outside income or liability considerations, the appropriate asset allocation can change substantially relative to considering the financial assets in isolation.[2] This suggests that, for an individual, institution or sovereign nation, there are considerations beyond the financial portfolio that should be taken into account when deciding an appropriate equity share. From the perspective of the Norwegian government, the return on the investments in the GPFG can be considered as an income source alongside income from tax revenues from various forms of economic activity. The government also has commitments in the form of fiscal spending each year, which is financed by combining its various income sources.

In this note, we explore the implications of expanding the equity share decision problem to include the financial portfolio, tax incomes and government expenditure together. We develop a stylised, reduced-form simulation framework where we model investment returns alongside simple characterisations of the Norwegian government's oil and non-oil tax revenues and fiscal spending commitments. We design our model to capture some core institutional arrangements where the government transfers net oil revenues into the GPFG and finances a yearly non-oil deficit from the fund, partly based on a spending rule.

In the first part of our analysis, we incrementally add components into a baseline model. We expand the equity share decision problem from looking at the financial portfolio in isolation, to including inflows from oil and a spending rule, to looking at the government's full problem including tax revenues and total government spending. As we do this, in order to be able to compare the implicit risk the government is running, we need a metric summarising this risk. We construct our framework in such a way that the risks facing the government across the problems are all captured in the *size distribution* of the fund at a given horizon. When looking at the financial portfolio in isolation, the size distribution reflects the financial market volatility that comes from the choice of an equity share. When we add oil revenues and a stylised spending rule, the size distribution reflects not only the volatility that comes from the choice of an equity share in the financial portfolio, but also the substantial volatility that comes from oil revenues. Finally, in the full model, fiscal spending is financed by non-oil tax revenues and through drawings on the fund. Given that it is unlikely that government

1 This approach was pioneered in the seminal work by Markowitz (1952) and extended in Merton (1969) and Samuelson (1969).

2 This idea has also been explored in a large body of academic work, where it is usually formalised in models where an individual investor receives a proportion of her total income from supplying labour (see, among others, Bodie, Merton and Samuelson (1992), Heaton and Lucas (1997), Viceira (2001), Cocco, Gomez and Maenhout (2005), Benzoni, Collin-Dufresne and Goldstein (2007) and Cochrane (2014)).

expenditure can vary too much from year to year, in the full model we assume that government expenditure is fixed (growing at a constant rate), while we allow the revenue sources that finance this expenditure to vary over time. As a result, it is possible that these income sources undershoot the required amount of spending in any given year. In this scenario, in order to maintain a fixed level of spending, the government needs to draw additional financing from the fund. In the model, the risks associated with the financing of expenditure, and risks associated with investment returns and oil revenues, are hence all captured by the amount of capital in the fund at a given time. As a result, the size distribution of the fund is a sufficient statistic for capturing the implications of varying the equity share also in the broader government problem.

To characterise the fund's size distribution across the different problems, we simulate a large number of possible paths for equity and fixed-income returns, foreign exchange, oil prices and government income and spending, given a set of capital market and spending behaviour assumptions. We show that, when augmenting the equity share decision with income and spending considerations, there is a substantial increase in the risk being borne by the government for any given equity share, driven by correlations between the various components. This implies that analyses that do not account for the behaviour of spending and revenue are likely to understate the risk-return trade-off actually being faced. This is especially true when incorporating oil revenues into the model, since oil prices have varied substantially and have tended to co-vary quite strongly with financial market returns, at least over the past ten years. Our analysis also highlights how the nature and extent of these risks can change with different assumptions and inputs. Of particular importance is the level and flexibility of government expenditure: lower spending levels are naturally less likely to be underfunded, and the expected longevity of the fund is significantly extended as a result. Increased flexibility in spending in the short run ensures that there is less volatility in consumption in the longer run. In general, while an equity share may be chosen on the basis of given spending behaviour and revenues, and returns and correlations across revenues, spending and markets, a significant structural change in any of these input variables would imply that the equity share decision should be revisited.

In the second part of the analysis, we explore a number of issues within the full model. First, we analyse the implications of the large growth in the GPFG as a proportion of total government expenditure over the past decade. We show that, under a strict version of the current spending rule, this would imply a greater risk of variability in spending, all else equal. We then show the importance of the level that spending is set at, and that reducing the spending proportion of the fund reduces the risk of a fund shortfall substantially and increases expected fund longevity. Following this, we look at the importance of rebalancing in our model, and find that the choice of whether to rebalance is a much smaller consideration relative to the equity share decision itself. We then show that the choice of currency in which to measure the fund is important, and that using krone- instead of dollar-denominated values can change the analysis substantially. Finally, we examine the main drivers of fund growth over the past ten years, and explore how the fund might perform over the next ten years under adverse scenarios for oil prices and equity and fixed-income returns.

The note proceeds as follows. The next section provides background to the broader equity share discussion, setting out how the fund has evolved over recent history. Section 3 provides details of the model set-up, calibration and simulation methodology. Section 4 incrementally builds spending and revenue layers to illustrate how these considerations impact the equity share decision. In the final section, we explore a range of additional issues using the full model.

2. The GPFG and the Norwegian economy

Broadly speaking, in asset allocation models, the importance of revenues and spending streams is proportional to the magnitude and variability of these streams, as well as the co-movement between the streams and returns on investments in the financial portfolio. For Norway, basic figures suggest that revenues and spending are likely to be important considerations in choosing the equity share. They are significant in size and variability, and co-vary with returns on the assets held in the GPFG.

The fund structure was established by law in 1990 to meet key economic policy objectives. The separation of the spending decision from the income stream from oil production was key. This enabled the government to run high oil extraction rates and at the same time maintain economic stability and avoid "Dutch disease", i.e. the deindustrialisation of internationally competitive sectors in Norway. It was also recognised that oil resources represent wealth, and that the spending of oil revenues reduces the national wealth position. Saving part of the oil revenues through international financial investments would reallocate wealth and secure some level of intergenerational equity.

The fund is set up such that two sources of revenues – government oil revenues and the fund's return – are both transferred to the fund directly. An amount equalling the "non-oil" budget deficit, i.e. government spending in excess of "non-oil" income, is transferred from the fund. Throughout the period from 1996 until 2015, the oil revenue into the fund exceeded the transfer needed to cover the deficit, and due to this positive net transfer and the reinvestment of the fund's return, the accumulation of wealth in the fund has been uninterrupted. High oil prices, high production levels and high domestic economic activity led to the fund growing to around 2.5 times Norwegian GDP, as shown in Figure 1.

Figure 1: Fund value (billion NOK) vs Norway GDP

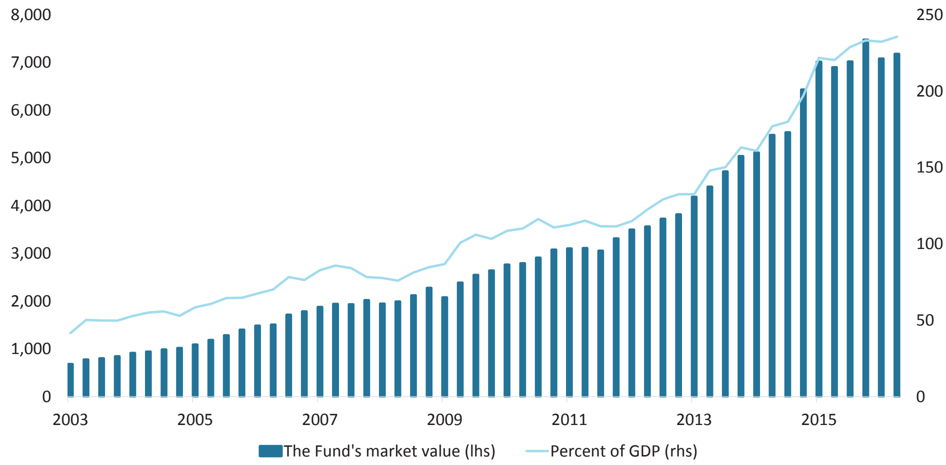

In 2001, a fiscal "spending rule" was introduced. According to the rule, spending should not exceed the expected real return of the fund, estimated to be 4 percent at the time. The spending rule is not a legal requirement, but rather a political economic yardstick which secures the original fund objectives and strengthens the intergenerational perspective. The growth in fund value has far exceeded the growth of the domestic economy and has thus rendered the spending rule a non-binding constraint. In recent years, transfers to cover the (structurally adjusted) deficit have averaged around 3 percent of the value of the fund, covering a growing budget deficit. The fund has become an increasingly important contributor to central government expenditure, as shown in Figure 2, making up around 16 percent of expenditure in 2016.

Figure 2: Transfers to finance the (structurally adjusted) non-oil government budget deficit, share of budget and fund

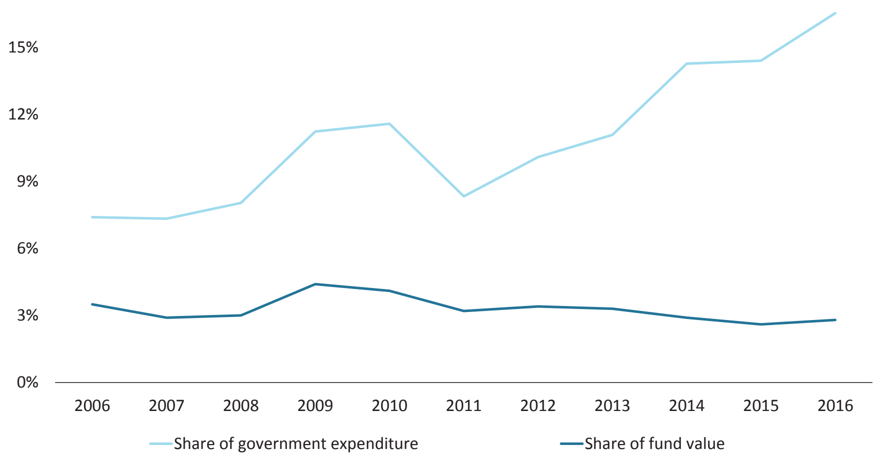

One of the most prominent activities of the Norwegian economy is the production of oil and gas. The remaining offshore petroleum wealth is estimated to be around 3,500 billion kroner, and a large domestic industry has developed around the production of oil and gas, accounting for around 22 percent of annual gross domestic product in 2015. Naturally, the revenue from these activities is strongly influenced by the level of oil and gas prices, which have tended to vary substantially. Over the past two decades, the net flows to the fund (oil revenues less spending) have been substantial, as shown in Figure 3, totalling around 3,700 billion kroner.[3] Following the recent fall in oil prices, and increased spending, this net inflow to the fund has become a net outflow.

Figure 3: Net inflow to the fund: petroleum revenues less non-oil budget deficit

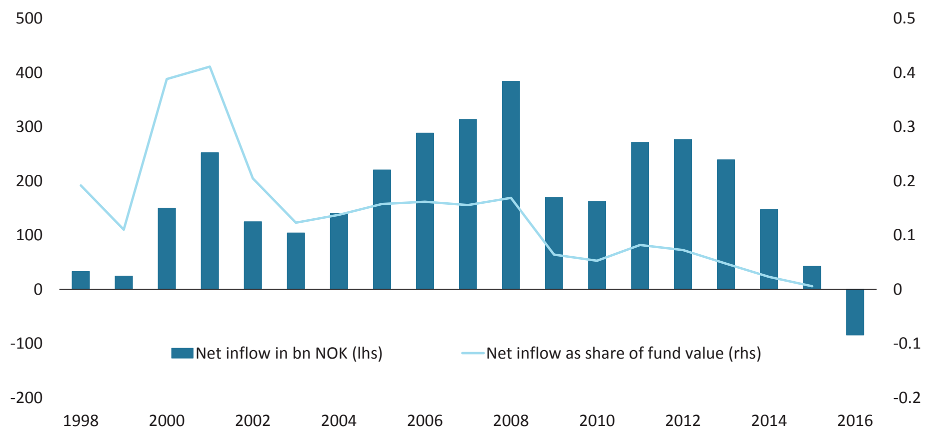

The sectors of the Norwegian economy outside the oil sector also form a large proportion of total GDP, and the income generated from these sectors can also vary substantially over time. The "non-oil" tax income from these sectors naturally depends on the overall health of the economy, and also co-varies with oil and gas sector income. The diverse nature of the fund's investments means that it also has exposure to a large range of non-oil related securities, the returns of which are likely to co-move with non-oil tax revenues.

3. Framework and simulation set-up

This section describes the framework we use to explore the equity share decision with government revenue and spending. We first outline the reduced-form processes that describe the dynamics of financial market returns, foreign exchange rates and oil prices. We then outline simple processes for government spending and revenues. Within our framework, all variables are in real terms, and there are two currencies, domestic currency denominated in Norwegian kroner (NOK) and foreign currency denominated in US dollars (USD). We focus mostly on the portfolio measured in foreign currency, as one of the main aims of the fund is to safeguard and build its international purchasing power.

i. Financial market processes

We model real log returns on dollar-denominated equity (r_{t+1}^{EQ}) and fixed-income (r_{t+1}^{FI}) markets. Returns are a function of a risk-free rate, a mean-reverting state variable and a noise term:

3 Measured in 2013 values.

$$r_{t+1}^{EQ} = r_t^f + X_t + \varepsilon_{t+1}^{EQ}$$

$$X_{t+1} = \varphi_0 + \varphi_1 X_t + \varepsilon_{t+1}^{x}$$

$$r_{t+1}^{FI} = r_t^f + Z_t + \varepsilon_{t+1}^{FI}$$

$$Z_{t+1} = \gamma_0 + \gamma_1 Z_t + \varepsilon_{t+1}^{z}$$

The state variables, X_t and Z_t, allow for predictability in equity and fixed-income markets, respectively. We also model oil prices and the USDNOK exchange rate. Both the log oil price ($p_t^{OIL} = \log(P_t^{OIL})$) and log USDNOK ($p_t^{FX} = \log(P_t^{FX})$) follow random walks:

$$p_{t+1}^{FX} = \alpha_{FX} + p_t^{FX} + \varepsilon_{t+1}^{FX}$$

$$p_{t+1}^{OIL} = \alpha_{OIL} + p_t^{OIL} + \varepsilon_{t+1}^{OIL}$$

The oil price and exchange rate drift terms (α_{FX} and α_{OIL}) are included to account for the Jensen's inequality term arising from modelling log prices and are set at 0.5 times the respective shock variance.

ii. Government processes

The government spends a krone-denominated amount each year, G_t, which grows at a rate of g_G, each year:

$$G_t = G_0(1 + g_G)^t$$

The government also receives revenues in the form of taxes on non-oil related activity, R_t^{NOIL}, denominated in kroner. The growth rate of this income, g_t^{NOIL}, follows a first-order autoregressive process with the same unconditional mean as the growth in government expenditure:

$$R_{t+1}^{NOIL} = R_t^{NOIL}(1 + g_{t+1}^{NOIL})$$

$$g_{t+1}^{NOIL} = \delta_0 + \delta_1 g_t^{NOIL} + \varepsilon_t^{NOIL}$$

$$E\left[g_{t+1}^{NOIL}\right] = g_G$$

Finally, the government receives proceeds from the production of oil and gas. For simplicity, we refer only to oil revenues within the model, which act as a proxy for the combined production of both oil and gas. Each year, a fixed amount of oil is sold at the prevailing dollar oil price, P_t^{OIL}. There are costs incurred by the oil sector, denominated in kroner and converted at the prevailing USDNOK exchange rate.[4] We assume an effective government

