ISSN 1503-299X

WORKING PAPER SERIES

No. 1/2022

Capital preservation and current spending with Sovereign Wealth Funds and Endowment Funds: A simulation study*

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

Haakon Andreas Trønnes
Department of Economics
Norwegian University of Science and Technology

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

Department of Economics

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

We are indebted to Espen Henriksen and Snorre Lindset for useful comments and early inputs to the study. Henriksen helped us obtain the financial data. Special thanks go to Gunnar Bårdsen for constructive comments on the econometric part of our work and to Knut K. Aase for introducing us to the literature on fractional Brownian motions. The responsibility for errors and omissions is solely our own. We have benefitted from financial assistance from Finansmarkedsfondet, the financial markets research fund under The Research Council of Norway, Project no. 294398.

Capital preservation and current spending with Sovereign Wealth Funds and Endowment Funds: A simulation study

Knut Anton Mork*, Haakon Andreas Trønnes**, and Vegard Skonseng Bjerketvedt***

February 2022

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

**PhD student, Norwegian University of Science and Technology (NTNU), Department of Economics, haakon.a.tronnes@ntnu.no

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

JEL classification: C63, E62, G11

Key words: Sovereign wealth funds
Long-term risk
Fiscal smoothing
Depletion risk

We are indebted to Espen Henriksen and Snorre Lindset for useful comments and early inputs to the study. Henriksen helped us obtain the financial data. Special thanks go to Gunnar Bårdsen for constructive comments on the econometric part of our work and to Knut K. Aase for introducing us to the literature on fractional Brownian motions. The responsibility for errors and omissions is solely our own. We have benefitted from financial assistance from Finansmarkedsfondet, the financial markets research fund under The Research Council of Norway, Project no. 294398.

ABSTRACT

We study the dynamic performance of a sovereign wealth fund with the dual purpose of preserving wealth to future generations and provide regular contributions to annual budgets. We do not seek to derive optimal results but use stochastic simulations to study the implications of existing rules. Because our immediate focus is on the Norwegian Government Pension Fund Global, popularly referred to as the oil fund, we study rules established for this fund and calibrate our model from data for this fund’s financial returns as well as the Norwegian fiscal policy for which fund withdrawals make up an integral part. However, our results should be of interest to other sovereign wealth funds as well as endowment funds for universities and other non-profit institutions. Withdrawals are limited to the fund’s expected real returns as a first approximation; however, deviations from this rule are allowed to fund automatic stabilizers as well as discretionary fiscal policy. On top of that, the rule also allows smoothing so as to avoid abrupt changes. We find that this combination of rules implies considerable uncertainty regarding the fund’s future development. It does not preserve the fund’s value for future generations, not even in expectation, and will eventually lead to depletion. Keeping risk (equity share) low and withdrawal rates below the expected rate of return are possible remedies. However, for the depletion risk to be eliminated, stricter bounds must be put on countercyclical fiscal policy than the pattern followed in recent decades.

1. Introduction

Owners and managers of sovereign wealth funds (SWF) and endowment funds face a number of challenges and tradeoffs. On the one hand, they typically seek to preserve the value of the fund's assets for future use; on the other, they want to be able to draw on the fund's financial returns to defray the costs of current services and/or policy measures. It is usually desirable to shield the stream of these services from the vagaries of the financial markets, as argued, for example, by Barro (1979).

The financial returns are rarely the single source of revenues for the institution in questions. Universities with endowment funds receive government funding and/or collect tuition. Governments with SWFs collect taxes. These revenues may vary, typically with the business cycle. For a sovereign government, this kind of smoothing is essentially equivalent to letting the automatic stabilizers work. Governments typically want to go further and use discretionary fiscal policy to help stabilize the country's business cycle. These sometimes conflicting demands create challenges in terms of the preservation of the fund for future generations. For example, occasions may arise where current needs require a sell-off of assets in the wake of price declines. In extreme cases, policy makers may face a tradeoff between stabilizing the economy and depleting the fund.

With a complete representation of fund-owner preferences regarding these and other tradeoffs we should be able to derive the optimal strategy for investing the fund and spending of the proceeds. Any such presentation would necessarily be complex and might not even be well defined considering that the relevant decisions are made by groups whose members might have conflicting interests. Furthermore, rational decision making is made difficult by the complexities of the consequences of many owner actions. We believe this makes a case for studying the implications of observed fund-owner behavior. In this paper, we do this by simulating the likely consequences of the rules and practices that have been implemented for the Norwegian Government Pension Global (GPFG). Because many of these rules and practices are shared by other endowment funds and SWFs, we believe our results should be interesting for a wide range of cases.

The Act of 1990 (with subsequent amendments) governing the fund stipulates that the purpose of preserving the government's temporary revenue from oil and gas extraction. for future generations. A regulation passed by Parliament in 2001 furthermore stipulates the rules for spending of the fund's proceeds. They can be summarized as follows:

- As a main rule, an amount corresponding to the fund's expected financial return can be added to each annual budget. The expected rate of return was stipulated as 4 percent in 2001 and lowered to 3 percent in 2017.
- The rule applies to the structural budget balance, meaning that the rule does not apply to the funding of automatic fiscal stabilizers, whether positive or negative.
- Additional fiscal spending can be financed by the fund if desirable for stabilization of domestic business cycles, provided similar tightening is undertaken in good times.
- Additional smoothing can be applied to avoid sharp changes in annual withdrawals.

This paper simulates the likely implications of the practices based on this set of rules over time. We base the simulations on an econometric model of fiscal behavior and the relevant financial variables, estimated on data for 1991–2020. Our simulations reveal a number of challenges that arise from the observed behavior.

First, although the main rule of spending only the expected return is intended to maintain the fund's value for future generations, does not meet that intention. At best, it preserves the expected value, but with a variance that rises with the time horizon. However, preservation of the expectation does not make preservation of the current value likely. To the contrary, by concentrating more and more of the probability distribution next to its zero lower limit, it eventually depletes the fund with a probability approaching unity.

Second, because the fund is invested globally, but the proceeds spent locally in the national currency, currency conversion is likely to introduce an element of negative serial correlation in the rates of return in the national currency. Then, the fund's future value is not even preserved in expectation.

Third, the use of the fund for fiscal policy tends to change the shape of the probability distribution of future fund values from lognormal to a hybrid of lognormal and a truncated normal with a spike at zero, thus raising the risk of future depletion.

Fourth, smoothing after a major loss means that withdrawals as a percent of fund value will increase at the same time as fund value falls. That raises the question of whether the fund can be caught in a "death spiral" that make depletion unavoidable.

Our results indicate that at least the first three of these challenges are quantitatively important for the Norwegian GPFG. However, we also propose remedies. Reducing financial risk taking by lowering the equity share to 50 percent from the current 70 reduces the need for smoothing by smoothing the development of the fund itself. It also delays the tendency for the probability mass to be concentrated near zero but does not eliminate it. A better remedy for that problem would be to reduce the withdrawal rate by 1.5 to 2 percent lower than the expected return. Both remedies would result in lower withdrawals in early years; but this effect may be reversed over time, especially with a lower withdrawal rate, because the fund then is allowed to grow faster.

Reducing or eliminating the depletion risk associated with fiscal policy would require stricter limitation of the scope for fiscal stimulation during the recessions. Whether or not such limitations would be worth the cost of sharper recessions depends on policy makers' preferences.

Fiscal sustainability is naturally at least as important for governments whose financial operations are mainly on the liability side of the balance sheet. Our paper is thus related to the vast literature on public debt, with Blanchard's (2019 Presidential Lecture as a recent example. However, the issues faced by governments with negative financial wealth lie beyond the scope of this paper. We also do not go into the issues that arise for a government that draws down its financial wealth and starts borrowing.

Our paper adds to the literature on the tradeoffs involved in the management of persistent, yet temporary revenues from non-renewable natural resources, such as van der Ploeg and Venables (2011) and van den Bremer et al. (2016). However, unlike the former, we do not focus on developing economies. Unlike the latter, we ignore the prospect of future fund deposits and focus instead on the management of and spending of revenues from the fund proper. Our approach is similar to the one taken by Lindset and Madsen (2018) but goes deeper into the consequences of the strategy set up by a government like Norway's. Unlike Aase and Bjerksund (2021) and Lindset and Mork (2019), we do not seek to derive optimal results, but seek instead to uncover the implications, perhaps unintended, of decisions actually made.

We furthermore add to the more general literature on SWFs, including Baldwin (2012), Alhashel (2014), Bernstein et al. (2013), Paltrinieri and Pichler (2013), Dreassi et al. (2017), and Johan et al. (2013). Finally, our results should also add insights to the literature on other endowment funds, which have been studied, for example, by Barber and Wang (2013), Brown et al. (2014, and Dahiya and Yermack (2018).

The rest of the paper is organized as follows. Section 2 gives a conceptual introduction to the challenges of long-term investing. Section 3 presents the relevant data and our time-series model of fiscal behavior and financial returns. Section 4 presents the simulation model, including our specification of the Norwegian Fiscal Rule. Section 5 presents the main simulation results, and Section 6 the implications of smoothing. Section 7 presents and discusses the possible remedies for the challenges that our analysis uncover; and Section 8 concludes.

2. Some Inconvenient Truths

The Norwegian fiscal rule allows annual withdrawals corresponding to the expected real return on the fund. Similar rules are practiced by a number of endowment funds. Such rules are intended to guarantee preservation of the fund for future users and generations. The intuition is as simple as it is appealing: By spending only the return, the principal is preserved.

In the absence of risk, this intuition is correct. With risk, returns will fluctuate around its expected value, but the fluctuations should average to zero over time, as suggested by the law of large numbers. It would then seem to follow that the value of the fund should be preserved over time as well. That turns out to be true only in a very limited sense. That is, provided the rate of return is serially uncorrelated, the fund value will be preserved *in expectation*; however, the preponderance of probability is for the fund to gradually lose value over time. Furthermore, this preponderance rises with the time horizon. Eventually, the fund will be depleted with a probability approaching unity.