4 Production is fixed at $y = 1$ billion barrels, where we convert the value of gas production into an equivalent number of barrels of oil by value, using figures from Norsk Petroleum. Total costs are fixed at $c = 200$ billion kroner. Costs are set broadly in line with figures from Norsk Petroleum. Revenues are denominated in dollars. While in practice the government receives krone and foreign currency oil revenues and transfers these to the fund, for simplicity we assume dollar revenues are transferred to the fund directly.

take from the profits of the sector of 89 percent, and the tax revenue is constrained to be non-negative.[5]

$$R_t^{OIL} = \max(89\% * [yP_t^{OIL} - c/P_t^{FX}], 0)$$

It should be noted that the modelling of the net cash flow from oil production is fairly crude, since one could reasonably assume that both costs and production levels should be a function of the price. Since the main aim of our note is to provide some simple illustrations of how the risk changes as we expand the portfolio choice problem, not to provide an accurate model of the Norwegian economy, we did not pursue this topic further. Given our simplistic modelling of the net cash flow from oil, and given that it is not obvious how the cash flow should be discounted, we also avoid presenting calculations of the net present value of the cash flow from oil.

iii. Calibration

The set of shocks $\{\varepsilon_{t+1}^{EQ}, \varepsilon_{t+1}^{X}, \varepsilon_{t+1}^{FI}, \varepsilon_{t+1}^{Z}, \varepsilon_{t+1}^{FX}, \varepsilon_{t+1}^{OIL}, \varepsilon_{t+1}^{NOIL}\}$ is drawn from a multivariate normal distribution. The volatilities are set as follows, where the values correspond to the volatility of total equity and fixed-income returns, oil prices and FX returns, and the growth rate in non-oil revenues:

$$\sigma^{EQ} = 16\%$$

$$\sigma^{FI} = 6\%$$

$$\sigma^{FX} = 10\%$$

$$\sigma^{OIL} = 30\%$$

$$\sigma^{NOIL} = 2\%$$

The volatilities and correlations between shocks, as shown in Table 1, are based on estimates from historical data, though the calibration requires a degree of judgement given likely variation in correlations over time.[6] There remains the risk that the volatility or correlation assumptions will not hold in the future and, for correlations, may not even be the appropriate sign. To partially address this issue, we later examine alternative correlation structures and their implications for our results.

5 This is in line with the estimate of the effective government take from the total net present value of the cash flow from oil and gas from the Ministry of Finance Revised National Budget 2016.

6 Volatilities and correlations are based on monthly data from January 1973 to April 2016 for the S&P 500, the Barclays Global Aggregate Fixed Income index, Brent crude and USDNOK, although we put more weight on the latter part of the sample when calibrating the volatilities and correlations. All series are in real terms, adjusted using the US or Norwegian consumer price index. Oil volatility is estimated from changes in the price of the first futures contract. A lower volatility estimate is obtained when using longer-dated (e.g. 12 months ahead) contracts, which may be appropriate to the extent that oil price risk is hedged. We examine the robustness of our findings to a lower volatility estimate in Appendix C. Non-oil revenue calculations are based on government total revenue figures adjusted for taxes on petroleum extraction and income from state direct interests. Following the discussion in Pastor and Stambaugh (2009), we assign a negative correlation between equity and fixed-income return innovations and their respective state variables.

Table 1: Market and revenue correlations

	EQ	X	FI	Z	OIL	FX
EQ						
X	-0.9					
FI	-0.3	0				
Z	0	0	-0.9			
OIL	+0.3	0	-0.2	0		
FX	-0.3	0	+0.2	0	-0.3	
NOIL	+0.2	0	-0.1	0	0.2	-0.1

The parameters of the state-variable processes are set to be consistent with a degree of predictability in fixed-income and equity markets. This is captured through the R^2 statistic from a regression of returns on a predictor variable. We set the constants, persistence parameters and volatilities of the state-variable innovations, such that the predictive regression R^2 is 10 percent for equity returns and 30 percent for fixed income (see Appendix A for more details). The constant terms in the state-variable processes are calibrated to target expected excess returns on bonds and equities. The expected excess bond return (term premium) is set at 0 percent, and the expected excess equity return (equity premium) at 3 percent. The initial value of oil is set at 50 dollars per barrel, and the initial value of the USDNOK exchange rate at 8. The risk-free rate, r_t^f, varies over time but follows a non-stochastic path.[7, 8, 9]

$$E(r_{t+1}^{FI}) - r_t^f = 0\%$$

$$E(r_{t+1}^{EQ}) - r_t^f = 3\%$$

$$P_0^{FX} = 8$$

$$P_0^{OIL} = 50$$

$$r_1^f = -0.7\% \quad , \qquad r_{10}^f = 1\%$$

The initial value of government expenditure is set equal to 1,200 billion kroner,[10] and the initial value of non-oil revenues is set equal to 990 billion kroner.[11] The growth rate of government expenditure and non-oil revenues

7 In year 1, the risk-free rate is set in line with the level implied by one-year Fed Funds futures, adjusted using Consensus inflation forecasts at the one-year horizon. The rate then linearly reverts to its long-term value of 1 percent. We use a long-term nominal rate of 3 percent, which is the median projection from the Federal Reserve Economic Projections, less a 2 percent long-term inflation rate. This implies an average real short rate of 15bps over the next ten years. Given a ten-year TIPS rate of around 15bps, we back out a real term premium of zero today. We also assume that the real term premium is zero going forward. This is lower than the historical average realised real term premium, but since the historical average was realised partly in periods when the correlation between equity and fixed-income returns was positive, we adjust the expected term premium down given our assumption of a negative correlation between equity and fixed-income returns. A non-stochastic path for the real interest rate is assumed purely for modelling convenience and means that we are likely to underestimate the risk associated with all investments in the model, especially from year 10 onwards.

8 The risk-free rate and excess return assumptions imply that the total expected real return on a 60/40 equity/fixed-income portfolio is 1.1 percent in the first year, which increases linearly to 2.8 percent for year 10 and beyond. The average total expected real return on the 60/40 portfolio over the ten-year period is 1.95 percent.

9 The equity premium is set at the conservative end of the 3-4 percent range estimated in NBIM Discussion Note 1/2016. Oil and FX prices are set approximately in line with current values at time of writing.

10 This is set roughly in line with central government total expenditure in the 2016 Fiscal Budget.

11 This is set such that the initial non-oil deficit is on average equal to 3 percent of the initial fund value, approximately equal to the cyclically-adjusted non-oil deficit in 2016.

is set at 1.5 percent per annum (on average for non-oil revenues). The persistence of the non-oil revenue growth process is set at 0.5:[12]

$$G_0 = 1{,}200$$

$$R_0^{NOIL} = 990$$

$$g_G = 1.5\%$$

$$E\left(g_{t+1}^{NOIL}\right) = 1.5\%$$

$$\delta_1 = 0.5$$

4. The asset allocation problem

The aim of our framework is to capture the main features of the current institutional arrangements surrounding the fund. Figure 4 shows a stylised representation of income and spending flows. Two of the income sources, fund return and net oil revenues, are transferred to the fund directly. Each year, an amount equalling the non-oil budget deficit is transferred from the fund. This amount is determined through a "spending rule" where, over time, these transfers should correspond to the expected real return of the fund. These transfers are combined with the non-oil tax revenues to finance government expenditure.

Figure 4: Stylised schematic of flows between government and the Government Pension Fund Global

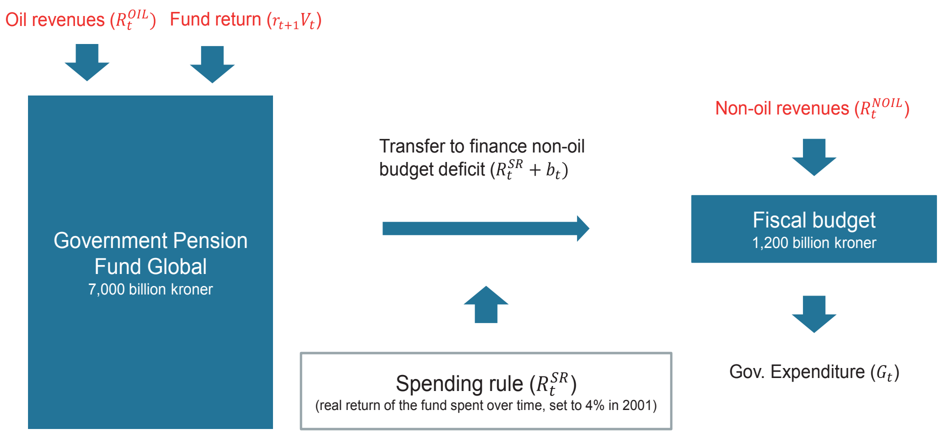

The government spends a krone-denominated amount each year, G_t, while receiving revenues in the form of taxes on oil-related activity, R_t^{OIL}, non-oil related activity, R_t^{NOIL}, and the spending rule amount from the fund, R_t^{SR}. The initial fund size is 875 billion dollars (7,000 billion kroner), which is allocated to equity and fixed-income markets, based on a choice of equity share. The fund generates income from returns on its investments. A target equity

12 The growth rate is a conservative forward-looking estimate of real GDP growth. The persistence parameter is based on an estimate of a first-order autoregression of annual mainland Norwegian real GDP growth.

share is chosen at the start of the simulation and remains unchanged over the simulation horizon. The fund is fully rebalanced to the target equity share each year. Each year, the government partially covers its total expenditure using non-oil tax revenues and a spending rule that is sourced from the fund. The spending rule income is defined as a proportion of the fund's value, V_t, set at 4 percent of the value in a given year:

$$R_t^{SR} = 4\% * V_t$$

The total transfer from the fund is equal to the spending rule income, R_t^{SR}, plus a "buffer" term, b_t:

$$G_t = R_t^{NOIL} + (R_t^{SR} + b_t)$$

This reflects the amount needed to finance the government spend when non-oil revenues and the spending rule over- or undershoot the required total. As shown above, R_t^{NOIL} is subject to shocks, and this can cause unexpected increases or decreases in the income available to finance total government expenditure. As a result, when the government undershoots its spending commitment, it needs to source additional income from the fund, reflected in b_t, where a positive value denotes an additional withdrawal above the 4 percent spending rule. Similarly, when total expenditure is lower, the unused amount under the 4 percent rule can remain in the fund (b_t is negative). The government hence uses the fund as a buffer to ensure that fixed spending commitments can be fulfilled. Taking the revenue sources, spending and fund together, the fund value, V_t, evolves according to the following process:[13]

$$V_{t+1} = V_t(1 + r_{t+1}) + R_t^{OIL} - R_t^{SR} - b_t$$

The value of the fund reflects the total return on its value in the most recent period, $r_{t+1}V_t$, plus the transfer of oil revenues, less the spending rule income stream and buffer, and is naturally a function of the target equity share.

Throughout the analysis, we present the probability that the fund falls significantly in value over the ten-year simulation horizon. Our framework is set out in such a way that that the risks facing the government are all captured in the fund size distribution. The government needs to withdraw/save additional funds from/in the fund in order to maintain a fixed level of spending, which implies that a persistent inability of the government to finance spending reflects directly in a drawdown in the level of the fund. As a result, the risks associated with the financing of expenditure, and risks associated with investment returns and oil revenues, are all captured by the amount of capital in the fund at a given time. The size of the fund is therefore a sufficient statistic for capturing the risks of varying the equity share in the broader government problem.

We capture this risk by calculating the probability of a 50 percent or greater decline in the fund relative to its initial value over the ten-year simulation

13 This is the dollar value. Currency conversions for spending rule income and the transfer terms are omitted for ease of notation.

horizon, assuming that an important objective is to maintain the international purchasing power of the fund. This metric captures the downside risks for the government, in particular the scenario where simultaneously low tax revenues, oil prices and market returns force the government to draw down heavily on the fund to avoid reducing its expenditure levels.[14]

In this section, we use the simulation set-up to examine the effect of introducing revenue and spending considerations into the asset allocation decision. Using the structure described in the previous section, we simulate data at a yearly frequency for equities, fixed income, oil, FX and government revenues. We simulate 100,000 alternative paths and track the value of the fund in each path under different equity share choices. We begin with a stripped-down version of the model, considering the fund's investments in isolation, then progressively build layers of income and spending, and demonstrate the effect of these additions.

We consider three cases:

i. *Case 1: Financial portfolio only*

$$V_{t+1} = V_t(1 + r_{t+1})$$

ii. *Case 2: Financial portfolio with oil revenues and spending rule (flexible spending)*

a. $$V_{t+1} = V_t(1 + r_{t+1}) + R_t^{OIL}$$

b. $$V_{t+1} = V_t(1 + r_{t+1}) + R_t^{OIL} - R_t^{SR}$$

iii. *Case 3: Financial portfolio with all revenues and spending (smoothed government spending)*

$$G_t = R_t^{NOIL} + R_t^{SR} + b_t$$

$$V_{t+1} = V_t(1 + r_{t+1}) + R_t^{OIL} - R_t^{SR} - b_t$$

In the first case, we consider the simulation of the fund in isolation, that is, without any revenue or spending considerations. We examine the distribution of possible fund values over the simulation horizon, and show how expected values and risks change under alternative target equity share assumptions. In the second case, we show how the distribution of outcomes changes when adding oil revenues and the spending rule into the model. We consider two sub-cases in order to illustrate the contributions from oil revenues and the spending rule, where only oil revenues are added in Case 2a, following which the spending rule is introduced in Case 2b. In Case 3, we show the effect of adding non-oil revenues into the model. We show how the impact of changes to the equity share differ within these three

14 In practice, there could be other options: the government may choose to spend less or borrow to finance a deficit, although the latter has a similar effect to increasing the equity share in the financial portfolio as long as the interest rate the government borrows at is similar to the average yield of the fixed-income part of the financial portfolio.

environments, and also how the results depend on the assumptions we have made in calibrating our framework.

While the calibration of the framework tries to be as realistic as possible, it remains a simple and stylised representation of reality, and the limitations of this modelling approach should be kept in mind. In particular, the calibration is based on estimated moments using historical data as a guide, and is deliberately reduced-form in its design. This means that the usual caveat, that the past may not be a good guide to the future, holds. The model also abstracts from complicated economic mechanisms that underpin the interaction and dynamics of fiscal spending, monetary policy, non-oil and oil revenues and exchange rates, which may have an influence on our findings. Our objective is to illustrate the trade-offs that ought to be considered by the fund owner, rather than to advise on a specific equity share. It is therefore more appropriate to focus on how the risks vary for a given change in the model, and/or how an increase or reduction in the target equity share or other policy measure would amplify or mitigate these risks.

i. Case 1: Financial portfolio only

The first case we consider omits any role for revenue or spending, to provide a benchmark upon which to build our analysis. Here, the value of the fund simply evolves in line with the returns generated through the equity and fixed-income investments:

$$V_{t+1} = V_t(1 + r_{t+1})$$

Figure 5 shows the distribution of fund values in each period over the simulation horizon, assuming a 60 percent equity share. The black line shows the expected value of the fund each period, the dark band covers the inner 50 percent of the probability mass of the distribution, the medium band together with the dark band covers 80 percent, and the light band together with the two inner bands covers the 90 percent interval.

Figure 5: Distribution of fund value over a ten-year period (60 percent equity share)

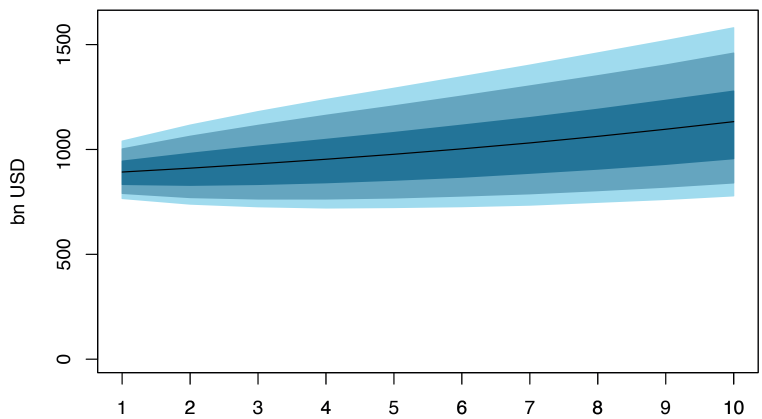

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year.

The expected value of the fund at the ten-year horizon is approximately 1,100 billion dollars, though there is considerable dispersion in fund values around this expectation. The expected value and dispersion of fund values naturally depend on the relative weights of equity and fixed-income investments. This can be seen in Figure 6, which shows how the distribution of fund values at the ten-year horizon changes depending on the target equity share. Clearly, a change in the target equity share has a marked impact on the terminal distribution.

As shown in Figure 7, for each incremental increase in the equity share, there is a corresponding increase in the mean value. Given the higher expected return on equities relative to fixed income, a higher allocation to equities is associated with a higher terminal level, on average. The higher expected return is not without its cost, however. The fund distributions show a clear widening of the distribution as the equity share is increased, reflecting the higher volatility associated with equity returns.[15] The case where the fund is considered in isolation provides a benchmark on which we can build government income and spending considerations. In the benchmark case, there are no transfers to or withdrawals from the fund, and so its performance purely reflects the expected returns and risk of the financial assets.

Figure 6: Distribution of fund value at ten-year horizon for alternative equity shares

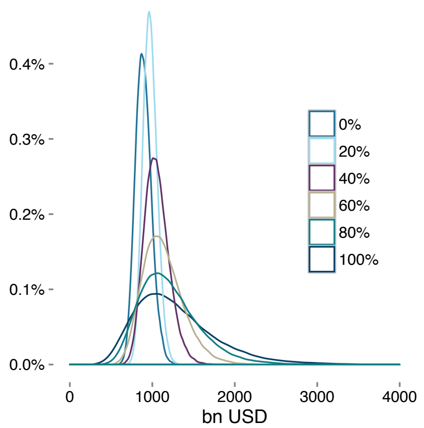

Figure 7: Expected total fund value at ten-year horizon by equity share

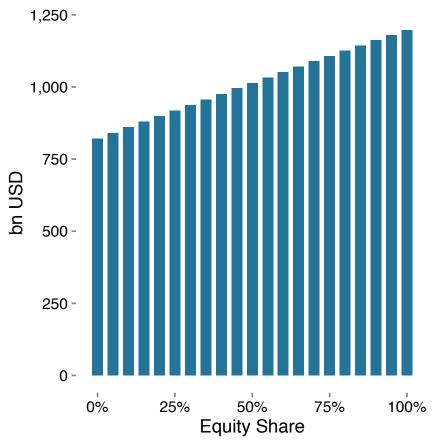

ii. *Case 2: Financial portfolio with oil revenues and spending rule (flexible spending)*

The previous case provides the foundation upon which to build the first layer of government structure and to examine how this changes the problem. Next, we show that adding oil revenues and the spending rule can increase uncertainty around future fund values substantially. We consider two sub-cases within Case 2. First, in Case 2a, we show the effect of including only the transfer of oil revenues into the fund, alongside the financial portfolio. In

15 The mode of the distributions does not shift rightward with the equity share in the same way as the mean. This is because the variance of the distributions does not scale in the same way with the equity share due to the assumed negative correlation between equity and fixed-income returns.

Case 2b, we then show the effect of adding transfers out of the fund defined by the spending rule $R_t^{SR} = 4\% * V_t$:

a. $$V_{t+1} = V_t(1 + r_{t+1}) + R_t^{OIL}$$

b. $$V_{t+1} = V_t(1 + r_{t+1}) + R_t^{OIL} - R_t^{SR}$$

Figure 8 shows the distributions of fund value at the ten-year horizon for Case 1, with only the financial portfolio, and Case 2a, where oil revenues are included and transferred into the fund. The distributions are based on a 60 percent equity share. Naturally, the inclusion of oil revenues paid into the fund increases the expected value of the fund over the simulation horizon, to around 1,400 billion dollars. On average, the tax revenue from oil production is around 25 billion dollars each year, and this is reflected in the rightward shift of the Case 2a distribution relative to Case 1.

Figure 8: Distribution of ten-year fund value, Case 1 vs Case 2a, 60 percent equity share

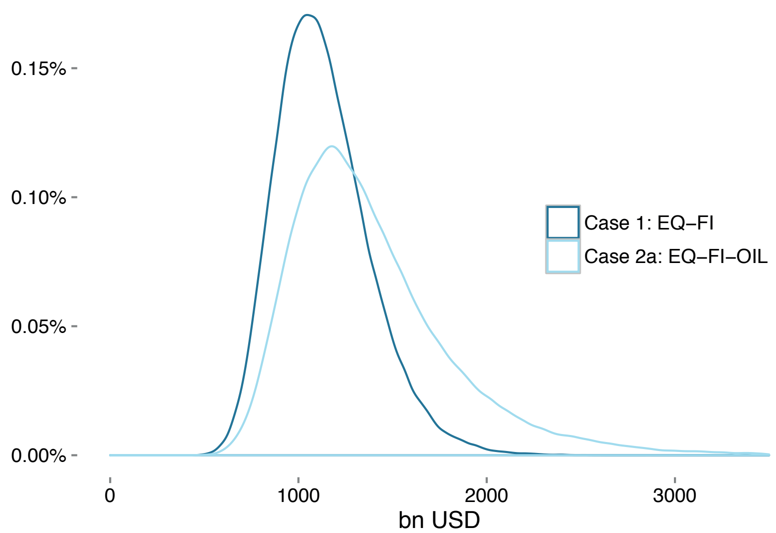

There is also a large increase in the dispersion of values, which can be more easily observed in Figure 9, where the cases are defined in a such a way that the mean value is aligned across cases. This is done by scaling the Case 1 distribution by the ratio of Case 2a and Case 1 means.[16] The increased width results from the addition of the volatile oil income stream into the evolution of the fund's value. The 30 percent volatility in oil prices feeds directly into volatility in tax from oil revenues, leading to a much more dispersed set of outcomes for the fund compared to the financial assets case. This dispersion is further increased due to the positive correlation between equity returns and oil price changes.[17] Overall, Case 2a shows that there are potentially large effects on the distribution of future fund values and hence the risk profile of the government's total portfolio by taking oil revenues into account.

16 The addition of an inflow into the base case will increase the mean value of the fund, and so the width of the distribution in dollars will be naturally bigger. Through scaling by the ratio of means, we can show the distribution adjusted for this effect.

17 The effect is offset slightly by the negative correlation between fixed-income returns and oil prices, though the lower volatility of fixed-income returns means that the equity-oil correlation is more prominent.

Figure 9: Distribution of ten-year fund value, Case 1 vs Case 2a, 60 percent equity share, adjusted means

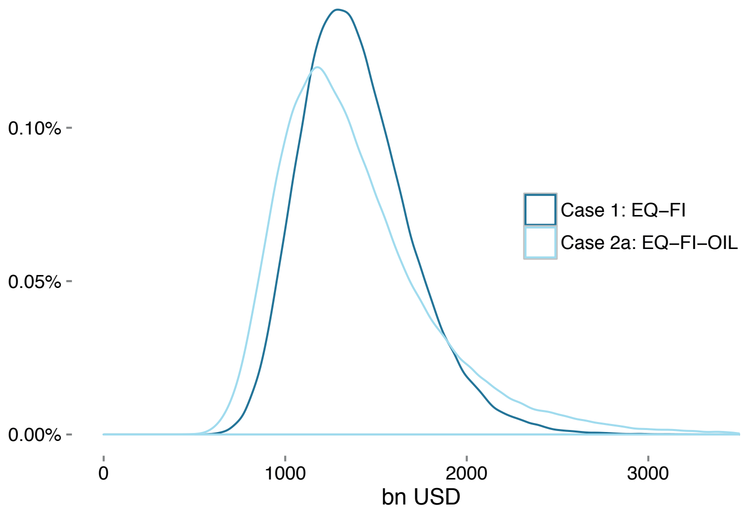

Next, we add the spending rule, defined as 4 percent of the fund's value in a given year, together with oil revenues. The spending rule takes approximately 40 billion dollars from the fund on average each year. This means that around 15 billion dollars is withdrawn in net terms on average across the simulations.[18] Given there is now an outflow from the fund in Case 2b, the distribution shifts to the left relative to Case 1 and Case 2a, where the expected value in Case 2b is approximately 970 billion dollars. Given the assumption that oil prices are flat on average, the net withdrawals from the fund persist in our simulations, and so fund values tend to be lower over the simulation horizon.

Figure 10: Distribution of ten-year fund value, Case 1 vs Case 2a vs Case 2b, 60 percent equity share

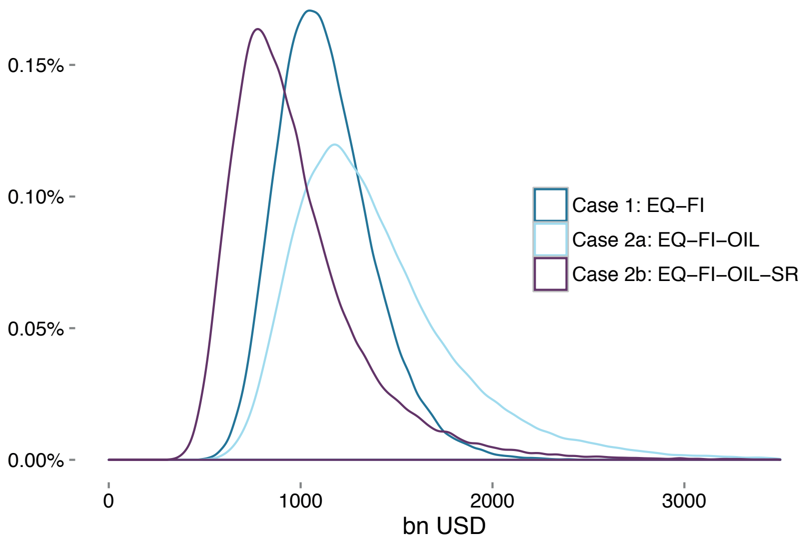

We again adjust the distributions to ensure the same means. When comparing the dispersion of the distributions accounting for the net outflow, however, there is little change when moving from Case 2a to Case 2b. The

18 This is somewhat higher relative to forecast withdrawals from the fund in 2016, due to our assumption of a higher spending rate of 4 percent.

effect of this is shown in Figure 11, where the means of the distributions have been aligned using the same methodology as before. There is only a slight further increase in the dispersion of the distribution, reflecting the small extra variability introduced by the spending rule. It should be noted, however, that the *dollar* volatility is significantly smaller when adding the spending rule, as seen in Figure 10.

Figure 11: Distribution of ten-year fund value, Case 1 vs Case 2b, 60 percent equity share, adjusted means

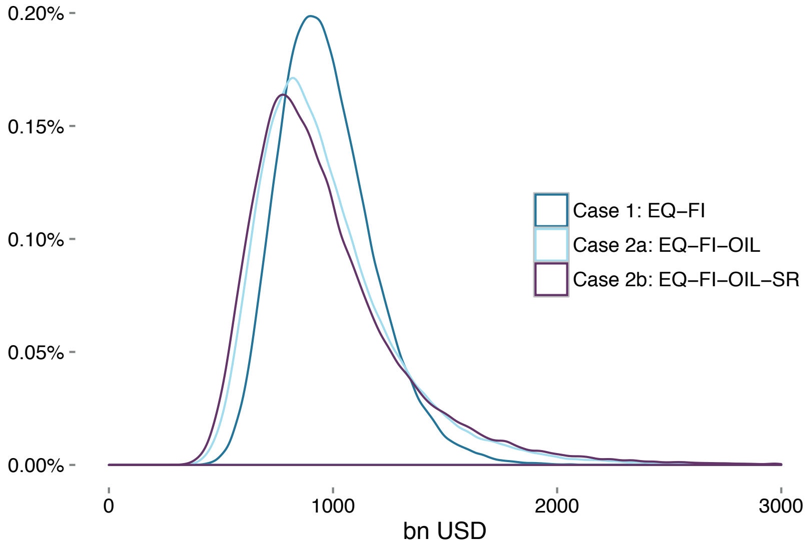

Overall, the introduction of oil revenues and the spending rule has a large effect on the mean and variance of the fund distribution. The strict use of the spending rule in Case 2b implies quite procyclical transfers from the fund, however, and has perhaps somewhat unrealistic implications for government expenditure. Within the framework outlined earlier, we showed that government expenditure is financed through non-oil revenues and the spending rule. The 4 percent rule used in Case 2b therefore implies that government expenditure is quite volatile.[19] For example, if large falls in oil prices and revenues coincided with poor fund performance, perhaps due to a weak global economy, then government expenditure would need to be reduced substantially in line with lower spending rule income. It is arguably more realistic to model a less variable government expenditure stream, and we address this in the next section where we add the final layer of the government problem.

iii. Case 3: Financial portfolio with all revenues and spending (smoothed government spending)

In this section, we include the remaining components that complete the full model specification. We introduce government expenditure, G_t, that is financed through the spending rule as well as non-oil revenues:

$$G_t = R_t^{NOIL} + R_t^{SR} + b_t$$

19 Non-oil revenues are positively correlated with the spending rule amount, with a correlation around 0.2, and are therefore unlikely to offset changes in the spending rule.

$$V_{t+1} = V_t(1 + r_{t+1}) + R_t^{OIL} - R_t^{SR} - b_t$$

$$G_t = G_0(1 + g_G)^t$$

The evolution of the value of the fund is similar to before: oil revenues are transferred into the fund, and the spending rule draws from the fund. An importance difference in Case 3 is the introduction of the "buffer" term, b_t . This reflects the residual of government expenditure when non-oil revenues and the spending rule over- or undershoot the required total. The buffer spending $b_t = G_t - R_t^{NOIL} - R_t^{SR}$ ensures that G_t grows in line with *expected* GDP growth (i.e. spending values are fixed), with the volatility coming from non-oil tax revenues being directly reflected in the fund size distribution. This allows us to offset the procyclicality of transfers from the fund that follows from a strict application of the spending rule, and reduce the counterfactually large variation of government expenditure that this would lead to, albeit at the cost of perhaps smoothing government expenditure too much. Since the current level of government expenditure is around 3 percent rather than 4 percent of the fund's value, b_t starts out as negative in our simulations.