The main explanation for this result, which was proved by Dybvig and Qin (2021) and presented in the context of the Norwegian GPFG by Aase and Bjerksund (2021), is that the law of large numbers applies to sums, not products; and the dynamic development of an investment fund depends on products of the type

$$A_{t+1} = (1 + r_t)(A_t - D_t),$$

where A_t denotes the fund value, D_t withdrawals, and r_t the rate of return. Simply put, a 50 percent loss is not canceled out by a subsequent 50 percent gain. Although the effects of gains and losses will outweigh each other in expectation, the probability distribution of future fund values become increasingly skewed to the left the further one looks into the future. The increasing uncertainty is a case of a mean-preserving spread. Whereas spreads to the right (toward higher values) are unlimited, spreads to the left (towards lower values) are limited by the lower bound of zero, an absorbing barrier[1]. So, as the time horizon increases, more and more of the probability mass will agglomerate at values very close to zero. The expectation will be preserved by the possibility of extremely high values, but with low probabilities. In the limit, the entire probability mass collapses into a spike next to zero.

A simple proof for the case of normally distributed and serially uncorrelated returns with constant means and variances in continuous time, so that $r(t) \sim \mathcal{N}(\mu, \sigma^2)$, goes as follows. From Itô's lemma, the law of motion for the fund value is

$$d \ln A = \left(\mu - \frac{1}{2}\sigma^2 - \delta\right) dt + \sigma dB, \tag{2.1}$$

[1] We ignore the possibility of negative asset values, e.g. debt.

where δ denotes withdrawals as shares of the fund value, and B is a Brownian motion, so that

$$\ln\left(\frac{A(t)}{A_0}\right) = \left(\mu - \frac{1}{2}\sigma^2 - \delta\right) t + \sigma^2 \int_0^t dB. \tag{2.2}$$

Then, clearly,

$$\mathbb{E}_0 \ln\left(\frac{A(t)}{A_0}\right) = \left(\mu - \frac{1}{2}\sigma^2 - \delta\right) t \tag{2.3}$$

and

$$\mathbb{V}_0 \ln\left(\frac{A(t)}{A_0}\right) = \sigma^2 \mathbb{V}_0 \left(\int_0^t dB\right) = t\sigma^2, \tag{2.4}$$

where the last equality follows from the fact that B is a Brownian motion.

So, because $A(t)/A_0$ is lognormally distributed,

$$\mathbb{E}_0 \left(\frac{A(t)}{A_0}\right) = e^{(\mu-\delta)t}, \tag{2.5}$$

which is unity provided $\delta = \mu$.

Thus, withdrawing an amount corresponding to the expected return does indeed preserve the fund in expectation. However, the probability of the future fund value falling short of an arbitrary positive number a is

$$P(A(t) - a) = \frac{1}{2}\left[1 + \mathrm{erf}\,(\frac{\ln a - \left(\mu - \frac{1}{2}\sigma^2 - \delta\right)t}{\sigma\sqrt{2t}}\right], \tag{2.6}$$

where erf denotes the Gauss error function

$$\mathrm{erf}(z) = \frac{2}{\sqrt{\pi}}\int_0^z e^{-t^2}dt.$$

Clearly, for $t \to \infty$,

$$P(A(t) - a) \to \begin{cases} 1\ if\ \delta > \mu - \frac{1}{2}\sigma^2 \\ \frac{1}{2}\ if\ \delta = \mu - \frac{1}{2}\sigma^2. \\ 0\ if\ \delta < \mu - \frac{1}{2}\sigma^2 \end{cases} \tag{2.7}$$

If withdrawals equal expected returns, so that $\delta = \mu > \mu - \frac{1}{2}\sigma^2$, the value of the fund will ultimately fall below any positive number with a probability approaching unity (i.e. almost surely). Figure 1 shows the probability density functions for this distribution with

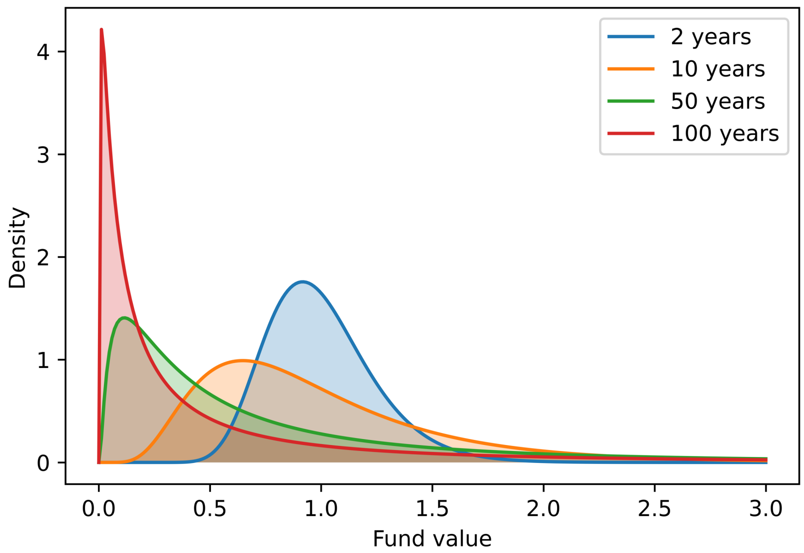

Figure 1: Probability distribution of fund value in t years starting from 1 in year 0.

Log-normal distribution, continuous time
Withdrawal rate equal to expected return
Standard deviation of rate of return: $\sigma = 0.17$

increasing time horizons t. The graph illustrates how the density becomes more and more skewed until most of its mass is concentrated near zero. On the other hand, (2.7) also shows that sustainability can be obtained, by withdrawing less than the expected return[2], i.e. $\delta \leq \mu - \frac{1}{2}\sigma^2$.

In practice, withdrawals are typically not made or decided as a continuous stream, but as an annual value, corresponding to some fraction of the value of the fund at the beginning of the year. That does not change the argument, however. Of course, compound

[2] Lindset and Matsen (2018) state as a condition for sustainability of the fund's value is a "spending polity in which the expected cumulative spending rate is less than or equal to the expected cumulative log-returns." Although their analysis appears to be correct, we believe this wording may give the unintended impression that spending the expected return is sufficient for value preservation.

interest will make the expected rate of annual return greater than the expected instantaneous return, in our example an expected annual return of $\bar{r} = e^{\mu + \frac{1}{2}\sigma^2} - 1 > \mu$. One might think that this fact would solve the problem just discussed. However, time-series estimates of annual returns would then be estimates of $\bar{r}$, not μ. Equating annual withdrawals to the expected annual return would then set withdrawals at the ratio $\bar{r}$ of the fund's value as well, whereas value preservation would require annual withdrawal rates of at most $\delta = e^{\mu} - 1 < \bar{r}$.

If the rates of return are negatively serially correlated, withdrawals equaling the expected return will not even preserve the fund's value in expectation. Mean reversion of stock prices, as claimed by Fama and French (1988) and analyzed by many others, is known to be advantageous to investors as it implies lower long-term uncertainty than with white-noise returns. What seems less well known is that negative serial correlation for returns also reduces the fund's expected future value for a given withdrawal rule.

For this result again, the explanation is to be found in compound interest. To see this, consider the simple case in discrete time where withdrawals are made at the rate $\delta = \bar{r}$, the expected rate of return. After two years, the fund will develop according to

$$\frac{A_2}{A_0} = (1 + r_1 - \bar{r})(1 + r_2 - \bar{r}),$$

so that

$$\mathbb{E}_0\left(\frac{A_2}{A_0}\right) = 1 + \mathrm{cov}_0(r_1, r_2) < 1$$

if the rates are negatively serially correlated.

Although we do not assume mean reversion for stock returns, negative serial correlation is introduced via another route, which is relevant for a fund that is invested globally, but with annual draws determined by the rate of return in local currency. This is the case for the Norwegian GPFG, which by law is restricted from investing in Norwegian companies or securities denominated in Norwegian kroner, but for which the rule for annual withdrawals is expressed as a fixed percentage of the fund's value in kroner. If r_t^{USD} is the fund's real rate of return in U.S. dollars and v_t the log real exchange rate, the rate of return in Norwegian kroner, r_t, is then

$$r_t = r_t^{USD} + v_t - v_{t-1}.$$

Clearly, if the real exchange rate is stationary, the conversion of the return into Norwegian kroner introduces an element of negative serial correlation.

The question of stationarity versus a unit root for real exchange rate has not been settled in the literature. The well-known difficulty of reliably forecasting future exchange rates, as noted by Meese and Rogoff (1983) and further analyzed by many others[3], supports a unit root, whereas the hypothesis of long-run purchasing power parity supports stationarity. In our data, a unit root can be rejected with borderline significance, as explained in the next section. On an extended sample, going back to the end of the Bretton Woods system in 1971, the rejection becomes barely significant at the 10 percent level. In our simulation exercises, we find that a unit-root assumption results in projections of future real exchange rates that in some cases are too high and in other cases too low to be

[3] Including Hungnes (2020) for the Norwegian krone

credible. When added to the long-sample econometric data and the theoretical purchasing-power argument, this convinces us to specify the real exchange rate as stationary, at least as our main case.

With log-normal, but serially correlated returns in continuous time, equations (2.1) and (2.2) continue to hold except that the Brownian motion B is replaced[4] by the fractional Brownian motion B_H with the Hurst coefficient H. A fractional Brownian motion is defined such that $H = 1/2$ indicates no serial correlation, $H > 1/2$ positive serial correlation, and $H < 1/2$ negative serial correlation, which is the case we consider. Then, the expected future fund value continues to be given by (2.3), whereas the variance in (2.4) is replaced by

$$\mathbb{V}_0 \ln\left(\frac{A(t)}{A_0}\right) = \sigma^2 \mathbb{V}_0\left(\int_0^t dB_H\right) = t^{2H}\sigma^2 < t\sigma^2 \text{ for } H < \frac{1}{2}. \tag{2.8}$$

Again using the formula for the expectation of a lognormal variable, we then obtain

$$\mathbb{E}_0\left(\frac{A(t)}{A_0}\right) = e^{\left(\mu - \frac{1}{2}\sigma^2 - \delta\right)t + \frac{1}{2}t^{2H}\sigma^2} < e^{(\mu-\delta)t} = 1 \tag{2.9}$$

for $\delta = \mu$. So, keeping withdrawals equal to the fund's expected return does not even preserve the fund value in expectation. This effect adds to the unsustainability of the popular rule of equating annual withdrawals to expected returns. Compensating for it requires cutting the withdrawal rate even lower. Figure 2 shows the probability distributions at varying time horizons for a fund with lognormally distributed gross returns with negative

[4] We are indebted to Knut K. Aase for introducing us the analyses of fractional Brownian motions by Yaglom (1958), Mandelbrot and van Ness (1968), and Biagni, Hu, Øksendal, and Zhang (2008).

serial correlation. As the graph suggests, the variance in this case even shrinks with the time horizon, partly because of the mean reversion implied by the negative serial correlation and partly because the variance of a lognormal distribution is proportional to the median, which declines as mass is moved towards zero[5].

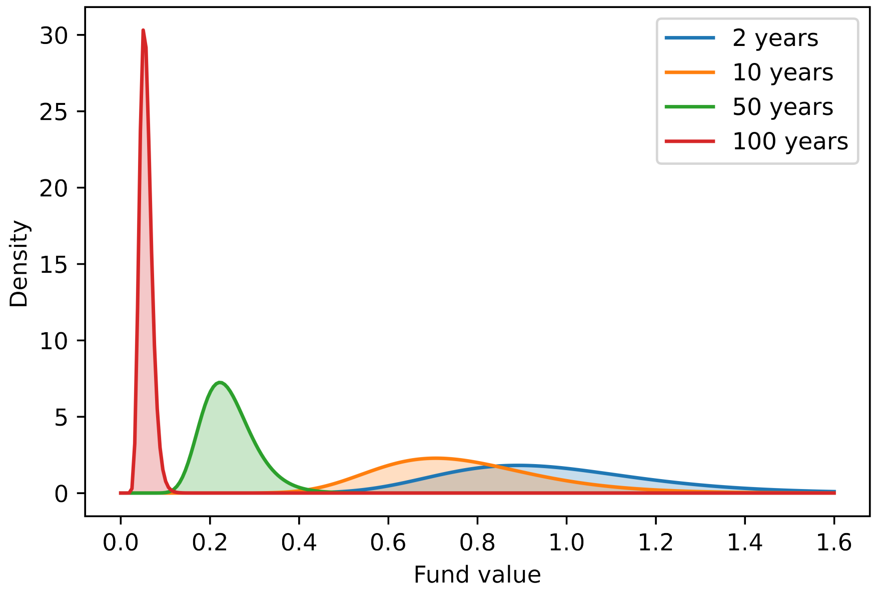

Figure 2: Probability distribution of fund value in t years.

Log-normal distribution, continuous time
Withdrawal rate equal to expected return
Standard deviation of rate of return: $\sigma = 0.17$
Autocorrelation: -50%

Smoothing adds yet another layer of complications. Endowment fund owners may, for example, permit larger draws at times where other sources of revenues, such as tuition or ticket sales, are lower than normal. SWFs may be used to maintain public services

[5] The variance of the fund value t periods hence is, in this case

$$\mathbb{V}_0\left(\frac{A(t)}{A_0}\right) = \left(e^{t^{2H}\sigma^2} - 1\right)e^{-\sigma^2 t\left(1-t^{2H-1}\right)}.$$

whenever tax revenues fall because of recessions, as codified in the Norwegian Fiscal Rule. Discretionary fiscal action to stabilize the economy may be added. A minimum requirement for such extraordinary draws to be sustainable would be that they are reversed in good times, statistically speaking that they are distributed symmetrically around zero. As an illustration of their effects, Figure 3 adds a normally distributed shock to the process presented in Figure 1.

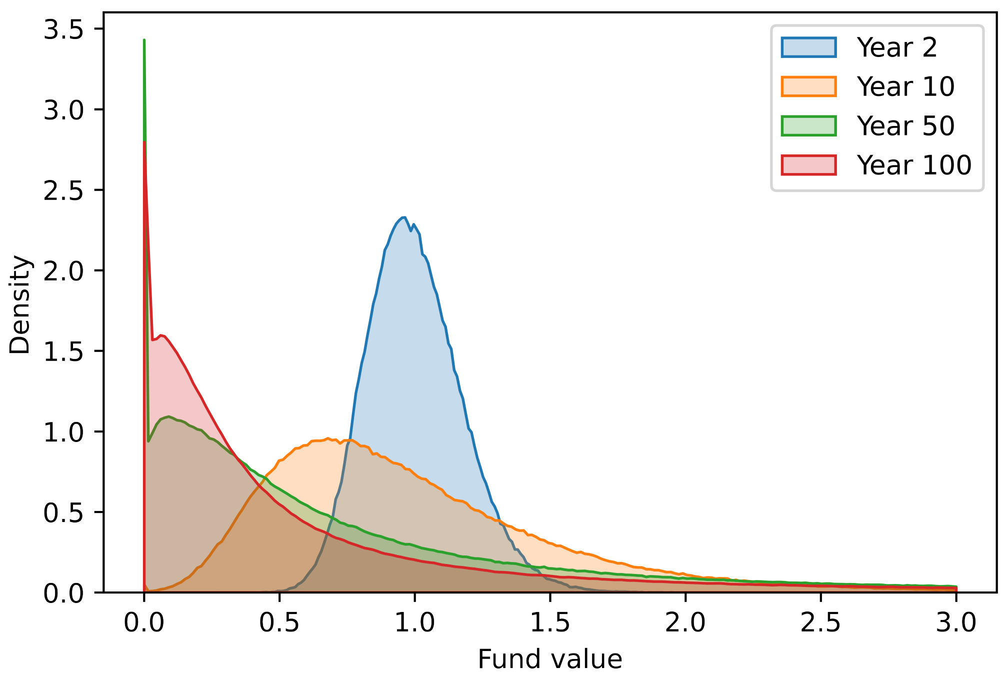

Figure 3: Probability distribution of fund value in t years.

Log-normal distribution, continuous time
Withdrawal rate equal to expected return
Standard deviation of rate of return: $\sigma = 0.17$
Including normally distributed shocks
Figure based on simulations

Even though these shocks are symmetric around zero, negative shocks have a much more dramatic effect because of the zero lower bound for the fund itself. The distribution for the future fund value becomes a hybrid of a lognormal and a truncated normal with a spike at zero. The height of the spike obviously depends on the variance of the normally distributed element; however, to eliminate it completely would require replacing the normal distribution with a support with finite bounds.

Many real-life spending rules usually allow for more explicit smoothing because the vagary of the market makes the fund value change over time, sometimes abruptly. Withdrawals made at a fixed ratio of the fund's value would then follow the same fluctuations, which would make smoothing attractive.

In discrete time, a natural form for smoothing would take the linear form

$$D_t = \lambda D_{t-1} + (1-\lambda)\delta A_t, 0 < \lambda < 1. \quad (2.10)$$

This formulation does not add to the depletion risk, provided it is applied symmetrically for positive and negative changes in the fund value. However, after a financial disappointment, a larger share than normally will then be drawn form an already low fund value. That raises the question of whether such situations could lead the fund into a "death spiral" that won't stop until the fund is depleted. The answer is, perhaps not surprisingly, that this could happen if the drop in the fund value is large enough, relative to the speed back to normal $(1-\lambda)$ implied by the smoothing rule. Furthermore, in the absence of bounds for the financial returns, that can indeed happen and will in fact happen eventually with probability one.

Dybvig and Qin propose a modification to the linear smoothing rule above, whose main purpose is to make sure that the speed back to normal is higher the further away from target the spending rule has gone. As the linearity of the above formula is mainly chosen out of convenience, the addition of a non-linear term can hardly be considered a fundamental blow to the desirability of smoothing *per se*. However, we suspect that actual decision makers would be tempted to make the speed back to normal slower rather than faster in cases of large financial losses. The rational solution to this dilemma might perhaps be to move slowly at first, but then significantly faster once agents have had time to prepare for changes.

However, after catastrophic losses, other concerns may sensibly take priority over fund value preservation. For this reason, we will not attempt to model reactions to such events explicitly. Instead, we will use simulations as illustrations of the possible problems that may result from linear smoothing.

3. Data and Time-Series Model

Simulation of fund behavior over time requires model specification of the time series behavior of financial returns as well as the fiscal variables that determine withdrawals. We estimate this behavior from historical data. Because the fiscal data are available on the annual frequency only, we use annual data for all our series.

For the financial data, we start by noting that the fund's equity and fixed-income decisions are required to closely follow the FTSE Global equity index and the Bloomberg bond index, respectively, data for which our colleague Espen Henriksen graciously shared with us. Unfortunately, these index series are only available from 1991 on, which limits our sample to the 1991 – 2020 period. We thus base our moment estimates for the fund distribution on data for the total-return development of these indices. For each series, we compute annual nominal returns in U.S. Dollars (USD) as the log difference between the end-of-year values for the respective index values. We then deflate them by the December-to-December log changes in U.S. CPI (downloaded from the FRED database) to obtain real USD returns. However, because what matters to the Norwegian government is real returns in Norwegian kroner (NOK), we also need data for the real exchange rate. In analogy with the asset-return data, we use year-end daily bilateral nominal exchange rates, obtained from the Norges Bank, and deflate them by the December values of the relative U.S.-Norwegian CPI, the latter downloaded from Statistics Norway.

Data for the relevant fiscal data were available in the Norwegian government accounts (*Statsrekneskapen*). By law, all non-oil deficits[6] on the government budget are required to be covered by fund withdrawals for as long as there is money in the fund. The time series behavior of the withdrawals is thus identical to that of the non-oil deficit.

[6] The Norwegian government accounts distinguish between total and non-oil deficits. By law, the government's net revenues from oil and gas extraction are to be deposited directly into the fund. The difference between all other expenditures and revenues are referred to as the non-oil deficit.

This deficit is furthermore decomposed into a structural and a cyclical component according to a set of algorithms developed by the Ministry of Finance. The cyclical component is supposed to represent the contribution—positive or negative—that is purely due to the cyclical movements in the overall economy and thus not caused by changes in tax rates or entitlement rules. Typically, recessions reduce revenues because incomes fall and increase expenditures on unemployment benefits, and vice versa. These movements are also referred to as automatic stabilizers. We denote this magnitude as X_t.