Figure 12 shows the mean-preserved distribution of fund values across Cases 1, 2b and 3. There is a further widening of the distribution, reflecting the added volatility from non-oil revenues, and the positive correlation between the non-oil revenues with both oil revenues and equity returns. Despite volatility and correlation not being set at high levels, there is a sizeable difference due to the use of the fund as a buffer.

Figure 12: Distribution of ten-year fund value, adding non-oil revenues, 60 percent equity share, preserved means

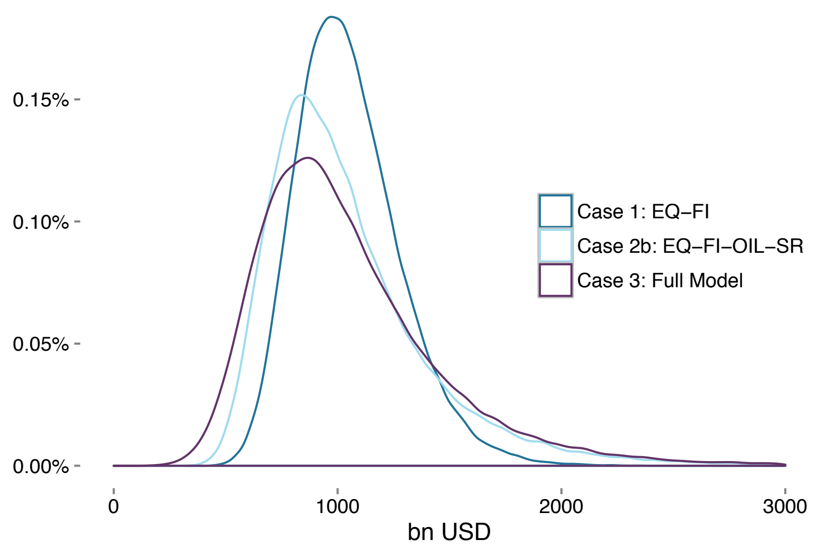

The effect of smoothing government expenditure results in a wider distribution of fund size, and this effect accumulates over time. Figure 13 shows the distribution of fund values in each period over the simulation horizon, assuming a 60 percent equity share. Compared to Figure 5 which considered the financial portfolio in isolation, the fan is wider and widens at a faster rate each period. While this reflects a number of features in the model, in particular the random-walk assumption of oil and FX prices, it also reflects the increased long-term volatility associated with the shorter-term fixing of government consumption.[20]

Figure 13: Distribution of fund value over a ten-year horizon, 60 percent equity share

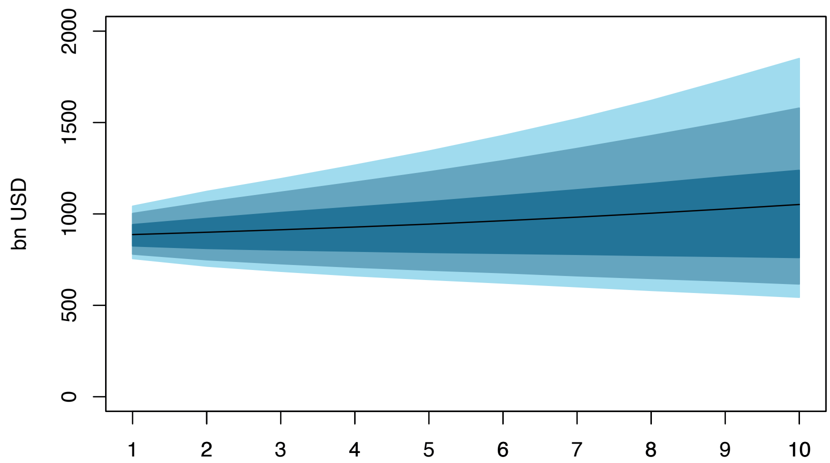

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year.

We can show the differences across distributions for different equity shares by examining the likelihood of a large reduction in the value of the fund over the ten-year horizon. Figure 14 shows the probability that the fund value falls below 50 percent of its starting value at the ten-year point for different choices of equity share. In Case 1 with only the financial portfolio, this probability is negligible, only rising above zero for high equity share choices. When the equity share reaches high levels, the likelihood of a shortfall increases due to the added volatility from the large equity allocation.

When adding oil revenues and the spending rule into the allocation problem in Case 2b, the shortfall probability curve remains low for equity shares of 50 percent and below, reflecting the procyclicality of the spending rule. For higher equity shares, the shortfall probabilities increase substantially relative to Case 1, reflecting the additional volatility from oil prices combined with the positive correlation between equity returns and oil price changes.

There is an upward shift in the shortfall curve when moving to Case 3, the most realistic case. This is the result of the additional volatility introduced from variation in non-oil revenues, and the positive correlation across equities, oil revenues and non-oil revenues. In Case 3, government expenditure is supported by additional withdrawals when the fund falls to low levels, whereas in Case 2b the rule implies that spending from the fund would be reduced. Therefore, in Case 3, there is an increased risk of being unable to finance the required level of fiscal spending, through falls in non-oil revenues occurring at the same time that oil revenue transfers into the fund also diminish, and also tending to coincide with low equity returns. These combined effects act to produce large fund drawdowns and to accentuate

20 The profile and width of the fan are maintained when turning off the volatility of the oil price and the exchange rate.

the left tail of the distribution. There is a pronounced U-shape of the shortfall curve.[21] This reflects the positive correlation between equity returns and oil income, which acts to reduce the level of the equity share that provides the lowest shortfall risk. If the equity share is set too low, however, the shortfall risk rises again, since the higher share of fixed-income investments is associated with a lower expected return and therefore lower expected fund value over the simulation horizon.

Figure 14: Probability of 50 percent fall vs initial value at ten-year horizon

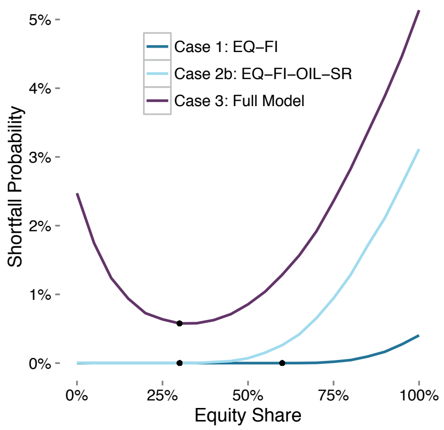

Figure 15: Probability of 50 percent fall vs initial value at ten-year horizon, negative oil-equity correlation

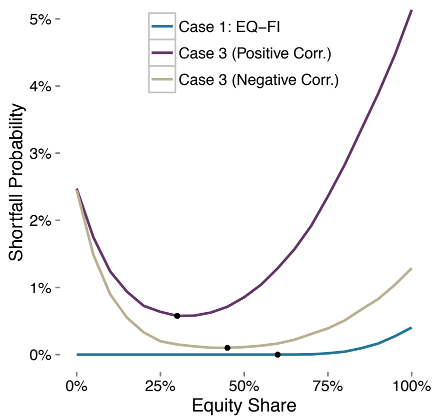

It is worth emphasising that the level and position of the shortfall curves and distributions, and the relative changes across cases, are a direct result of the chosen model calibration. One particular assumption that is important is the positive correlation between oil and equities. Naturally, given a different assumption, the introduction of oil into the problem could have different implications. We explore an alternative in Figure 15, which shows how the shortfall probability changes if we assume a correlation between oil and equity shocks of -0.3 instead of +0.3. There is still an upward shift in the shortfall curve due to the higher overall fund volatility, but now the effect of increasing the equity share changes. At a 0 percent equity share, the risk is naturally uninfluenced by the change in correlation. The effect of increasing the equity share from this point is different, however. The negative correlation means that equities can act as a mild hedge against oil volatility, and so the shortfall likelihood is reduced at a faster rate for initial increases in the equity share. The shortfall risk still increases for higher equity shares, however, due to added equity risk.

We can also show the full model distributions in cumulative probability space. Figure 16 shows the cumulative probability distribution at the ten-year horizon, for the full model, for a range of equity shares. The y-axis measures the likelihood of being at a given fund value, shown on the x-axis, at the ten-year point. As shown in the shortfall charts, the probability of large declines is highest for high equity shares, though this probability is also non-zero for lower equity share choices as well. The figure also shows the increase

21 The results are robust to the choice of percentage decline used to calculate the shortfall probability. Appendix B shows the probability by equity share for a range of percentage choices.

in expected return as the equity share is increased, where the slope of the curve decreases as the equity share is increased.

Figure 16: Cumulative probability distribution, full model, ten-year horizon, by equity share

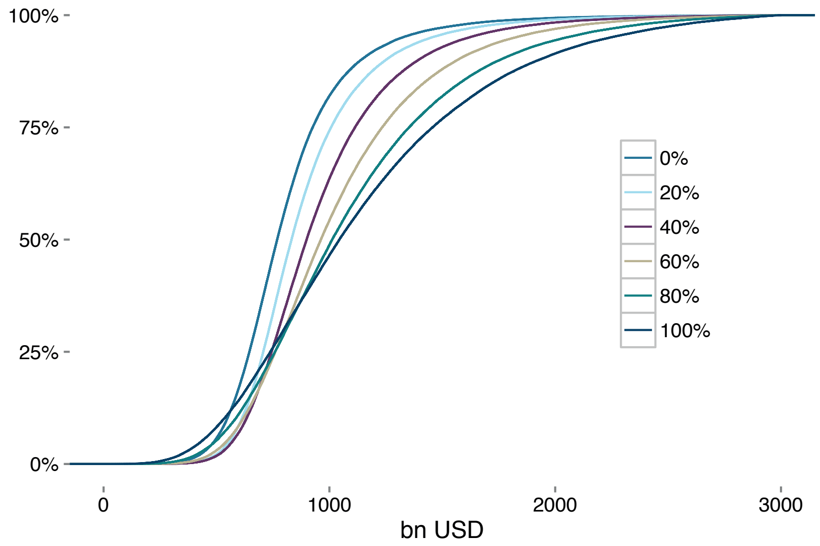

Finally, we can show the implications of adding revenues and spending by estimating how long the fund would be expected to last for a given equity share. We do this by extending the simulations to 100 years and obtaining distributions of fund "longevity". We measure the longevity of the fund as the number of years before the fund runs out.[22] It is possible, and indeed in many cases not unusual, for the value of the fund to remain above zero for the entire 100-year horizon. For simulations where this happens, we set the life of the fund equal to 100 years.[23] As part of the extension in horizon, we include an additional assumption for oil revenues, where production levels (and the associated costs) decline over the long term. Specifically, we impose a path where production levels and costs decline linearly towards zero from year 10 to year 50, such that oil revenues are zero from 50 years onwards.[24] Aside from changes to the path for oil revenues, we use the same calibration as described in Section 3. It should be noted that extending the horizon using the same assumptions as before is not entirely innocent, as there might be co-integrating effects between many of the variables in the framework, such as GDP and equity returns, which are not modelled here. Depending on the simulation horizon, this could change the risk characteristics of the revenue streams, and as a result the modelled fund longevity.

22 We do not capture the possibility that the fund runs out before oil reserves are exhausted, where in theory the fund could start to be replenished by subsequent oil revenues, though the likelihood of this occurring is small.

23 The expected life of the fund is directly relatedly to where the limit is set. In many simulations, the fund would last far beyond the 100-year point, so the expected life is likely to be significantly underestimated. We therefore focus on medians and interquartile ranges rather than expectations in the graphs below.

24 This is based on estimates of remaining reserves from the Norwegian Petroleum Directorate and Norwegian Ministry of Petroleum and Energy (for oil, gas, condensate and NGL measured in standard cubic metre oil-equivalents). We increase this estimate by a factor of two to account for proven and undiscovered reserves, and divide by the average production over the past five years to obtain an estimate of the number of years these reserves will last. Oil production in Norway is likely to last longer than 50 years, but perhaps with a more non-linear trend towards zero.

Figure 17 shows the median value and the 25th and 75th percentiles of the longevity distribution for Cases 1-3, assuming a 60 percent equity share. For Case 1 and Case 2b, the median and interval values all lie at the 100-year point, and the fund survives to 100 years with certainty, which follows from our assumptions on the return processes and spending being a fraction of the fund value in every period (Case 2b).

Figure 17: Fund longevity median and 25th and 75th percentiles. Case 1, Case 2b, Case 3 and Case 3 with negative oil-equity correlation, 60 percent equity share

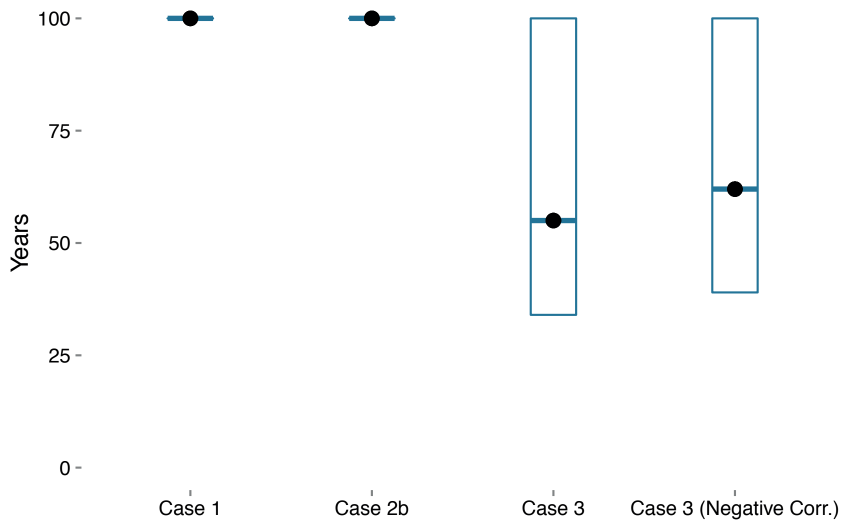

In line with the increased width of the ten-year distributions and the increased shortfall risk, the intervals around expected fund longevity increase when moving to Case 3. This again reflects the additional risk of combined adverse outcomes across equities, oil and non-oil revenues. The effect of adding non-oil revenues, while maintaining fixed government expenditure, is to increase the likelihood that the fund runs out completely over the very long term. If assuming a negative correlation between equity returns and oil price changes, there is a slight increase in the median fund longevity and the lower interval increases, in line with the reduced shortfall risk seen earlier.