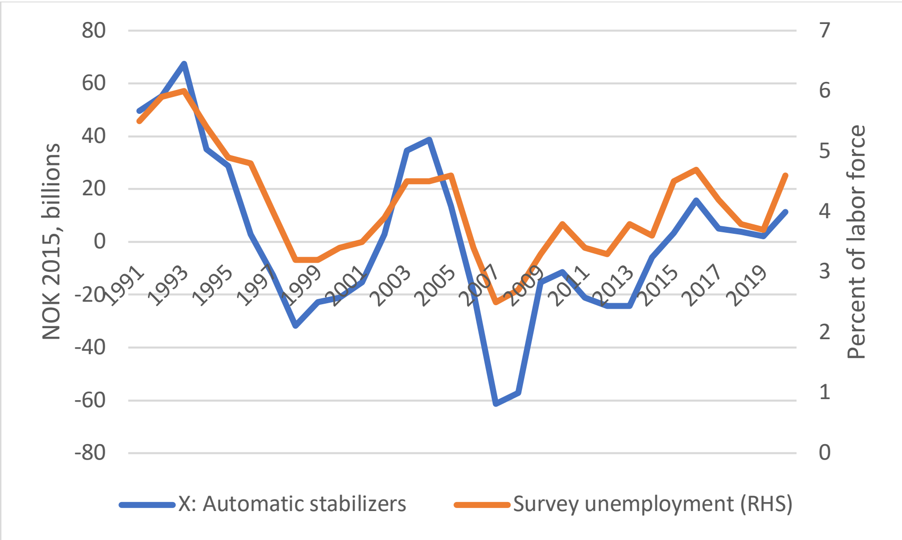

Figure 4: Automatic stabilizers and unemployment

Data defined as in the text.
Sources: Government accounts and Statistics Norway

Data for this variable, which may take on negative as well as positive values, can thus be taken directly from the government accounts[7]. As shown in Figure 4, the movements of this variable closely parallels those of the ILO-compatible survey unemployment. It can thus be interpreted as a fair (though inverse) representation of the Norwegian business cycle.

[7] In additional to this cyclical component, the difference between the structural and actual non-oil deficit includes three other items: (i) transfers from Norges Bank, deviations from trend, (ii) extraordinary items, and (iii) changes in accounting standards. We ignore these components because, in our analysis, they mainly represent noise.

The automatic stabilizers can be interpreted as the cost of maintaining a smooth flow of government services in the face of cyclical variations in tax revenues and expenditures such as unemployment compensation. Foundations face similar challenges, wanting to maintain services and other activities, like museum exhibitions or university teaching and research at times when revenues from sources other than the endowment fund temporarily fail to arrive in their usual amounts. Recessions may, for example, hurt ticket revenues, charitable giving, student enrollment, or government funding. Thus, the presence of our variable X does not limit the validity of our analysis to governments.

Fiscal spending or tax cuts, aimed at stabilizing the overall economy, is specific to governments, however. The specification of the rules governing Norwegian fiscal policy requires some care. The Fiscal Rule, introduced in 2001, allows the government to normally run a structural, non-oil deficit corresponding to the expected return on the fund, defined as the fund's value at the beginning of the year multiplied by the expected rate of return. Letting $\bar{r}$ denote this rate and A_t the beginning-of-year value of the fund, this normal, structural, non-oil deficit is defined as

$$\bar{S}_t = \bar{r}A_t. \tag{3.1}$$

However, the *Storting* (Parliament) may exceed this limit if the cyclical situation calls for discretionary fiscal stimulus, provided this excess is compensated by subsequent tightening as the economy normalizes. Letting F_t denote the discretionary spending (or tax cut), the realized structural, non-oil deficit can then be written as the sum

$$S_t = \bar{S}_t + F_t, \tag{3.2}$$

where F_t can have either sign.

The government accounts contain data for S_t but not F_t. Instead, we construct them by using the above formula for $\bar{S}_t$ and subtract the resulting series from S_t. For 1991 – 2001, before the Fiscal Rule went into effect, we stipulate $\bar{S}_t = 0$, so that $F_t = S_t$. For the subsequent years, we use $\bar{S}_t = \bar{r}A_t$. However, the official estimate for $\bar{r}$, which was estimated as 4 percent for the 2001 decision, was cut discretely to 3 percent in 2017 for budgets from 2018 on. However, this cut came after a lengthy public debate starting with a speech[8] by the Central Bank Governor in 2012, where he argued that the persistent decline in global interest rates called for a lower estimate. Although the Governor's proposal at first was rejected by elected officials, the debate continued. During this debate, structural, non-oil deficits were persistently kept well below the bar set by the 4 percent criterion, by increasing margins, which we interpret as a gradual lowering of the effective criterion until the official confirmation of the new 3 percent. In our calculations of $\bar{S}_t$, we have thus let $\bar{r}$ decline linearly by 0.2 percentage points starting in 2014 until it reached 3 percent in 2018.

Figure 3 displays the movements in the actual structural deficits and the permitted levels according to our smoothed series as well as the raw series implied by the abrupt drop in $\bar{r}$ from 4 to 3 percent in 2018. The upward trend in all series reflects the persistent growth of

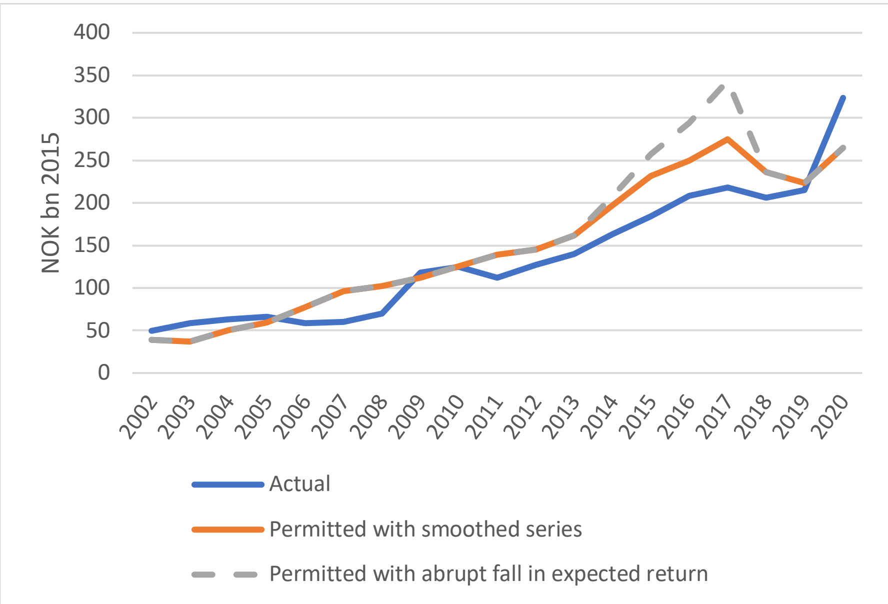

Figure 5: Actual and permitted non-oil, structural deficits

Variables defined as in the text.

[8] https://www.norges-bank.no/en/news-events/news-publications/Speeches/2012/16022012-Annual-address/, accessed Sep 21, 2021.

the fund as new oil and gas revenues have been deposited. The graph furthermore shows how actual structural deficits increasingly fell short of the upper limit between 2013 and 2018. In fact, they also stayed below our smoothed series. The drop in the smoothed series from 2017 to 2018 is due to the negative real return on the fund that year and is thus not a result of smoothing. The 2020 performance is naturally dominated by the Covid-19 pandemic.

In simulating fund withdrawals, we need the fiscal variables in natural units, i.e. NOK trillions at fixed prices. However, for modeling their time-series behavior, we note that the absolute magnitudes of fiscal actions tend to follow the same trend as real GDP. In Norway's case, we follow the common practice of using real mainland GDP, i.e. GDP except for all sectors except oil and gas extraction[9]. With this motivation, we specify our time-series model for the fiscal variables in units of percent of trend mainland GDP, with the latter estimated from the same sample as the model, i.e. 1991 – 2020.

The variables included in our model are thus the real USD-denominated equity return, denoted r_Q, the corresponding period return on bond investments, r_B, the log real exchange rate, v_t, real automatic stabilizers as a percent of trend mainland real GDP, x_t, and discretionary fiscal spending (and/or tax cuts), also as a percent of trend mainland GDP, f_t.

The movements in the Norwegian business cycle are contained in the residuals of the equations for the two fiscal variables. They thus summarize the combined effects of the various shocks driving this cycle. Because our main contribution lies in the simulation of the possible future movements of the fiscal variables, we find this parsimonious representation

[9] Norway's mainland GDP also includes the value added of offshore shipping, but this component is too small to matter for our purposes.

of the business cycle preferable to a more comprehensive specification of the underlying shocks and transmission mechanisms behind these movements.

As a preliminary exercise, we use univariate Dickey-Fuller tests to examine the variables' stationary properties. The results are presented in Table 1. For the two rates of return, we easily reject unit roots. For the exchange rate, a unit root cannot be rejected on the 10 percent level in our sample. However, as noted in Section 2, we obtain significance on this level when we extend the sample back to the end of the Bretton Woods system in 1971. Stationary is furthermore implied by the hypothesis of long-run purchasing power hypothesis. Because specifying the exchange rate as stationary also helps us avoid unreasonably high or low exchange rate projections in the simulations, we maintain stationarity as our main hypothesis.

Variable	$\hat{\boldsymbol{\rho}}$	$(\hat{\boldsymbol{\rho}}-\mathbf{1})/\boldsymbol{s.e.}(\hat{\boldsymbol{\rho}})$
$\boldsymbol{r_Q}$	-0.0943	-5.7011
$\boldsymbol{r_B}$	-0.0897	-5.8739
$\boldsymbol{v}$	0.6980	-2.2487
$\boldsymbol{x}$	0.7887	-2.2164
$\boldsymbol{f}$	0.8689	-2.0542

Table 1: Dickey-Fuller test results

Based on OLS equations including constant terms, but no time trends
Variables defined as in the text

For the detrended fiscal variables, we are unable to reject unit roots on statistical grounds. However, because non-stationarity would be economically unsustainable, we treat them as stationary as well. As in the case of the exchange rate, the stationarity assumption saves our simulations from blowing up because of unreasonably large deviations from the historical experience.

Informal inspection of the data suggests that the Norwegian business cycle is somewhat correlated with the global financial markets. This correlation appears to be especially strong during periods of crisis; however, our sample is too short to allow modeling of time-changing correlations. As a substitute, we look for possible asymmetries by distinguishing between the effects of positive and negative rates of return.

We thus specify a VAR model for our five variables with the modification that the coefficients of lagged rates of return are allowed to be asymmetric. We limit the lag length to two years on *a priori* grounds. For the equity and bond rates of return, r_Q and r_B, respectively, we are unsurprised to find no significant coefficients for their own lags or for those of any other variables.

The exchange rate shows signs of positive dependence on the two-year lags of stock returns. The implication would be that changes in the exchange rate, apart from the slow movement towards the long-term equilibrium discussed above, could be forecast two years in advance, a contradiction of the well-known difficulty of outperforming random-walk type forecasts. Now, a two-year lag effect seems unreasonable for any asset price. Furthermore, the coefficient that we find is unstable across subsamples. When tested on rolling samples going back to the end of the Bretton Woods system in 1971, with the S&P 500 as a proxy for our equity-return variable, it seems to have arisen mainly during the bull market and subsequent dollar strengthening of the late 1990 and to have faded somewhat during the 2010s. When added to our *a priori* reasons, this instability convinced us to treat this in-sample correlation as spurious and thus restrict all the lagged variables for the exchange rate to zero, except only for its own lag as implied by the stationarity.