Figure 18 shows how longevity varies in Case 3 under different equity share assumptions. The median longevity increases with the equity share, given the increase in the total expected return on the fund. An important factor for the longevity simulations is the long-term difference between spending from the fund and the expected return on the fund. If the expected return, together with oil revenues, exceeds the spending proportion, then the value of the fund can be sustained indefinitely *in expectation*, since expected inflows exceed expected outflows. Similarly, if the expected return is not high enough, there will be a deficit in expectation which, over long horizons, will significantly diminish the value of the fund.

Figure 18: Fund longevity median and 25th and 75th percentiles. Case1 and Case 3 by equity share

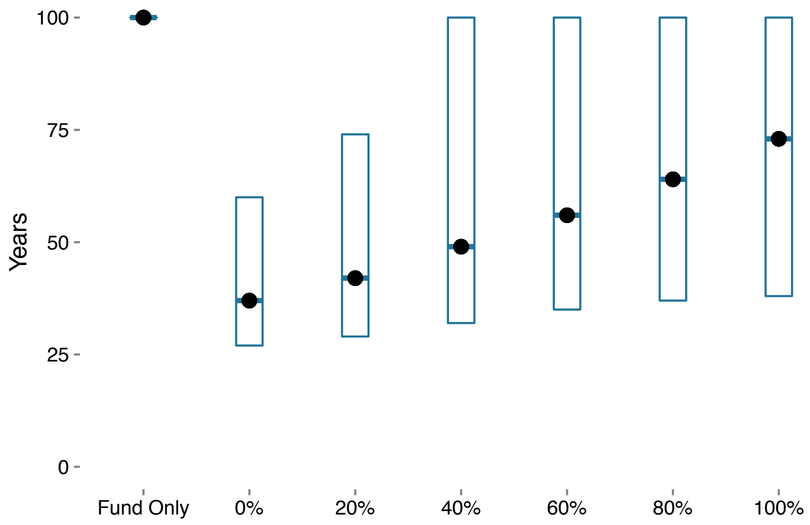

Under our baseline assumptions, the long-term expected return on the fund with a 60 percent equity share is 2.8 percent. The assumption of spending from the fund of 3 percent of the *initial* fund value, growing at an expected rate of 1.5 percent, together with the expected return and oil revenues adding to the growth of the fund over time, implies that spending will average 3 percent initially but decline over time.[25] As a result, at least for 60 percent equity shares and above, the expected return is around the level that will cover the financing of spending. Changes to the equity share directly influence the fund's expected return and, if set too low, can mean that expected returns do not sufficiently cover the spending rule withdrawals. We observe a downward drift in the intervals when moving to equity shares below 40 percent. This highlights an important issue: while shortfall risk may in some cases be reduced by targeting lower equity shares, this is associated with a reduction in expected return on the fund, and the expected longevity is diminished as a result.

Overall, the added revenue and spending layers bring with them additional variability in the fund and downside risks, and changes the asset allocation problem substantially. As a result, any analysis that does not account for spending and revenue behaviour is likely to be understating the risks actually being taken. Clearly, these additions can have a large effect on the possible trajectories for the fund, and the decision regarding an appropriate equity share can look markedly different as a result. There is an increase in risk relative to looking at the fund in isolation, though the nature of the increase and the risk profile by equity share depend on the correlation between oil prices and the investment portfolio. This highlights how the distributions of the fund can change with different assumptions and inputs. While an equity share may be chosen on the basis of given spending behaviour and revenues, and returns and correlations across revenues, spending and markets, a significant structural change would imply that the equity share decision should be revisited.

25 This is a result of oil inflows leading to fund growth. Without oil inflows, the spending proportion is an increasing share of the fund, since it grows at 1.5 percent in expectation.

5. Additional questions

This section uses the fully-constructed model to answer a number of questions. First, given the significant growth in the fund over the past decade, in particular the increase in its relative contribution to government expenditure, we ask what effect this has had on the equity share problem. We show that the increased proportion of expenditure implies additional required volatility in either government expenditure or transfers to and from the fund. We then examine the importance of the level of spending. We show that reducing the spending rule proportion can reduce the shortfall risk substantially and increases expected fund longevity.

We then look at the rebalancing rule within the model and ask how the distribution of outcomes changes if the equity share is allowed to drift over time relative to the choice of equity share. We find that choosing to rebalance has a relatively small impact on the distribution. Next, we ask how the analysis we have presented would change if evaluating the fund in krone rather than dollar terms. Noting the different properties of dollar and krone returns, we show that the risk-reward trade-off can look very different when returns are denominated in kroner.

Finally, we ask how realised paths over the past ten years compare with the distributions generated from the simulations, and consider how the fund would perform under adverse scenarios. We show that the realised path of the fund has been better than expected based on the model assumption we use, and that "stranded asset" and "lost decade" scenarios could reduce the value of the fund substantially.

i. *Question 1: What is the consequence of the large fund growth since 2006?*

As outlined in Section 2, the fund has experienced significant growth over the recent past and is now more than 2.5 times larger than mainland GDP in Norway. Relatedly, transfers from the fund used to finance central government expenditure are now in excess of 16 percent of total expenditure, compared with around 7 percent in 2006.[26] This changing proportion has a meaningful impact on the equity share decision at the current juncture, and we use the full model to demonstrate this. Throughout the analysis so far, we have focused on the variability in the level of the fund as summarising risks borne by the government, assuming that the government keeps its expenditure on a fixed trajectory. We can instead look at the variability in government expenditure that would be required in order to avoid making transfers to or withdrawals from the fund over and above the spending rule. By approaching the problem from this angle, we are able to capture the changing size of the fund over time and its role within government expenditure. Compared to a decade ago, the proportion of expenditure covered by the transfers from the fund has increased roughly 2.5 times over. The implication is that a given level of volatility in the fund will now play a larger role in the financing of fiscal commitments. The consequence of this is that any additional volatility in the fund must now be

26 These values are structurally adjusted, see Figure 2.

manifested either through volatility in government expenditure, or through additional volatility in transfers from the fund.

We use the simulation framework to illustrate this effect. We use a version of the simulation mimicking the situation in 2006, where the fund is a smaller proportion of the fiscal spend, such that the size is 2.5 times smaller than the size in the benchmark calibration. We then allow government expenditure to vary over time in line with non-oil revenues and the spending rule, which we assume is equal to 4 percent of the fund value in a given year:

$$G_t = R_t^{NOIL} + 4\% * V_t$$

This implies that government spending is tied to total non-oil revenues and the spending rule (where the proportion is converted into kroner at the prevailing exchange rate) and inherits any variation in these inputs.

Figure 19: Volatility in expenditure: 2006 vs 2016 fund size for 60 percent equity share

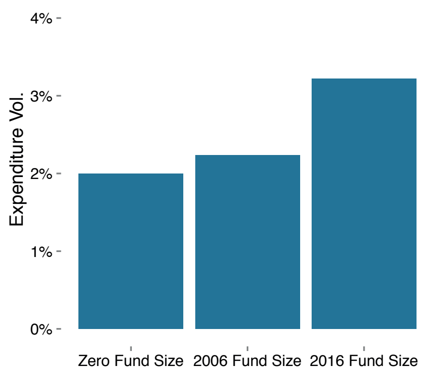

Figure 20: Volatility in expenditure: 2016 fund size, by equity share

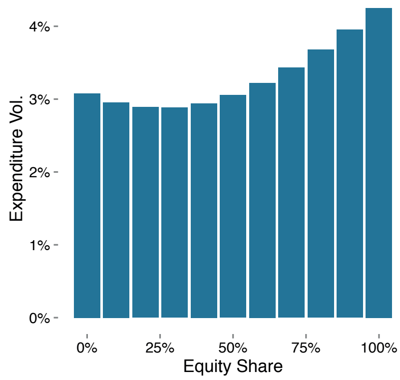

Figure 19 shows the expenditure volatility for the 2006 and 2016 fund sizes against the volatility assuming a fund size of zero, assuming a target equity share of 60 percent.[27] The change in fund size alone implies that the government would need to endure an additional percentage point of expenditure variability. If we were to assume that the 60 percent equity share was the appropriate value in 2006, then holding all else equal, variation in expenditure is higher. This implies that a lower equity share would be desirable if the government did not wish to take on this extra risk. Figure 20 shows that this is not fully possible, however, where the lowest volatility in expenditure is only slightly below 3 percent.

The volatility level and profile by equity share are sensitive to the assumed covariance structure. In particular, the negative correlation across equity and fixed-income returns, and equity and the exchange rate, creates the U-shaped profile shown in Figure 20. Figure 21 shows how the profile changes when we turn off FX volatility, i.e. fully currency-hedging the portfolio to the krone and so removing the effect of correlations of exchange rate movements with equity returns (negative) and fixed-income returns

27 The volatility is measured by the standard deviation of the distribution of expenditures in the first year of the simulation, scaled by the mean value of expenditure in the first year.

(positive). Although hedging the currency risk might seem optimal from this perspective, it should be noted that the assumed correlation between currency and equity returns means that a non-currency-hedged portfolio is a better consumption hedge in the longer run.

Figure 21: Volatility in expenditure by equity share – constant FX

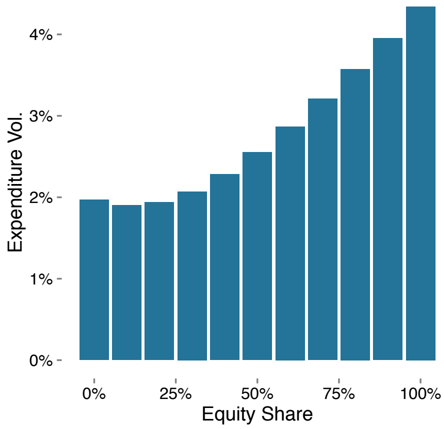

ii. ***Question 2: How does different spending behaviour impact our findings?***

In the fully specified model in Case 3, we assumed that the spending rule draws 3 percent of the initial value of the fund in the first period, and that this amount grows at a rate of 1.5 percent each year. In this section, we ask how the analysis would change if we assume a different level of spending. We revisit our findings when assuming that the spending rule is set at a lower level of 2 percent of the initial fund value, growing at the same rate:

$$R_t^{SR} + b_t = 2\% * V_0 * (1 + g_G)^t$$

Figure 22 compares the shortfall curve for Case 3 assuming a 2 percent initial spend, to the shortfall curves building from Cases 1 to 3 under a 3 percent initial spend, shown earlier in Figure 14. The shortfall curve under the 2 percent spending assumption lies in a lower position, closer to Case 2b under the 3 percent spend assumption. One way to interpret this is that, given our assumed framework and calibration, a reduction in spending from 3 percent to 2 percent can largely offset the added shortfall risk from incorporating non-oil revenues into the equity share problem. The profile of the Case 3 curve remains the same, where higher equity shares tend to steeply increase the shortfall risk due to the positive correlation between the revenue streams and equity returns. The effect of the changed spending assumption again highlights the importance of spending considerations in the equity share decision.

Figure 22: Probability of 50 percent fall in fund vs initial value, 2 percent vs 3 percent initial spend

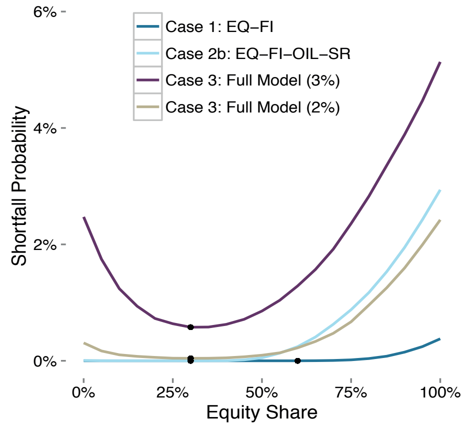

The effect of lower spending also shows clearly when estimating fund longevity. Figure 23 and Figure 24 show the median and interquartile range of fund longevity under 3 percent and 2 percent initial spending, respectively. There is a large effect when spending a lower proportion of the fund. Naturally, the median longevity increases for all equity shares, since withdrawals from the fund are smaller. In Figure 24, under the 2 percent spending assumption, the gap between the expected return of the fund and the spending rate for higher equity shares is sufficiently high that the median longevity is equal to 100 years.[28] Again, the difference between the long-term expected return and spending rate is the dominant factor in determining fund longevity. This difference is increased such that the sustainability of the fund is substantially improved.

Figure 23: Fund longevity median and 25th and 75th percentiles, 3 percent initial spend

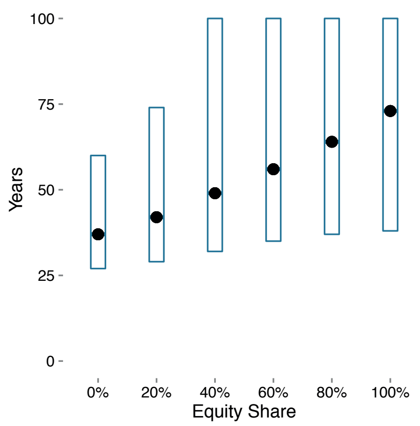

Figure 24: Fund longevity median and 25th and 75th percentiles, 2 percent initial spend

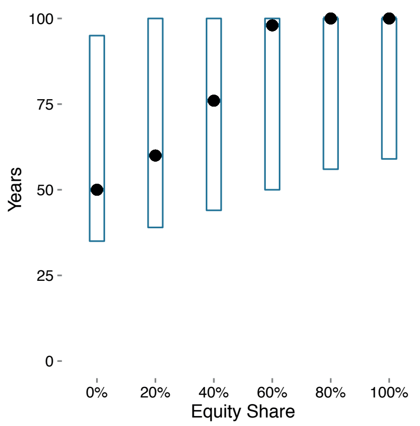

Overall, this section highlights that, while the spending rule decision cannot change the implications of adding outside incomes and spending into the asset allocation framework, it can act to change the degree of these changes. Of course, while a reduction in spending leads to a reduction in shortfall risk and an improvement in expected fund longevity, an increase in spending will have strong effects in the opposite direction.

28 Given that we limit longevity to 100 years, this actually implies that the fund would last somewhat longer.

iii. *Question 3: How important are rebalancing considerations?*

As described earlier, the asset owner chooses a target equity share at the beginning of each simulation, and each year the total fund value is fully rebalanced to this target. This is an approximation of the actual behaviour of the fund, where rebalancing currently only occurs if the equity share drifts beyond 4 percentage points of the 60 percent target share at the end of a given month. The theoretical motivation and empirical significance of choosing to rebalance has been well documented.[29] The simulations provide a different angle for assessing the implications of rebalancing rules, however, where we can examine its importance in a broader context of choosing an equity share over a long horizon. The simulations suggest that, while the decision to rebalance remains a significant one, it is a relatively small consideration compared to choosing the equity share itself, or the spending and income factors described above.