For the fiscal variables, we hold no such prior notions. Thus, for the lag coefficients of these equations, we eliminate insignificant coefficients one by one, starting with the ones with the highest p-values, and continuing until all p-values are 0.1 or less. We then use bootstrapping to estimate 90 percent confidence intervals for the contemporaneous correlations of the residuals, constrain to zero those whose intervals contained zero, and reestimate the model as a system of seemingly unrelated equations with the resulting residual covariance matrix. The estimation results are presented in Table 2.

Here, we note that the automatic stabilizer variable, x_t, which most closely (though inversely) parallels the Norwegian business cycle, rises whenever the global stock market declines. We interpret this effect as driven by the international transmission of the global business cycle, with the lag possibly reflecting the fact that the stock market tends to move before the real economy. Such impulses as well as domestic shocks then set off highly persistent movements in the left-hand variable. Perhaps surprisingly, we find no effect on this variable of lagged changes in discretionary fiscal policy, f_t, nor are the two significantly correlated contemporaneously.

The discretionary fiscal variable f_t displays a more complex dependence on the lagged values of stock and bond returns. Although the structural meaning of each individual coefficient may not be immediately clear, the main point seems to be a somewhat similar sensitivity to global stock market movements as the automatic stabilizers. Persistence is high for this variable as well. The positive dependence on lagged movements in the automatic stabilizers, x_t can be interpreted simply as fiscal action to stabilize the business cycle.

We note that discretionary fiscal spending tends to increase when the krone declines and vice versa. Because Norway does not have a specific exchange-rate target (at least not after 1999), we believe this coefficient reflects reactions to movements in the price of oil, as an oil price decline tends to weaken the krona while also giving rise to policy measures to counteract the negative effects on the domestic oil sector.

The negative contemporaneous correlation between the real bond return r_B and the real exchange v_t appears to reflect the weakening effect on the dollar of a softening in the global business cycle (recall that high period bond returns tend to reflect falling yields to maturity). The positive contemporaneous correlations between the exchange rate and both fiscal variables seem instead to reflect the likely weakening of the krone in response to weakening of the Norwegian business cycle.

	r_{Qt}	r_{Bt}	v_t	x_t	f_t
Constant	0.0600 (0.0328)	0.0286 (0.0106)	0.8901 (0.2891)		-0.0594 (0.0253)
$r^+_{Q,t-1}$					0.0384 (0.0146)
$r^-_{Q,t-1}$				-0.0309 (0.0060)	-0.0466 (0.0061)
$r^-_{B,t-1}$					0.0962 (0.0267)
$r^-_{B,t-2}$				0.0828 (0.0214)	0.0547 (0.0192)
v_{t-1}			0.6883 (0.1036)		0.0195 (0.0089)
x_{t-1}				1.1278 (0.1142)	
x_{t-2}				-0.3716 (0.1172)	0.1528 (0.0705)
f_{t-1}					0.7760 (0.0471)
SEE	0.1766	0.0570	0.1054	0.0049	0.0045

Residual correlation matrix:

	r_Q	r_B	v	x	f
r_Q	1	0	0	0	0
r_B		1	-0.4395	0	0
v			1	0.3566	0.4750
x				1	0
f					1

Table 2: Estimation results, annual data 1991 – 2020

Variables and estimation method defined as in the text
Heteroscedasticity-robust standard errors in parentheses

4. Simulation Model

Our simulations are based on the model of the preceding section. For the vector of innovation terms, $\varepsilon_t = \left(\varepsilon_{Qt}, \varepsilon_{Bt}\varepsilon_{vt}, \varepsilon_{xt}, \varepsilon_{ft}\right)'$

we assume

$$\varepsilon_t \sim \mathcal{N}(0, \Sigma),$$

where the elements of Σ are as estimated in Table 2. From this distribution, we make one million different pseudo-random draws, each for 40-year or 80-year periods. To avoid sampling errors to cause differences between the scenarios we use the same draws for all the cases we simulate. Spot checks indicate that no significant distortion has been introduced this way.

The normality assumption means that we ignore third and fourth order moments. This means, in particular, that we ignore possible fat tails in the distribution of the financial returns. We do this out of convenience; however, we note that ignoring fat tails tends to bias our results against finding extreme results[10]. Readers may find our results extreme enough.

Because the fiscal variables in our estimation model are defined as ratios relative to trend mainland GDP, we project future values of the same variables in NOK 2021 trillion by multiplication by the projected future mainland GDP, using the Government's assumption of 1.4 percent growth per year[11] from the actual 2020 trend level. We also consider alternative growth rate assumptions as robustness checks.

[10] Bjerketvedt, Henriksen, and Lindset (in progress) seek to match the empirical third and fourth moments of these variables.

[11] https://www.regjeringen.no/en/dokumenter/meld.-st.-14-20202021/id2834218/, accessed on Sep 22, 2021.

Because the Norwegian fiscal rule is based on the fund value in Norwegian kroner, we need to simulate the future development of the fund in the same unit. The annual log real return in real Norwegian kroner becomes, as indicated in Section 2:

$$r_t = \omega r_{Qt} + (1-\omega) r_{Bt} + v_t - v_{t-1},$$

where ω is the equity share, initially fixed at $\omega = 0.7$.

The fiscal rule requires the annual withdrawals to equal the expected annual return, which we denote $\bar{r}$, currently officially estimated at 3 percent. We do not necessarily endorse this estimate, which is lower than implied by our estimates in Table 2. However, empirical estimates of first-order moments of financial variables are unlikely to be precise in general, as noted by Maenhout (2004), Pastor and Stambaugh (2012), and Sargent and Stachurski (2021). For this reason, and to analyze the internal consistency of the fiscal rule, we accept the official estimate for the purpose of our simulation exercises.

Letting r_{Q0} and r_{B0} denote the estimated expectations of r_{Qt} and r_{Bt} in table 2, respectively, we modify the rates of adjustment for equity and bond returns, respectively, while maintaining the estimated equity premium, by choosing the number z such that

$$\omega\left(r_{Q0} - z\right) + (1-\omega)(r_{B0} - z) + \frac{1}{2}\sigma_r^2 = \ln(1+\bar{r}), \tag{4.1}$$

where σ_r^2 is the variance of r_t implied by the estimated parameters, and $\bar{r} = 0.03$. Then, the annual gross return,

$$R_t = e^{r_t}$$

has the expectation

$$\mathbb{E}R_t = \mathbb{E}e^{r_t} = e^{\mathbb{E}r_t + \frac{1}{2}\sigma_r^2} = e^{\ln(1+\bar{r}) - \frac{1}{2}\sigma_r^2 + \frac{1}{2}\sigma_r^2} = 1 + \bar{r}.$$

However, implementing this modification in the equations for the two fiscal variables would modify their unconditional expectations as well, leading to simulated variables unreasonably far removed from the historical experience. To avoid that problem, we omit the modification in (4.1) when simulating the values of the equations for the two fiscal variables.

For a given draw of r_t, the fund develops according to

$$A_{t+1} = \max\{A_t e^{r_t} - D_t, 0\}, \tag{4.2}$$

where the “max” specifies zero as an absorbing barrier. That is, if the withdrawal D_t for any simulated scenario in any particular period t exceeds the fund’s value before the draw is made, we constrain the fund’s value to zero for tall subsequent periods in that scenario. This assumption means that we do not consider the possibility of borrowing.

Because our model has lags of up to two years, we assign values to the lagged variables equal to their respective expectations. For our first actual simulation year, we fix the initial fund value at 12 (2021 NOK trillion), roughly its magnitude at the time the research was done.

4.1. Fiscal Rule Implementation

The specification of the withdrawal D_t is key to our analysis. Because, in the Norwegian case, the law requires all non-oil government deficits to be drawn from the fund (as long as there is money in the fund), each annual draw equals the sum of the structural non-oil deficit and the automatic stabilizers:

$$D_t = S_t + X_t.$$

As explained in the preceding section, the structural, non-oil deficit can be separated into the structural, non-oil deficit permitted under the Fiscal Rule and discretionary fiscal spending that may make the actual structural, non-oil deficit deviate from that rule:

$$S_t = \bar{S}_t + F_t.$$

Thus, the annal withdrawal can be decomposed into three parts:

$$D_t = \bar{S}_t + F_t + X_t. \tag{4.3}$$

In the simulations, we compute $\bar{S}_t$ as $\bar{r}A_t$. The draws of the model residuals, the definition of r_t, and the lag coefficients then yield values for x_t and f_t, which in turn are multiplied by the projected future mainland GDP to obtain values for X_t and F_t, respectively.

Strictly speaking, the assumption of zero as an absorbing barrier means that a country that depletes its sovereign wealth fund must forsake discretionary fiscal policy as well as the automatic stabilizers. Instead, such a country would more likely join the crowd of countries that finance deficits by issuing bonds. However, we the modeling of such regime shifts lie beyond the scope of this paper.

4.2. Smoothing

Fund owners typically prefer smooth withdrawals to choppy ones. Fiscal rules like the Norwegian one aim to facilitate such smoothness by allowing withdrawals to be made according to expected rather than actual returns. Even so, the Fiscal Rule can produce large, abrupt changes in the structural, non-oil deficit in cases of substantial changes in the fund's value. Similarly large changes could result if discretionary fiscal spending in one year are cut to zero the following year so as to bring the structural, non-oil deficit back to its stipulated value $\bar{S}_t$.

Such abruptness can be mitigated by one or both of the fiscal variables. However, the Norwegian fiscal rule allows for even more flexibility by permitting—indeed mandating—explicit smoothing of the structural, non-oil deficit[12]. We model this smoothing as a partial adjustment:

$$S_t = \lambda S_{t-1} + (1-\lambda)\bar{S}_t + F_t, 0 \leq \lambda \leq 1. \tag{4.4}$$

The last term in equation (4.4) implies that discretionary fiscal spending in period t comes on top of the structural, non-oil deficit implied by the smoothing rule. We consider this formulation realistic because we believe that decisions on this item will be made based on stabilization needs as perceived in period t and are thus not part of a return to normal. The continuation of stimulation measures into 2022 because of the advent of the Omicron strain of Covid-19 illustrates this point. However, this formulation also means that the discretionary fiscal spending in our simulations may differ from the definition implicit in (3.2), which we have applied to construct the data for F_t used in our estimations. This difference is logically consistent if we also assume that smoothing was not practiced during our sample period. We return to this question in the next section.