We first show how the distribution of ten-year total fund values changes when varying the rebalancing approach. Figure 25 shows the distribution under the benchmark rebalance case (Case 3), the threshold case where 64/56 percent bands must be hit for a rebalance to occur, and a drifting case with no rebalancing. The different cases are based on a target equity share of 60 percent. The rebalanced distributions differ only slightly compared to the drifting distributions. The rebalancing effect, where the average terminal value is slightly higher, but with a slightly thicker left tail and thinner right tail, is difficult to see. The similarity in distributions partly reflects the addition of income and spending streams, which widen the distribution and make rebalancing effects smaller in relative terms. The difference in distributions remains small when considering the fund in isolation, however, and reflects the volatility calibration used in our model.

Figure 25: Distribution of ten-year fund value under different rebalancing approaches (60 percent equity share)

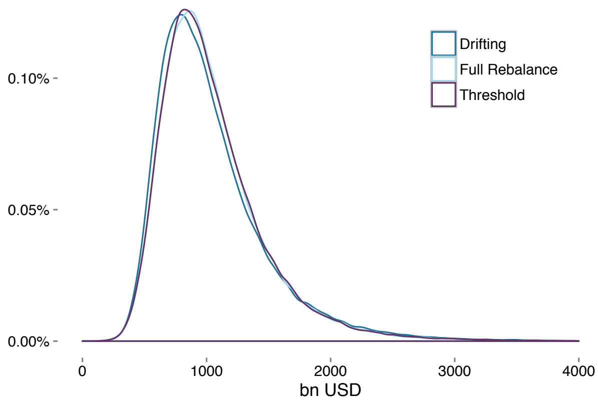

This result does not depend on the level of the target equity share. Figure 26 shows the same set of distributions when the equity share target is set at 10 percent and 90 percent. Overall, while choosing to rebalance can favourably influence the distribution of fund value, the effect is relatively small

29 NBIM Discussion Note 3/2012 provides detailed analysis of rebalancing rules for the fund.

compared to the equity share itself, where the distributions are relatively insensitive to the choice of rebalancing regime.

Figure 26: Distribution of ten-year fund value under different rebalancing approaches (10 percent and 90 percent equity share)

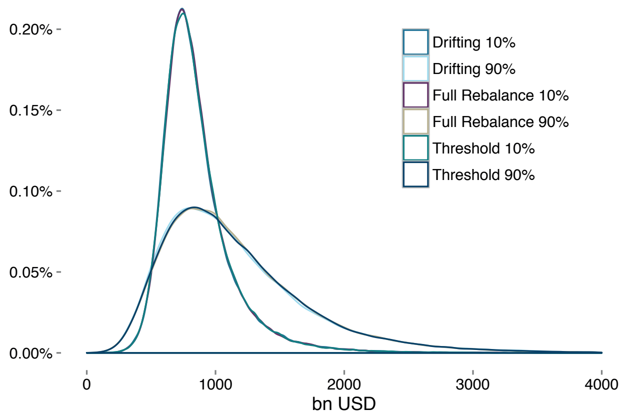

iv. *Question 4: How do currency considerations change the allocation problem?*

Throughout the analysis so far, we have been examining the distribution of fund values denominated in dollars. The effect of this is to largely abstract from exchange rate risk, though the risk is present through krone-denominated spending commitments and krone costs in the oil revenue function. In this section, we briefly show that this decision is significant, since evaluating fund returns in krone terms changes the volatilities and correlations of returns. We can first show this by comparing the Sharpe ratios of equity and fixed-income investments in dollars and kroner. Figure 27 and Figure 28 show how Sharpe ratios vary by equity share. The pronounced arc in Figure 27 results from the reasonably large difference in return volatilities across equities and fixed income, and their negative correlation. As the equity share increases further, the total volatility of returns increases in line with the high equity volatility, which acts to lower the ratio quite sharply.

Figure 27: Sharpe ratio, dollar fund returns

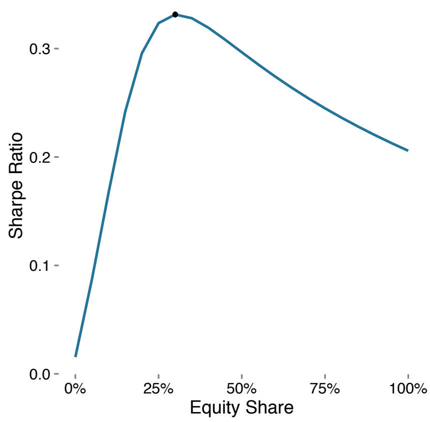

Figure 28: Sharpe ratio, krone fund returns

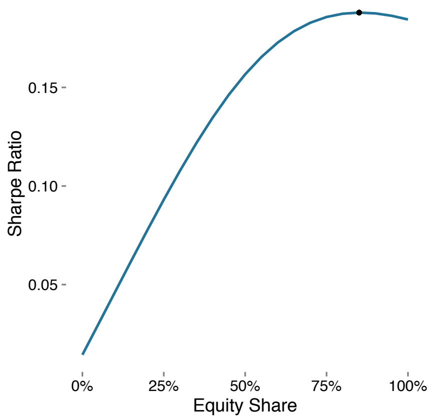

The profile of Sharpe ratios looks quite different when measuring in kroner, as shown in Figure 28, where the curve shifts lower. This reflects the correlation structure defined earlier, where equities and fixed income correlate with FX returns in opposing ways. The correlation with equities is negative, while the fixed-income correlation is positive. This means that, when measuring returns in kroner, the variability of returns is a function of the asset return volatility (in dollars), exchange rate volatility and the correlation between asset (dollar) returns and FX changes. For both equities and fixed income, the addition of FX volatility implies higher volatility of krone-denominated returns, which lowers the Sharpe ratio for all equity shares. However, the negative correlation between equities and FX mostly offsets the increase in volatility, while the positive correlation between fixed income and FX increases the volatility of krone-denominated fixed-income returns further. This makes fixed income much less attractive, given that expected returns remain the same (the exchange rate is flat on average over the ten-year horizon). While the Sharpe ratio still starts to fall as the equity share increases, the effect is much less pronounced compared to the dollar case.

The measurement issue also impacts shortfall probabilities. Figure 29 shows how the probability of a 50 percent shortfall changes whether measuring in dollars or kroner. Naturally, given the higher variance associated with krone returns, the probability is higher for all values of the equity share. There is also a significant change in the slope of the curve, such that the shortfall-minimising equity share is somewhat higher in the krone case, again reflecting the relative unattractiveness of fixed-income investments in the krone case. It is worth noting that the findings in the rest of the note do not change materially when using krone returns. The direction of movement in shortfall probabilities across the different cases is the same, but the changes tend to occur at a higher level compared to the dollar case. In summary, it is common to discuss and report the fund value in krone terms, and using this denomination for the fund can change the problem substantially. As a result, any change to the currency measurement basket of the fund might have consequences for the appropriate investment strategy.

Figure 29: Probability of 50 percent fall in fund vs initial value under dollar and krone measurement, full model

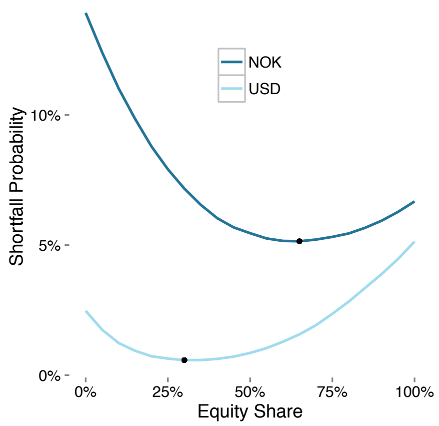

v. *Question 5: How do the ten years from end-2006 compare with simulated values?*

In this section, we examine how the realised path of the fund over the past decade compares to the distributions generated from our model. We compare the realised paths for the total value of the fund, and other variables, with the distributions of simulated values, using the Case 2b framework described in Section 4. The comparison gives an indication of how the evolution of the fund over the past decade compares with our simple set of assumptions, rather than a forecast evaluation or any test of the usefulness of the model framework or calibration.[30] We make only small changes to the calibration, adjusting the starting value of the fund, the initial USDNOK rate and oil price, the initial values of total government expenditure and non-oil revenues, and the risk-free rate path, to be in line with their end-2006 values:[31]

$$V_0 = 280 \qquad r_1^f = 2.8\% \quad , \qquad r_{10}^f = 2.0\%$$

$$P_0^{FX} = 6.25 \qquad P_0^{OIL} = 60$$

We do not change the processes describing the dynamics of the variables within the model, nor the mean values or shock covariance matrix. Throughout, we assume a 60 percent equity share, as an approximation of the target equity share of the fund over the 2006-2016 period. Figure 30 and Figure 31 show the realised paths of the dollar and krone values of the fund, respectively, against the distribution of simulated values from the 2006 version of the model.[32]

Figure 30: Realised vs simulated paths of dollar fund value 2006-2016 (60 percent equity share)

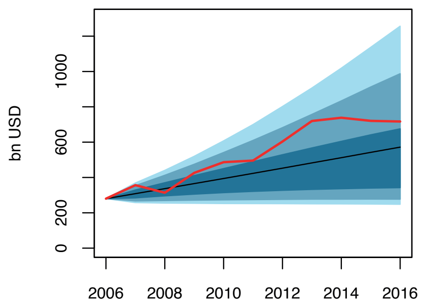

Figure 31: Realised vs simulated paths of krone fund value 2006-2016 (60 percent equity share)

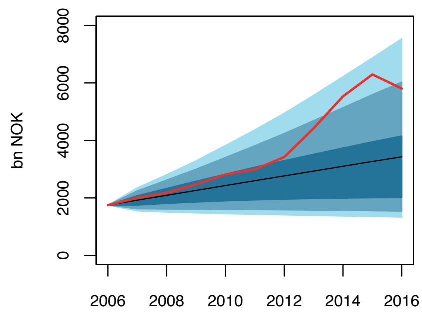

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year. Red line shows path of year-end values.

30 In many cases the calibration of the model either implicitly or explicitly incorporates information from after 2006 and so suffers from a look-ahead bias that invalidates a proper out-of-sample forecast evaluation.

31 The risk-free rate in year 1 is set in line with one-year Fed Funds futures, adjusted by the Consensus one-year inflation forecast. We assume the rate linearly reverts to a long-term value of 2 percent, together with a zero real term premium, implying an average short rate of 2.4 percent over the ten-year period, in line with the ten-year TIPS rate at end-2006.

32 The realised fund and other variables shown in this section are deflated to be in real terms. Figures show year-end values.

For both the dollar and krone values, the fund value is in the upper area of the simulated distribution, and the krone value has reached around the 90th percentile of the ten-year distribution. This suggests that, at least relative to the assumptions of the model, the fund has followed an unexpectedly strong course over the past ten years. There are many factors underlying this growth, and our framework can shed light on the main factors to which it is attributable. These factors include equity and/or fixed-income returns, revenues from oil production, spending relative to non-oil revenues and the spending rule, and currency effects.

Figure 32 and Figure 33 show realised equity and fixed-income returns, respectively, over 2006-2016, in red, relative to the simulated distributions. While there was a large negative return on equities in 2008 amid the financial crisis, returns have been positive on average at around 6 percent per annum over the ten-year period, although the cumulative return has slightly undershot the expected path from the model.[33] This implies that, while realised equity returns have contributed meaningfully to fund growth since 2006, they will not have generated fund growth in excess of the *expected* path. Realised fixed-income returns, on the other hand, have been stronger relative to the expected path over this period. Fixed income has consistently outperformed the expected path, with positive returns in each year apart from 2013. As a result, over 2006-2016, fixed-income returns also have led to the fund growing above its expected path implied by the model.

Figure 32: Realised vs simulated cumulative equity returns 2006-2016 (60 percent equity share)

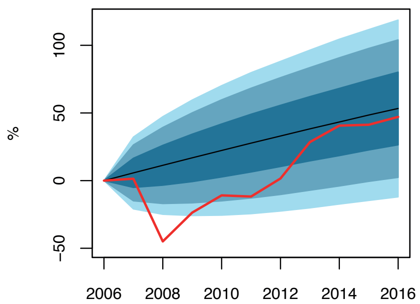

Figure 33: Realised vs simulated cumulative fixed income returns 2006-2016 (60 percent equity share)

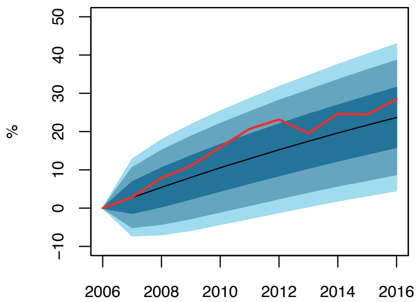

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year. Red line shows path of year-end values.

Next, we compare the realised oil price path to the model distribution. Figure 34 and Figure 35 show model-generated distributions of the oil price and net oil revenues, respectively, and their realised paths. The oil price was, in general, above the model expected value for much of the period up until 2014, despite falling sharply around 2009. This led to increased revenues from oil production and transfers into the fund over this period, as shown in Figure 35. In the past couple of years, however, the oil price has fallen quite sharply, and this will have led to a decline in the amount of oil revenues transferred to the fund. Overall, the unexpectedly strong revenues over the

33 Our model assumes that log returns are normally distributed. This is a strong assumption, as asset returns often have fatter tails than implied by a normal distribution. An indication of this is given in Figure 32, where the 2008 realised returns are outside the model-implied 90 percent confidence interval.

majority of the period led to large inflows into the fund and are one of the main factors underlying the better-than-expected growth in the fund over 2006-2016.

Within our framework, the rise in the fund can also result from relatively low spending levels. In Figure 36, we compare the spending rule fan, defined as 4 percent of the fund value in a given year, to the realised path of transfers from the fund to cover the non-oil budget deficit.[34] For much of the ten-year period, the government has spent less than a 4 percent rule would imply, meaning that less has been transferred from the fund, allowing it to grow at a faster rate. Finally, we look at the effect of currency movements on the fund, where movements in the USDNOK rate are primarily reflected in the krone value of the fund, rather than the dollar value. Figure 37 shows the realised path for USDNOK compared to the model distribution. The changes in the rate were relatively small until 2012, after which there was a large depreciation in the krone, contributing to the steep increase in the value of the fund over this period shown in Figure 31.

Figure 34: Realised vs simulated path of oil price 2006-2016 (60 percent equity share)

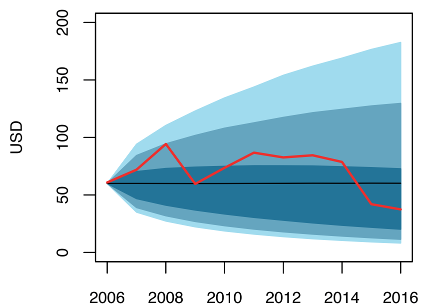

Figure 35: Realised vs simulated paths for cumulative oil revenues 2006-2016 (60 percent equity share)

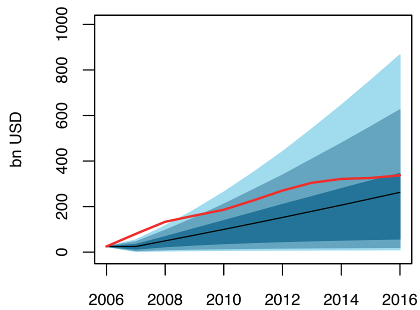

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year. Red line shows path of year-end values.