It is worth noting that smoothing of the structural, non-oil deficit does not translate directly into smoothing of the actual withdrawals. The difference is made up by the automatic stabilizers in addition to new fiscal spending. Thus, the actual withdrawals with smoothing become

[12] The official Norwegian text reads, "Ved særskilt store endringer i fondskapitalen eller i det strukturelle, oljekorrigerte underskuddet fra ett år til det neste, må endringen i bruken fordeles over flere år, basert på et anslag på størrelsen på realavkastningen av Petroleumsfondet noen år fram i tid.» https://www.regjeringen.no/contentassets/64c3ac1292b04349b4f8e097dfce6c9c/no/pdfa/stm200020010029000dddpdfa.pdf, accessed Sep 23, 2021. Our translation: "In case of unusually large changes in the fund's value or in the structural, non-oil deficit from one year to the next, the change should be distributed over several years, based on an estimate of the real return of the Petroleum Fund some years into the future."

$$D_t = \lambda D_{t-1} + (1-\lambda)\bar{S}_t + X_t - \lambda X_{t-1} + F_t. \quad (4.5)$$

5. Simulation Results

We simulate our model 40 years into the future, starting from a fund value of 12 (NOK trillions 2021). Table 3 displays the quantiles of the fund value after 40 years, as well as its mean value and the percentage of scenarios where the fund is depleted by the end of the 40-year period. As a simplest first case, we constrain both fiscal variables to zero throughout the simulation period. Then, we add the automatic stabilizers and lastly also the

	Stationary exchange rate			Exchange rate random walk		
	No X. no F	X, not F	X & F	No X. no F	X, not F	X & F
5%	1.7	1.3	0.4	1.2	1.0	0.8
25%	3.8	3.5	3.1	3.4	3.3	3.1
50% (median)	6.6	6.4	6.2	6.9	6.8	6.7
75%	11.5	11.4	11.2	14.0	14.0	13.8
95%	25.3	25.5	25.1	38.9	38.9	38.6
Mean	9.2	9.0	8.7	12.0	11.9	11.7
Depletion rate	0.0%	0.3%	3.5%	0.0%	0.4%	1.8%

Table 3: Simulation results, main scenarios

Distribution of fund value 40 years hence
No explicit smoothing

discretionary fiscal policy. No further smoothing is added at this stage. Although we maintain the hypothesis of stationarity for the real exchange rate, we also include, as an illustration, simulation results for an alternative model where the exchange rate is specified as a random walk. The parameter estimates for this model are presented in the Appendix.

The first point to note in these results is the median fund value, which is little more than half of its current level under either specification. This is the very uncomfortable implication of the theorem by Dybvig and Qin discussed in Section 2. The claim, often made by Norwegian policy makers, that adherence to the Fiscal Rule is “responsible” in the sense that it preserves the fund for future generations, simply does not hold water.

We secondly note the large difference between the mean fund value under the two specifications. This difference reflects the result, also discussed in Section 2, that negative serial correlation of financial returns implies an expected future fund value that is lower than its current value even if draws are restricted to the expected return. We are surprised to see the magnitude of this effect, which cuts the expected fund value by almost 25 percent. Because we believe the stationarity assumption to be realistic, we feel this result raises a flag not previously raised in the policy debate or the literature.

The risk of depletion seems less concerning than the median fund value in these cases. It remains minor even when the automatic stabilizers are added. However, discretionary fiscal policy raises the risknoticeably, especially for the case where the real exchange rate is assumed to be stationary. With both fiscal variablesincluded, the depletion probability of 3.5 percent seems more than high enough to raise concern.

The effects of the fiscal variables clearly depend on the amplitude of the Norwegian business cycle. In our model, we assume that it rises over time at the same rate as the trend growth of mainland GDP. Table 4 presents the sensitivity to varying growth assumptions, where we look at cases with half or double the government’s official forecast of 1.4 percent. The case without fiscal variables is not included here because it is not affected by assumptions about the growth rate. Paradoxically, the depletion risk rises with the growth rate. The reason is that, when higher growth makes the economy larger in the later years, the fiscal demands on the fund during recessions become greater. And with withdrawals equal to the expected return, on average, the fund will not grow with the economy, making the fiscal strains on the fund more severe over time.

	Trend growth 0.7%		Trend growth 1.4%		Trend growth 2.1%	
	X, not F	X & F	X, not F	X & F	X, not F	X & F
5%	1.4	0.7	1.3	0.4	1.2	0.0
25%	3.6	3.2	3.5	3.1	3.5	3.0
50% (median)	6.4	6.2	6.4	6.2	6.4	6.1
75%	11.4	11.2	11.4	11.2	11.3	11.2
95%	25.4	25.1	25.5	25.1	25.5	25.1
Mean	9.1	8.8	9.0	8.7	9.0	8.6
Depletion rate	0.2%	2.2%	0.3%	3.5%	0.6%	5.3%

Table 4: Simulation results with varying trend growth rates

Distribution of fund value 40 years hence
No explicit smoothing
Real exchange rate assumed stationary

Because this effect grows over time, it may be useful to look at the same cases simulated over 80 rather than 40 years. The results of this exercise are presented in Table 5. Here, we note first of all how the mean and median shrink as the time horizon grows, even without the fiscal variables. These are the results of the phenomena described in Section 2. With the fiscal variables, the depletion risk becomes substantial. We note in particular that the depletion risk over 80 years is high even in the absence of discretionary fiscal policy. If that policy is included, and trend growth proves higher than assumed, depletion results in almost half of the scenarios.

	Trend growth 0.7%		Trend growth 1.4%			Trend growth 2.1%	
	X, not F	X & F	No X, no F	X, not F	X & F	X, not F	X & F
5%	0.0	0.0	0.5	0.0	0.0	0.0	0.0
25%	1.2	0.0	1.7	1.0	0.0	0.2	0.0
50% (median)	3.2	2.8	3.6	3.1	2.4	3.0	1.2
75%	7.6	7.2	7.8	7.5	7.3	7.5	7.5
95%	23.9	23.2	23.9	23.9	23.3	24.1	24.0
Mean	6.6	6.1	7.0	6.5	6.0	6.4	5.8
Depletion rate	8.7%	25.0%	0.0%	15.1%	35.1%	23.7%	45.9%

Table 5: Simulation results, main scenarios

Distribution of fund value 80 years hence.
No explicit smoothing.
Real exchange rate assumed stationary.
The case of “No X, no F” is unaffected by assumptions about the trend growth rate.

The business-cycle amplitude can also change independently of the trend growth rate. Our model has been estimated on data from a period where the dominating industries in Norway were oil and gas extraction, with important additions of oil supply and service[13]. The cyclical movements of these industries have at times deviated from those of neighboring and other developed economies. Because fiscal policy typically aims to stabilize business cycles, Norwegian fiscal policy may behave differently from our estimated model once the country’s petroleum era comes to an end. Predictions of how this difference may look will necessarily be speculative, however, as long as we do not know the industry structure of the Norwegian post-oil economy.

[13] On the importance of oil and gas in the Norwegian economy, see, for example, Bjørnland and Thorsrud (2016) and Mork (2022).

	Estimated business cycles			Wider business cycles		
	No X. no F	X, not F	X & F	No X. no F	X, not F	X & F
5%	1.7	1.3	0.4	1.7	1.2	0.0
25%	3.8	3.5	3.1	3.8	3.5	3.1
50% (median)	6.6	6.4	6.2	6.6	6.5	6.3
75%	11.5	11.4	11.2	11.5	11.5	11.5
95%	25.3	25.5	25.1	25.3	25.6	25.7
Mean	9.2	9.0	8.7	9.2	9.1	8.8
Depletion rate	0.0%	0.3%	3.5%	0.0%	0.8 %	5.7 %

Table 6: Simulation results, wider business cycle fluctuations

Distribution of fund value 40 years hence
Wider business cycles: Residual standard deviation in equations for X and F raised 85%
Trend growth rate: 1.4%
No explicit smoothing
Real exchange rate assumed stationary

What does seem clear, though, is that the Norwegian business cycle has been significantly smoother over the last 30 years than that of neighboring Sweden and Denmark, as expressed, for example, by the GDP gaps estimated by the IMF. Based on the differences between these estimates we add a simulation case where the residual standard deviations of the fiscal variables both are increased by 85 percent, while keeping the trend growth rate at 1.4 percent. The results for the 40-year horizon are displayed in Table 6 along with the original results in Table 3. Although the higher variability of fiscal policy increases the variability of the future fund value, the differences are much more modest than the ones in Table 5. When the wider business-cycle fluctuations hit earlier, the fund is larger in most scenarios; and the fund size acts as a cushion against greater losses.

Neither means nor medians change much. The slight increases in the median may perhaps be explained by the fact that the presence of the fiscal variables introduces an element of normality in a distribution that is otherwise lognormal, as noted in Section 2. However, the risk of depletion does increase from 3.5 percent to 5.7 percent when both fiscal variables are included. Although the results in Table 6 do not identify changes in the business cycle as a major risk, the higher risk of depletion may serve as an admonition for policy makers to watch for changes in the business cycle when oil and gas cease to be part of the picture.

6. Smoothing

As mentioned in Section 4.2, the Fiscal Rule itself implies a kind of smoothing, in that draws are calibrated to the expected financial return rather than realized returns. The fiscal variables provide additional smoothing. Even so, the text of the Norwegian Fiscal Rule allows a third layer of smoothing to avoid large, abrupt changes in the structural, non-oil deficit. We specified a formula for such smoothing in equation (4.4).

For that equation to be consistent with the rest of our model, we need to assume that smoothing has not influenced the data behind our estimated model in any important way. If it did, the observed non-oil, structural deficits would already have been smoothed, so that adding another layer of smoothing would make the model internally inconsistent.

One might hope to resolve this question empirically by estimating the value of the smoothing parameter λ implied by the available data. An estimate close to zero would then imply little or no smoothing in the data, so that adding the smoothing feature in equation (4.4) would allow simulation of future fund performance if behavior is changed from no smoothing to smoothing.

An estimation equation would have $(S_t - \bar{S}_t) \equiv F_t$ on the left and $(S_t - \bar{S}_{t-1})$ on the right. For stationarity and correction of heteroscedasticity one would furthermore want to divide both sides by potential mainland GDP. However, by definition, the error term would be identical to the left-hand variable and thus correlated with the right-hand variable if $\lambda > 0$. Furthermore, even if this problem is ignored, we find the estimate of λ to depend crucially on a few outliers in the sample. If they are ignored, the estimate would be close to unity; if they are included, it would be essentially zero.

Less technically, we would argue that the rapid growth of the fund's value during our sample period has made smoothing superfluous even on occasions when it has been intended. The political process behind the 2010 budget may serve as an illustration of this point. When presented to the Storting in October of 2009, the government's budget proposal implied a structural, non-oil deficit corresponding to 5.7 percent of the expected fund balance as of January 2010, significantly higher than the 4 percent that was then the norm[14]. The government projected this deficit to stay constant in real kroner for another seven years. The fund was not expected to grow enough for this deficit to be consistent with the fiscal rule until 2017. As things turned out, the economy recovered faster than expected, resulting in a 5 percent lower structural, non-oil deficit for 2010. The fund also grew much faster than expected, partly because of higher-than-expected oil and gas revenues, partly because of favorable exchange-rate movements, and partly because of a strong stock-market recovery after the global financial crisis. As a result, the structural, non-oil deficit stayed within the 4 percent norm already in 2010.

We interpret this episode as suggesting that our data are consistent with no significant realized smoothing, but that the political willingness to smooth nevertheless is strong enough to justify simulations with the smoothing in (4.4) added.

That leaves the question of an appropriate numerical value for the smoothing parameter λ. In the absence of reliable estimates, we treat $\lambda = 0.5$ a reasonable guess, implying that half of a gap would be closed the first year. However, we also consider $\lambda = 0.8$

14 https://www.regjeringen.no/contentassets/9c40deee7756462dad2ea5c3e3c7ef6b/no/pdfs/stm200920100001000dddpdfs.pdf, accessed January 5, 2022.

as an alternative case. This much slower smoothing, where only one fifth of a gap is closed the first year, may serve as an approximate description of the decision making in the fall of 2010. It is furthermore consistent with the so-called MIT-Tobin rule, which reportedly is practiced by a number of U.S. university endowment funds[15].

Table 7 present the implications of this smoothing behavior. In all cases, smoothing

	$\lambda = 0.5$			$\lambda = 0.8$		
	No X, no F	X, not F	X & F	No X, no F	X, not F	X & F
5%	1.6	1.2	0.0	1.2	0.7	0,0
25%	3.6	3.3	2.3	3.1	2.7	0,0
50% (median)	6.4	6.2	5.7	5.8	5.6	4,2
75%	11.3	11.2	11.0	10.9	10.8	11,2
95%	25.6	25.7	25.3	26.4	26.6	28,1
Mean	9.1	8.9	8.3	8.9	8.7	7,8
Depletion rate	0.0%	0.5%	10.9%	0.0%	1.6%	31,4 %

Table 7: Simulation results with smoothing

Distribution of fund value 40 years hence
Real exchange rate assumed stationary
Trend growth rate: 1.4%

shifts the distribution to the left and increases the risk of depletion. However, In the absence of discretionary fiscal policy, the effects are minor, even with the very persistent smoothing with $\lambda = 0.8$. The reason is that, in these cases, changes in the fund value are the only source of smoothing—recall that the rule applies to the *structural*, non-oil deficit. When discretionary policy is added, so that the structural deficit is smoothed, large changes in fiscal actions will also call for smoothing. This feature raises the depletion risk substantially even with the modest smoothing of $\lambda = 0.5$. With $\lambda = 0.8$, the fund is depleted within the 40 year horizon in almost a third of the scenarios.

[15] http://web.mit.edu/fnl/volume/205/alexander_herring.html, accessed on Sep 24, 2021. See also Tobin (1974).

This high risk of depletion raises the question of possible death spirals, as analyzed by Dybvig and Qin and referred to in Section 2. We analyze the possibility by imposing a substantial equity loss in the second year of all simulated scenarios after establishing the common starting point in the first year, while keeping the real exchange rate unchanged between the first and the second year. The subsequent development could then be called a death spiral if all subsequent scenarios lead to depletion. Table 8 presents the results when both fiscal variables are included.

The results demonstrate, first of all, that large losses leave long shadows. That

	$\lambda = 0.5$			$\lambda = 0.8$		
Shock size	10%	50%	90%	10%	50%	90%
5%	0.0	0.0	0.0	0.0	0.0	0.0
25%	2.0	0.1	0.0	0.0	0.0	0.0
50% (median)	5.1	2.8	0.5	3.7	0.0	0.0
75%	10.1	6.4	3.4	10.3	5.0	0.0
95%	23.1	15.2	9.4	26.1	16.2	8.4
Mean	7.5	4.5	2.3	7.1	3.5	1.2
Depletion rate	12.2 %	24.7 %	46.5 %	33.4 %	55.9 %	81.8 %

Table 8: Simulation of development with smoothing after initial equity shock

Distribution of fund value 40 years hence
Real exchange rate assumed stationary
Both fiscal variables included
Trend growth rate: 1.4%

follows from the fact that withdrawals are calibrated to match the expected return on average, so that the fund will not be automatically replenished after a loss. In this sense, all losses are permanent (as are all gains). Second, the risk of depletion is high in all cases[16].

[16] We assume the shocks to be concentrated in the stock market because the large negative shocks in our sample, the dotcom bubble and the global financial crisis, mainly take this shape. Bond-market shocks of similar magnitudes as those that we assume for stocks would have involved several standard errors of deviation from the mean. In the interest of full disclosure, we nevertheless want to report that these results change dramatically if the loss is extended to affect stocks and bonds equally. Because of the positive coefficients for lagged values of r_B^- in the equations for x and f, a large drop in bond returns in our model translates into a large boom in the Norwegian economy, implying negative automatic stabilizers as well as a

However, for losses as low as 10 percent, comparison with the third and sixth column in Table 7 suggests that the depletion risk is driven more by the smoothing than by the initial loss. Furthermore, no case results in 100 percent depletion, and it takes a 90 percent loss and a smoothing parameter of $\lambda = 0.8$ to raise the depletion risk above 80 percent. Obviously, the probability of a 90 percent loss in one year is low. When not even that leads to 100 percent depletion, we dare conclude that the risk of true death spirals is low. We also want to add that, after an event that is dramatic enough to cause such a large, global stock-market decline, other concerns are likely to take priority over the preservation of the fund. That said, the risk of large losses should probably not be taken lightly.

7. Remedies

Despite the many challenges facing sovereign wealth funds and endowment funds, some of which have been studied in the preceding sections, they are not doomed to depletion. Good management can reduce the risk of depletion and even avoid it completely. Naturally, any such measure will carry costs. Whether or not the gains justify the costs will be a matter of investor preferences, which we do not specify in this paper.

substantial discretionary fiscal tightening. Both effects translate into huge budget surpluses that subsequently are deposited into the fund, to a surprising extent making up for the loss. We do not find this result credible and thus do not report the details of it. We interpret it instead as a warning against relying on linear (or, in our case, quasi-linear) models to predict the effects of events far outside the range of the data on which the model has been estimated.

One such step would be to take down risk by reducing the portfolio's equity share. That would smooth out movements in the fund's value and hence in withdrawals without explicit smoothing. Table 9 shows the results of reducing the equity share from 70 percent

	No smoothing			$\lambda = 0.5$			$\lambda = 0.8$		
	No X, no F	X, no F	F & X	No X, no F	X, no F	F & X	No X, no F	X, no F	F & X
5%	2,9	2,4	1,6	2,8	2,3	0,0	2,5	2,0	0,0
25%	5,1	4,8	4,4	5,0	4,7	3,6	4,7	4,4	0,1
50% (median)	7,7	7,5	7,3	7,6	7,4	6,8	7,3	7,1	5,5
75%	11,6	11,5	11,3	11,5	11,4	11,2	11,4	11,3	11,7
95%	20,8	20,9	20,6	20,9	21,0	21,0	21,3	21,4	23,9
Mean	9,3	9,1	8,7	9,2	9,0	8,3	9,0	8,8	7,7
Depletion rate	0,0 %	0,0 %	1,0 %	0,0 %	0,0 %	5,6 %	0,0 %	0,1 %	24,6 %

Table 9: Simulation of 50 – 50 portfolio

Distribution of fund value 40 years hence
Real exchange rate assumed stationary
Trend growth rate: 1.4%

to 50 percent in our model portfolio. The results in the first three columns, with no explicit smoothing, can be compared to the three first columns of Table 3. Although the depletion risk is modest even with the 70–30 portfolio, we note that it is even lower with a less risky one. Furthermore, the medians are much higher and the expected values almost identical. The reason is that the lower volatility of the portfolio returns delays the concentration of probability mass near zero discussed in Section 2. It does not eliminate it, however.

Inclusion of the fiscal variables also do not change results nearly as much as with the 70–30 portfolio. Finally, comparison with Table 7 shows that explicit smoothing has less of an unfavorable effect as well, especially if the smoothing is moderate with $\lambda = 0.5$. The reason is that the withdrawals are smoother to begin with so that explicit smoothing does not change them very much.

The obvious cost of this strategy is that withdrawals are reduced with the expected portfolio return from 3 percent to 2 percent. Even this effect is cushioned somewhat by the fact that the fund's value tends to be preserved better over time. However, lowering the equity share does not solve the fundamental problem of the almost sure eventual depletion if withdrawals are calibrated to the expected portfolio return.

To overcome this problem, the size of withdrawals must be lower than the expected return. The question is how much. Or, put another way, how high must the expected rate of return be for 3 percent withdrawals to be sustainable? We prefer to pose the question that way because a common defense of the 3 percent rule is that it is based on a conservative estimate of the expected return. So, we may reformulate our question as whether or not the 3 percent estimate is conservative enough to preserve the fund for future generations.

For the simple case of no fiscal variables and no serial correlation of financial returns, the answer can be read directly off formula (2.6), namely, that the difference between the rate of return and the withdrawal rate must at least equal one-half of the variance of the portfolio rate of return, which in our model would be 1.4 percentage points for the 70–30 portfolio. However, in the presence of serial correlation and fiscal policy, this quantity cannot simply be derived analytically. In particular, the normal distribution of the fiscal variables changes the shape of the distribution of the future fund value from lognormal to a hybrid of the lognormal and a truncated normal with a spike at zero. Thus, almost sure preservation of the fund may not be feasible. Indeed, our simulations strongly suggest that it is not, especially in the presence of smoothing.