Figure 36: Realised vs simulated cumulative transfers from fund 2006-2016 (60 percent equity share)

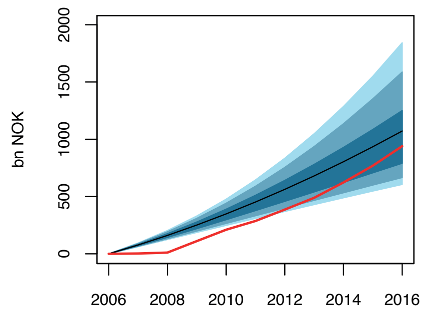

Figure 37: Realised vs simulated paths for USDNOK 2006-2016 (60 percent equity share)

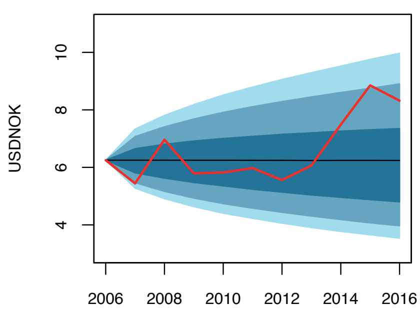

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year. Red line shows path of year-end values.

34 This differs from the chart shown in Section 2 of the fund's contribution to government expenditure. Here, we compare with the realised transfer to cover the non-oil budget deficit, whereas the earlier chart shows the proportion of the structural (cyclically-adjusted) deficit.

Overall, strong oil revenues, low spending and currency effects have been the main factors behind the *unexpectedly* strong fund growth since 2006, more so than returns on the fund's investments. While equity and fixed-income returns have still played a vital role in overall fund growth, these findings once again highlight the importance of the revenue and spending considerations.

vi. ***Question 6: How would the fund perform under adverse scenarios?***

In this final section, we show what path the fund could follow given particularly adverse scenarios for oil prices and investment returns. First, we model a "stranded asset" scenario for oil prices, where we impose large continuous declines in the price of oil over the simulation horizon. We then model a "lost decade" for equity returns, where we impose a path for returns that is in line with what was observed for Japanese stock market during the 1990s. Finally, we impose an adverse path for fixed-income returns, in line with real returns on ten-year US Treasury bonds during the 1970s. The same assumptions for revenues and spending behaviour are used as earlier in the note.

In the stranded asset scenario, oil prices fall steadily and quickly towards zero over the ten-year horizon. We impose an oil price path where each year the oil price falls by 5 dollars, reaching zero in year 10, which is naturally accompanied by quickly declining oil revenues. Figure 38 shows how the path for the oil price in the scenario compares to the distribution of oil prices in the benchmark model. Despite the large variance of oil prices in the simulation, the low terminal value under the stranded scenario is in the lower tail of the distribution of oil prices at the ten-year horizon.

Figure 38: Stranded asset scenario oil price vs benchmark distribution

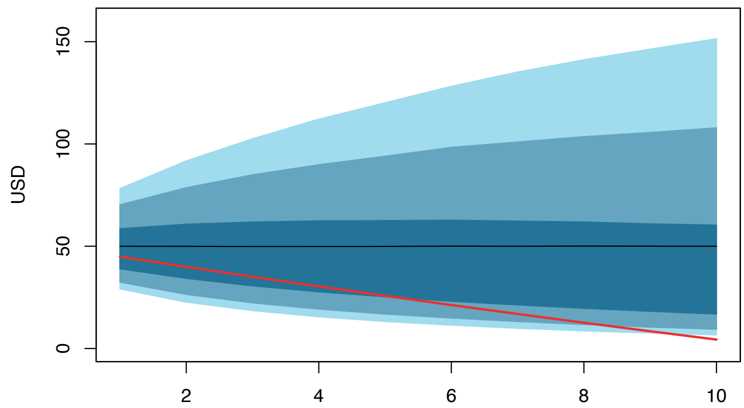

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year.

With the sharp declines in oil prices and the associated revenues, we can show how the total fund performs on average over the simulation horizon. Figure 39 shows the average trajectory of the total value of the fund compared to the benchmark model fan chart. At the ten-year point, the average fund value is almost 10 percent lower than its starting value. The scenario suggests that the fund would not be devastated by falling oil prices, though of course this is an extremely stylised representation of a stranded asset oil price path.

While the decline in the oil price is rapid in the scenario, the fund still receives positive oil revenues on average for many of the years during the simulation horizon. Naturally, a faster decline in the oil price would mean that the fund falls more on average in the scenario. In addition, we do not impose any corresponding equity, fixed-income or exchange rate movements associated with the declining oil price, nor any changes to non-oil revenues. Market changes would likely depend heavily on the nature of the decline in oil prices, which could plausibly generate both upward and downward movements in asset prices. Given the reliance of the Norwegian economy on oil production, we might be more inclined to predict a depreciation of the krone in line with the declining oil price, which would act to increase the krone value of the fund, holding all else equal. These price movements are extremely difficult to predict, making the impact of such a scenario very uncertain.

Figure 39: Distribution of fund value over ten-year period, benchmark distribution vs expected fund path under stranded asset scenario (60 percent equity share)

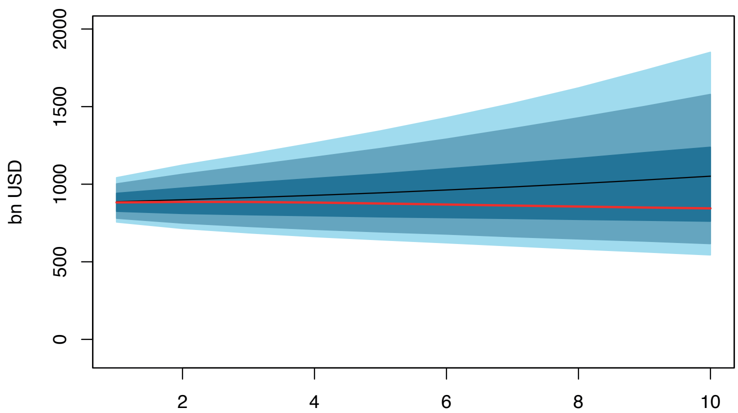

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year.

Next, we consider a lost decade scenario, where we impose equity returns in line with realised values for the Japanese equity market over the period 1990-2000. We assume that fixed-income returns follow the process defined earlier.[35]

Figure 40 shows equity returns relative to the return distributions generated by the model, with realised equity returns shown in red. The realisations lie relatively far into the tails of the model distribution, reflecting significant volatility over the 1990s, with an annual average return of around -3.5 percent.

35 We do not impose realised Japanese government bond returns for the same period, since this would subtract around 400 bps from the ten-year yield and would imply extremely negative yields if taken from today's levels.

Figure 40: Realised vs simulated paths of equity returns under lost decade scenario

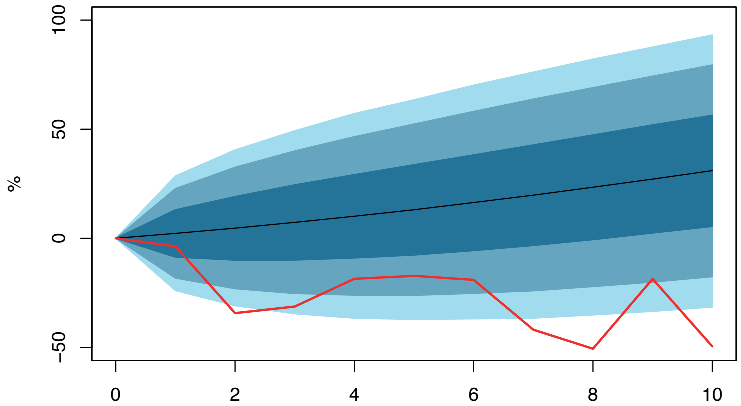

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year.

Figure 41 shows how the value of the fund would evolve based on the imposed equity returns, assuming a 60 percent equity share. Naturally, the value of the fund decreases given the sizeable equity allocation and weak equity performance, where positive fixed-income returns are not enough to offset these falls. This exercise shows that investment performance remains a key consideration alongside spending and revenue considerations, and that large movements in asset values will still lead to large changes in the value of the fund.

Figure 41: Distribution of fund value over ten-year period, benchmark distribution vs expected fund under lost decade scenario (60 percent equity share)

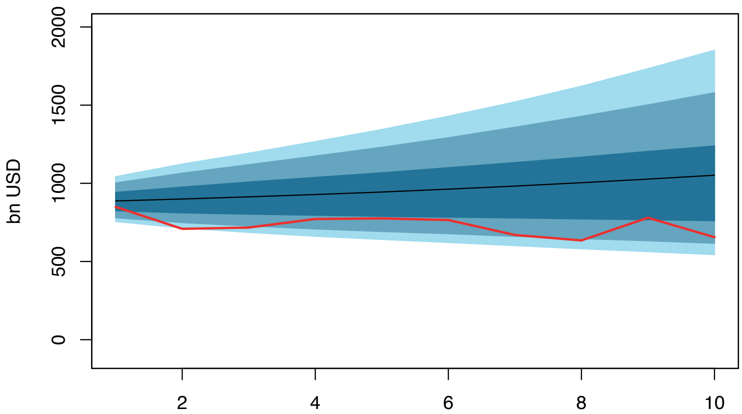

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year.

Finally we model the value of the fund under an adverse scenario for fixed-income returns, again assuming a 60 percent equity share. We impose a path for fixed-income returns in line with those observed on ten-year US Treasuries in the 1970s. Figure 42 shows the returns relative to the fixed-income distribution generated by the model, where the total return over the ten-year period is around minus forty percent. Within the fixed-income scenario, we choose to keep the risk-free short rate path fixed, which implicitly assumes that the rise in the bond yield is driven entirely by an increase in the term premium.[36] Figure 43 shows the mean value of the fund under the fixed-income scenario. The fund value naturally falls on average under the scenario, but the impact of the fixed-income scenario is relatively mild, due to the lower share of the fixed-income portfolio.

Figure 42: Realised vs simulated paths of fixed-income returns under fixed income scenario

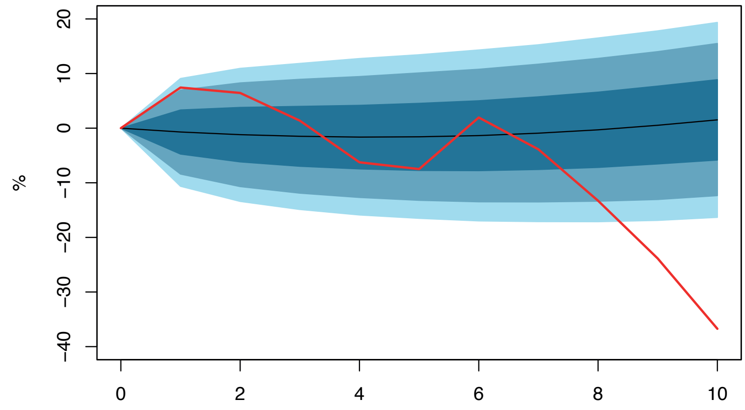

Figure 43: Distribution of fund value over ten-year period, benchmark distribution vs expected fund under fixed income scenario

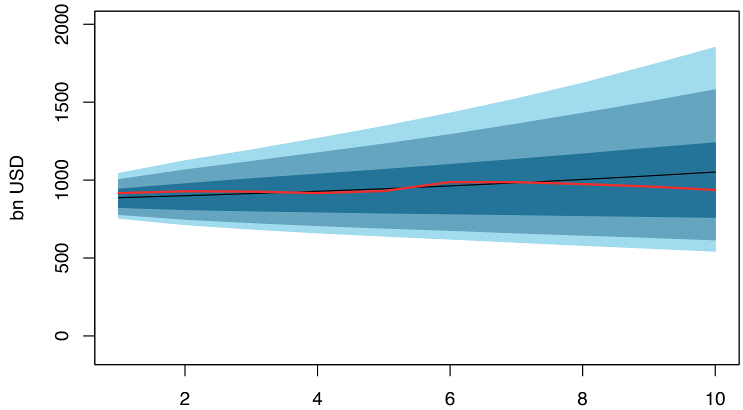

Dark band covers 50 percent probability, medium band 80 percent, and light band 90 percent. Black line shows the expected value each calendar year.

6. Summary

It is common to consider an equity and fixed-income allocation problem in terms of expected returns and risk. These remain important inputs, but there are additional considerations that should be taken into account, in particular other incomes and liabilities, which can change the problem significantly. We develop a simple framework for exploring how the asset allocation problem changes when expanding to include government revenues and spending. Through simulating a large number of paths for the fund, we are able to assess how the expected value and risks change when taking into account income in the form of oil and non-oil revenues, and liabilities in the form of

36 While this is an unrealistic assumption, it simplifies the analysis by avoiding modelling implications of short rate changes for equity returns.

fiscal expenditure commitments. We show that the additional uncertainty introduced by including revenue sources increases the risks to the fund for any equity share, and this implies that any analysis that does not account for spending and revenue behaviour is likely to be understating the risks actually being taken. Our analysis also highlights how the distributions of the fund can change with different assumptions and inputs. In particular, a reduction in the spending liability of the fund can reduce risk substantially and improve the expected longevity of the fund.

References

Benzoni, L., Collin-Dufresne, P. and Goldstein, R. (2007): "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets are Cointegrated", *Journal of Finance*, 62(5), 2123-2167.

Bodie, Z., Merton, R. and Samuelson, W. (1992): "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model", *Journal of Economic Dynamics and Control*, 16(3-4), 427-449.

Cieslak, A. and Povala, P. (2016): "Information in the Term Structure of Yield Curve Volatility", *Journal of Finance*, 71(3), 1393-1436.

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

Cochrane, J. (2014): "A Mean-Variance Benchmark for Intertemporal Portfolio Theory", *Journal of Finance*, 69(1), 1-49.

Cochrane, J. and Piazzesi, M. (2005): "Bond Risk Premia", *American Economic Review*, 95(1), 138-160.

Dahlquist, M. and Hasseltoft, H. (2013): "International Bond Risk Premia", *Journal of International Economics*, 90(1), 17-32.

Ferreira, M. and Santa-Clara, P. (2011): "Forecasting Stock Market Returns: The Sum of the Parts Is More Than the Whole", *Journal of Financial Economics*, 100(1), 514-537.

Heaton, J. and Lucas, D. (1997): "Portfolio Choice in the Presence of Background Risk", *Economic Journal*, 110, 1-26.

Kelly, B. and Pruitt, S. (2013): "Market Expectations in the Cross-Section of Present Values", *Journal of Finance*, 68(5), 1721-1756.