Rather than looking for a strategy that can ensure fund preservation, we then use the weaker criterion of preserving the median of the future distribution at least equal to the initial fund value, in our case 12. A simple binary search that this criterion is satisfied if the expected real 70–30 portfolio return is 4.76 percent. The results are presented in Table 10. This target is reached with both fiscal variables included and smoothing done at the moderate rate of $\lambda = 0.5$. In the absence of smoothing, but the same fiscal variables present, it would be preserved even after 80 years, though not with smoothing.

	40 years hence				80 years hence			
	No X, no F		X & F		No X, no F		X & F	
	No smoothing	$\lambda = 0.5$	No smoothing	$\lambda = 0.5$	No smoothing	$\lambda = 0.5$	No smoothing	$\lambda = 0.5$
5%	3.4	3.3	2.1	0.4	2.2	2.0	0.0	0.0
25%	7.6	7.4	6.7	5.8	6.6.	6.3	4.0	0.2
50% (median)	13.1	13.0	12.5	12.0	14.3	14.0	12.2	10.5
75%	22.7	22.8	22.2	22.1	31.2	31.2	28.9	27.9
95%	50.0	51.4	49.6	50.7	95.0	98.3	92.3	93.8
Mean	18.2	18.4	17.5	17.2	27.8	28.2	25.3	24.4
Depletion rate	0.0%	0.0%	1.0%	4.3%	0.0%	0.0%	12.4%	24.7%

Table 10: Simulation of 70–30 portfolio with expected rate of return of 4.76% and withdrawal rate 3%

Real exchange rate assumed stationary
Trend growth 1.4%

When we include our two fiscal variables, the spike at zero is harder to get rid of, even with moderate smoothing, especially on the 80-year horizon. In contrast, if fiscal policy is given up entirely, depletion is truly avoided—these zeroes are not the results of rounding. Smoothing does not change that conclusion. This result clearly suggests that depletion risks cannot be avoided just by keeping regular withdrawals sufficiently below the expected return. Limits would have to be put on fiscal policy as well. The extent to which that is desirable must be a matter of political preferences. The risk of depletion must be weighed against current needs for stabilization of the economy.

Would there not be warning signs along a path towards depletion? There probably would. If the fund loses much of its value over, say, the first 20 years, a correction could be implemented in the form of limits on fiscal spending, including regular withdrawals, until the fund has rebuilt itself. We suspect that such a strategy would be less efficient than setting a sustainable course from the beginning; however, to make such a judgement, we would again have to make assumptions about investor preferences.

Preferences would naturally also need to be specified in order to make combined judgments about withdrawal rates and portfolio composition. We leave such tradeoffs to the readers.

8. Conclusions

Long-term investors face a number of challenges. Among the most important is the tradeoff between withdrawals to finance current spending needs and the preservation of the fund for future generations. The rule often used of allowing draws corresponding to the expected financial return is intuitively appealing. A board of supervision, a parliament, or voters at large may find it convenient to use as a disciplinary tool or as a yardstick to judge the soundness of spending decisions or proposals. In the context of the Norwegian Government Pension Fund Global, it is the essence of the Fiscal Rule. Its main appeal is that it can be claimed to preserve the mathematical expectation of the fund's future value, seen from the present, at a level equal to the current fund value. It does not, however, guarantee the fund's long-term survival even if strictly adhered to. In fact, it leads almost surely to the fund's eventual depletion. This is not an hypothesis; it is a mathematical theorem.

Further complications are likely to pile up on top of this one. Small-country investors who diversify their assets in global markets, but make withdrawals for local use, typically face negatively correlated returns after conversion to the local currency. With such correlation, the Fiscal Rule does not even preserve the expected future value of the fund but biases it towards zero.

In practice, decision makers want flexibility to deviate from the rule. A seemingly innocent deviation is to allow countercyclical variations to ensure smooth funding of the services that the fund withdrawals are supposed to fund. In the Norwegian context, this flexibility is codified by defining the Fiscal Rule as applying to the structural rather than the actual budget deficit, so that the automatic fiscal stabilizers are allowed to work freely. In practice, draws on a sovereign wealth fund will also be used to fund discretionary fiscal policy. The effects of such deviations from the main rule may not only be large. They also change the shape of the probability distribution of future fund values. Because policy innovations are likely to be more or less normally distributed, their inclusion transforms the distribution of future fund values from lognormal to a hybrid of lognormal and truncated normal with a spike at zero. As a result, depletion becomes a real risk over relatively modest horizons.

Explicit smoothing on top of this flexibility, which is often favored by policy makers and codified in the Norwegian case, add to this risk. The result can easily become a distribution of future fund values with a mean less than the current value, an even lower median, and a substantial probability of depletion.

The mean and median problems can be remedied by requiring withdrawals at substantially lower rates than the expected rate of portfolio return. The increase in depletion risks caused by fiscal policy can furthermore be reduced by choosing less risky assets, e.g. by lowering the equity share. However, eliminating depletion risks further requires bounds on fiscal policy.

Further advice on what to do would require some kind of specification of investor preferences regarding the various tradeoffs involved. A particularly sticky point is how to specify the desire for smoothness and discretionary fiscal policy. We plan to take up that issue in a forthcoming paper.

The experience of the Norwegian Government Pension Fund Global so far has been benign with steady growth in the fund's value and no occasions where draws from the fund had to be curtailed, not even after the Global Financial Crisis. However, this 20-year period has been unusual in that sizeable deposits from the government's oil and gas revenues have been made every year. Real tests of smoothness and sustainability are much more likely to come once these revenues start to dry up.

Appendix: Parameter estimates with exchange rate assumed to be a random walk

	r_{Qt}	r_{Bt}	Δv_t	x_t	f_t
Constant	0.0600 (0.0328)	0.0309 (0.0096)			-0.0035 (0.0010)
$r^{+}_{Q,t-1}$				-0.0296 (0.0085)	
$r^{-}_{Q,t-1}$				0.0115 (0.0084)	-0.0175 (0.0075)
$r^{+}_{B,t-1}$					0.0708 (0.0239)
$r^{-}_{B,t-2}$				0.0624 (0.0264)	
Δv_{t-1}					0.0171 (0.0077)
x_{t-1}				1.1046 (0.1475)	
x_{t-2}				-0.3544 (0.1791)	
f_{t-1}				0.03306 (0.1773)	
f_{t-2}				-0.3155 (0.1720)	
SEE	0.1766	0.0570	0.1126	0.0046	0.0048

Residual correlation matrix:

	r_Q	r_B	v	x	f
r_Q	1	0	0	0	0
r_B		1	-0.4404	0	0
v			1	0.1904	0.1048
x				1	0
f					1

Table 2: Estimation results, annual data 1991 – 2020

Variables and estimation method defined as in the text
Heteroscedasticity-robust standard errors in parentheses

References

Aase, K.K. and P. Bjerksund, 2021. "The Optimal Spending Rate versus the Expected Real Return of a Sovereign Wealth Fund." *Journal of Risk and Financial Management*, 14 (425): 1 – 35.

Alhashel, B. 2015. 2Sovereign Wealth Funds: A literature review." J*ournal of Economics and Business* 78: 1–13.

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

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

Barro, R.J., 1979. "On the determination of the public debt." *Journal of Political Economy* 87: 940–971.

Bernstein, S., J. Lerner, and A. Schoar. 2013. "The Investment Strategies of Sovereign Wealth Funds." *Journal of Economic Perspectives* 27: 219–37.

Biagni, F., Y. Hu, B. Øksendal, and T. Zhang, 2008. *Stochastic Calculus for Fractional Brownian Motion and Applications*, London: Springer-Verlag

Bjerketvedt, V.S., E. Henriksen, and S. Lindset: "Fooled by the Average: Long-term Risks for Endowmet and Sovereign Wealth Funds". In progress.

Bjørnland, H.C. and L.A. Thorsrud, 2016. "Boom or Gloom? Examining the Dutch Disease in Two-speed Economies." *Economic Journal*, 126 (598): 2219 – 2256.

Blanchard, O., 2019. "Public Debt and Low Interest Rates." Presidential Lecture, *American Economic Review*, 109(4): 1197–1229.

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

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

Dreassi, A., S. Miani, and A.Paltrinieri. 2017. "Sovereign Pension and Social Security Reserve Funds: A Portfolio Analysis." *Global Financial Journal* 34: 43–53.

Dybvig, P.H. and Z. Qin, 2021. "How to Squander Your Endowment: Pitfalls and Remedies." Available at SSRN: https://ssrn.com/abstract=3939984 or https://dx.doi.org/10.2139/ssrn.3939984.

Fama, E.F. and K.R. French, 1988. "Permanent and temporary component of stock prices." *Journal of Political Economy* 96: 246–273.

Hungnes, H., 2020. "Predicting the exchange rate path. The importance of using up-to-date observations in the forecast." Statistics Norway: Discussion Papers no. 934.

Johan, S. A., A. Knill, and N. Mauck. 2013. "Determinants of Sovereign Wealth Fund Investment in Private Equity vs Public Equity." *Journal of International Business Studies* 44: 155–72.

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

Lindset, S. and E. Matsen, 2018. "Institutional spending policies: Implications for future asset values and spending." *International Markets and Portfolio Management* 32(1): 53–76. Managing Oil and Sovereign Wealth." European Economic Review 82: 113-131.

Maenhout, P.J., 2004. "Robust Portfolio Rules and Asset Pricing." *Review of Financial Studies* 17(4): 951–983.

Mandelbrot, B.B. and J.W. van Ness, 1968. "Fractional Brownian Motions, Fractional Noises and Applications. "SIAM Review 10(4): 422-437.

Meese, R.A. and K. Rogoff, 1982. "Empirical Exchange Rate Models of the Seventies: Do they fit out of sample?" *Journal of International Economics* 14: 3–24.

Mork, K.A., 2022. "The Long Norwegian Boom: Dutch Disease After All?" *The Energy Journal*, forthcoming.

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

Pastor, L. and R.F. Stambaugh, 2012. "Are Stocks Really Less Volatile in the Long Rund?" *Journal of Finance* 67(2): 431–478.

Sargent, T.J. and J. Stachurski, 2021. "Two Modifications of Mean-Variance Portfolio Theory". In *Advanced Quantitative Economics with Python*, https://python-advanced.quantecon.org/intro.html

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

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

van der Ploeg, Frederick, and Anthony J. Venables. 2011. “Harnessing Windfall Revenues: Optimal Policies for Resource-Rich Developing Economies.” The Economic Journal 121: 1–30.

Yaglom, A.M., 1958. “Correlation Theory of Processes with Random Stationary nth Increments.” In *Twelve Papers on Function Theory, Probability and Differential Equations*, American Mathematical Society Translations: Series 9, Volume 8: 87–141.