Markowitz, H. (1952): "Portfolio Selection", *Journal of Finance*, 7(1), 77-91.

Merton, R. (1969): "Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case", *Review of Economics and Statistics*, 51(3), 247-57.

NBIM Discussion Note 3/2012: Empirical Analysis of Rebalancing Strategies.

NBIM Discussion Note 1/2016: The Equity Risk Premium.

Norwegian Ministry of Finance Revised National Budget 2016.

Pastor, L. and Stambaugh, R. (2009): "Predictive Systems: Living with Imperfect Predictors", *Journal of Finance*, 64(4), 1583-1628.

Samuelson, P. (1969): "Lifetime Portfolio Selection by Dynamic Stochastic Programming", *Review of Economics and Statistics*, 51(3), 239-246.

Van Binsbergen, J. and Koijen, R. (2010): "Predictive Regressions: A Present-Value Approach", *Journal of Finance*, 65(4), 1439-1471.

Viceira, L. (2001): "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income", *Journal of Finance*, 56(2), 433-470.

Appendix A: Predictability, volatility and correlation calibration

We calibrate the excess return volatility of equities to 16 percent, and fixed-income returns to 6 percent. Using assumptions about the degree of predictability in bond and equity markets, and the persistence of state variables, we can solve for the appropriate shock variances to obtain the desired volatility calibration.

Equity returns are a function of the risk-free rate, state variables and noise, and these components are therefore the sources of volatility. The relative volatility between components can be calibrated by making an assumption about the degree of predictability in returns at an annual frequency. We assume 10 percent for equities (in line with studies such as van Binsbergen and Koijen (2010), Ferreira and Santa-Clara (2011) and Kelly and Pruitt (2013)) and 30 percent for fixed income (see Cochrane and Piazzesi (2005), Dahlquist and Hasseltoft (2013) and Cieslak and Povala (2016)).

$$R^2 = 1 - \frac{\sigma^2(\varepsilon_{t+1}^{EQ})}{\sigma^2(r_{t+1}^{EQ})}$$

$$R^2 = 0.10$$

$$\sigma^2(r_{t+1}^{EQ}) = 0.16^2 = 0.0256$$

Implies that:

$$\sigma^2(\varepsilon_{t+1}^{EQ}) = 0.02304$$

Equity returns total variance:

$$r_{t+1}^{EQ} = r^f + X_t + \varepsilon_{t+1}^{EQ}$$

$$\sigma^2(r_{t+1}^{EQ}) = 0.0256 = \sigma^2(\varepsilon_{t+1}^{EQ}) + \sigma^2(X_t) + 2 * \rho(\varepsilon_{t+1}^{EQ}, X_t) * \sigma(X_t) * \sigma(\varepsilon_{t+1}^{EQ})$$

Risk-free rate has zero variance. Assume that $\rho(\varepsilon_{t+1}^{EQ}, X_t)$. Use $\sigma^2(\varepsilon_{t+1}^{EQ})$ to solve for $\sigma^2(X_t)$

$$\sigma^2(X_t) = 0.00256$$

X_t follows a first-order autoregressive process, for which the variance is:

$$X_{t+1} = \varphi_0 + \varphi_1 * X_t + \varepsilon_{t+1}^x$$

$$\sigma^2(X_t) = \frac{\sigma^2(\varepsilon_{t+1}^x)}{1 - \varphi_1^2}$$

We assume $\varphi_1 = 0.8$, so:

$$\sigma^2(\varepsilon_{t+1}^x) = 0.0009$$

The same method is used for calibrating the components of fixed-income returns where $R^2 = 0.30$ and $\gamma_1 = 0.4$.

We also need to account for separate components of equity and fixed-income returns when calibrating the correlations of shocks. Using the bond-equity correlation as an example, we target the correlation:

$$\rho\left(r_{t+1}^{EQ}, r_{t+1}^{FI}\right) = -0.3$$

The correlation can be decomposed into scaled co-variances of the state variables and noise terms:

$$\rho\left(r_{t+1}^{EQ}, r_{t+1}^{FI}\right) = \frac{cov(Z_t, X_t)}{\sigma\left(r_{t+1}^{EQ}\right) * \sigma\left(r_{t+1}^{FI}\right)} + \frac{cov\left(Z_t, \varepsilon_{t+1}^{EQ}\right)}{\sigma\left(r_{t+1}^{EQ}\right) * \sigma\left(r_{t+1}^{FI}\right)} + \frac{cov\left(X_t, \varepsilon_{t+1}^{FI}\right)}{\sigma\left(r_{t+1}^{EQ}\right) * \sigma\left(r_{t+1}^{FI}\right)} + \frac{cov\left(\varepsilon_{t+1}^{FI}, \varepsilon_{t+1}^{EQ}\right)}{\sigma\left(r_{t+1}^{EQ}\right) * \sigma\left(r_{t+1}^{FI}\right)}$$

Assume that $cov\left(Z_t, \varepsilon_{t+1}^{EQ}\right) = cov\left(X_t, \varepsilon_{t+1}^{FI}\right) = 0$. $\sigma\left(r_{t+1}^{EQ}\right)$ and $\sigma\left(r_{t+1}^{FI}\right)$ have already been defined.

$$\rho\left(r_{t+1}^{EQ}, r_{t+1}^{FI}\right) * \sigma\left(r_{t+1}^{EQ}\right) * \sigma\left(r_{t+1}^{FI}\right) = cov(Z_t, X_t) + cov\left(\varepsilon_{t+1}^{FI}, \varepsilon_{t+1}^{EQ}\right)$$

We then assume entire covariance comes from $cov\left(\varepsilon_{t+1}^{FI}, \varepsilon_{t+1}^{EQ}\right)$ and solve for the correlation:

$$\frac{cov\left(\varepsilon_{t+1}^{FI}, \varepsilon_{t+1}^{EQ}\right)}{\sigma\left(\varepsilon_{t+1}^{EQ}\right) * \sigma\left(\varepsilon_{t+1}^{FI}\right)} = \frac{\rho\left(r_{t+1}^{EQ}, r_{t+1}^{FI}\right) * \sigma\left(r_{t+1}^{EQ}\right) * \sigma\left(r_{t+1}^{FI}\right)}{\sigma\left(\varepsilon_{t+1}^{EQ}\right) * \sigma\left(\varepsilon_{t+1}^{FI}\right)}$$

Appendix B: Shortfall Probabilities

Figure 44 and Figure 45 show shortfall probability curves under the full model, when varying the shortfall percentage between 10% and 90%.

Figure 44: Alternative shortfall probabilities 10% - 40% falls in fund vs initial value

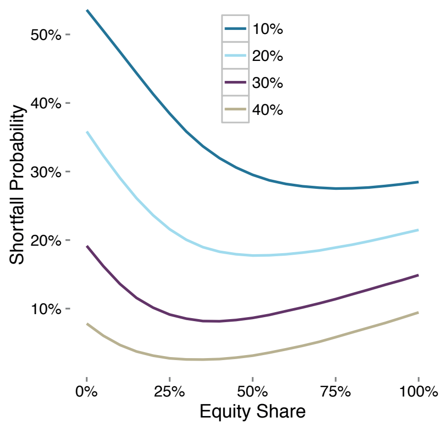

Figure 45: Alternative shortfall probabilities 50% - 90% falls in fund vs initial value

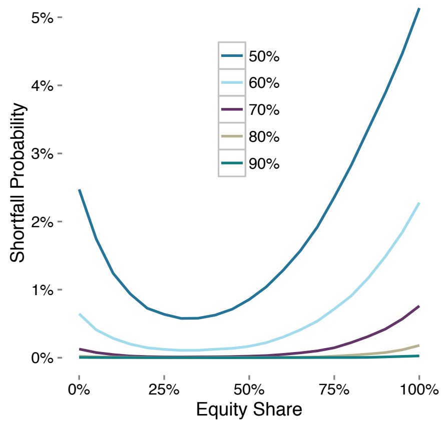

Appendix C: Alternative assumptions

In this section, we briefly review the findings in the main text when varying a number of assumptions. We consider the additions of income and spending – the progression from Case 1 to 3 – when assuming no predictability in equity and fixed-income returns, when the correlation between equity and fixed-income returns is zero, and using a lower oil volatility assumption of 15 percent.

i. ***No predictability***

We assume that equity and fixed-income returns are not predictable / expected returns do not vary over time. Figure 46 shows the change in the Case 3 distribution with and without predictability. Figure 47 shows how the addition of income and spending changes in Cases 1-3, with means preserved across the distributions. The impact of assuming predictability is relatively small; the lack of mean reversion generates a slightly wider distribution. The build from Case 1 to Case 3 remains the same as in the main text. Figure 48 shows how the probability of a 50 percent shortfall in fund value evolves over the three cases, where the progression remains the same.

Figure 46: Distribution of ten-year fund value, base vs no predictability

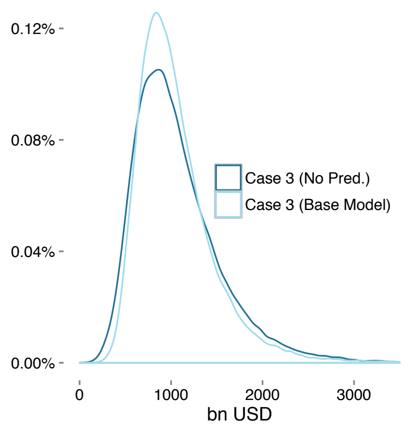

Figure 47: Distribution of ten-year fund value, Cases 1, 2 and 3, preserved means, no predictability

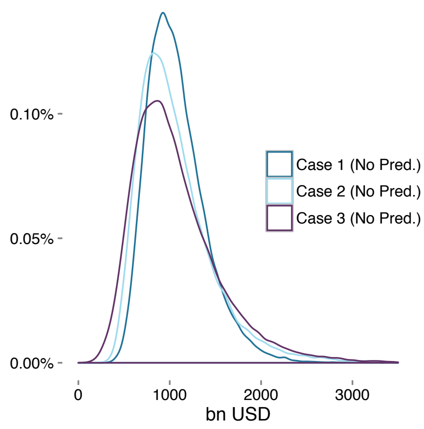

Figure 48: Probability of 50 percent fall in fund vs initial value at ten-year horizon, no predictability

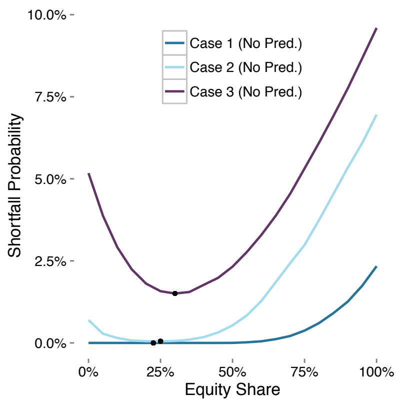

ii. *Correlation between equity and fixed-income returns*

We set the correlation between equity and fixed-income returns to +0.3. Figure 49 shows how the ten-year distribution changes given the change in correlation. The impact is relatively small. Figure 50 shows that the progression of mean-preserved distributions across the three cases remains the same under the zero correlation, and this is also the case in the shortfall curves in Figure 51.

Figure 49: Distribution of ten-year fund value, base vs zero correlation between equity and fixed income

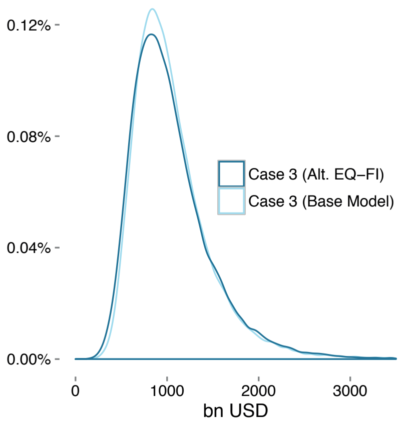

Figure 50: Distribution of ten-year fund value, Cases 1, 2 and 3, preserved means, zero correlation between equity and fixed income

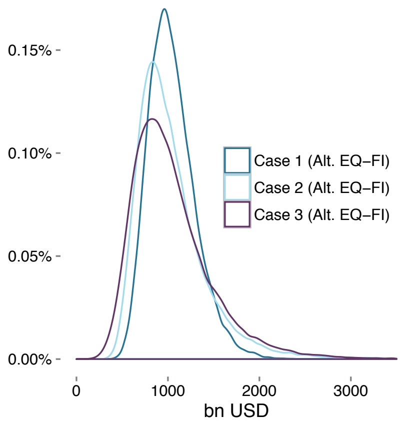

Figure 51: Probability of 50 percent fall in fund vs initial value at ten-year horizon, no predictability

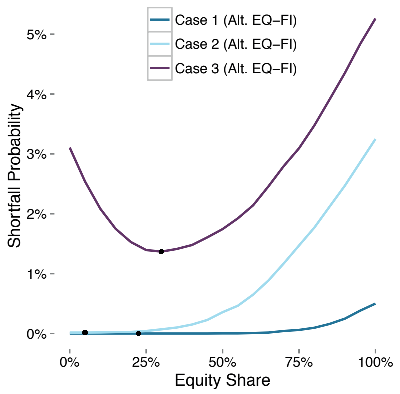

iii. *Lower oil volatility*

Here, we present our findings when assuming a lower volatility of oil price innovations – reduced from 30 percent to 15 percent. Figure 52 shows the effect on the ten-year distribution, which is slightly narrower. Figure 53 shows the build from Case 1 to Case 3. Figure 54 shows the shortfall probability curves under the lower volatility assumption.

Figure 52: Distribution of ten-year fund value, base vs lower oil volatility

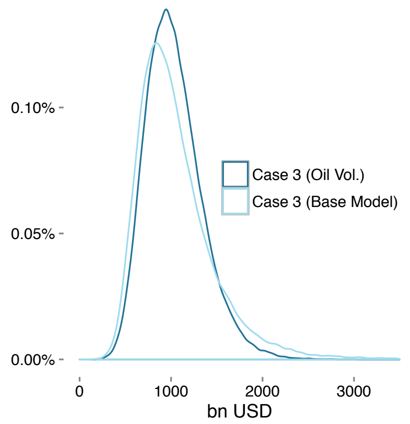

Figure 53: Distribution of ten-year fund value, cases 1, 2 and 3, preserved means, lower oil volatility

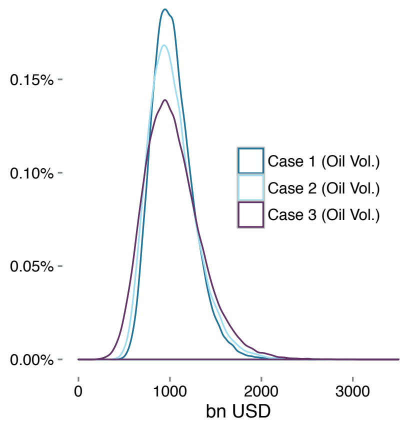

Figure 54: Probability of 50 percent fall in fund vs initial value at ten-year horizon, lower oil volatility