Optimal asset allocation for sovereign wealth funds

Received (in revised form): 29th April, 2008

Andreas Gintschel

is an executive director at the investment banking arm of JPMorgan, advising European pension funds and insurance companies on strategic issues related to risk and capital management and asset management. Previously, he held various positions with Deutsche Bank's Asset Management and Investment Banking Divisions, as well as Group Treasury, and was Assistant Professor of Finance at Emory University's Business School. Andreas holds a MS and PhD in Finance and Accounting from the University of Rochester, NY.

Bernd Scherer*

is MD and global head of Quantitative GTAA at Morgan Stanley. Prior to joining the firm, Bernd worked at Deutsche Bank Asset Management as head of the Quantitative Strategies Group's Research Center as well as Head of Portfolio Engineering in New York. He authored and edited six books on quantitative asset management and various articles in refereed journals. Bernd received Master's degrees in Economics from the University of Augsburg and the University of London and a PhD in Finance from the University of Giessen. He is a visiting professor at Birkbeck College (London) as well as WHU (Koblenz) and external adviser to the Swiss Finance Institute.

*Morgan Stanley, Investment Management, 25 Cabot Square, Canary Wharf, Floor 07, London E14 4QA, UK.
Tel: 44 20 7425 4016; E-mail: Bernd.Scherer@morganstanley.com

Abstract This paper develops a framework for partially hedging the market risk of oil reserves through appropriately allocating financial assets for Sovereign Wealth Funds, in particular so-called 'oil revenue' or 'petroleum' funds. Empirically, the hedge potential is substantial even when using relatively coarse partitions of the investment universe, such as Morgan Stanley Capital International (MSCI) country or industry indices. For example, if the market values of oil reserves and financial funds are equal, risk reduction is by as much as 50 per cent (10 per cent if short sales are not allowed) from original levels, translating into a certainty equivalent return of 3.26 per cent pa (48 basis points if short sales are not allowed). Moreover, choosing a portfolio along the efficient frontier, which is typically viewed as the key task in asset allocation, is relatively unimportant compared to the hedge decision.

Journal of Asset Management (2008) 9, 215–238. doi:10.1057/jam.2008.19

Keywords: *sovereign wealth funds, asset allocation, nontradeable asset, conditional Value at Risk, portfolio optimisation, oil price*

Introduction: Oil revenue funds as subclass of sovereign wealth funds

For the purpose of this paper, we define Sovereign Wealth Funds (SWF) as sovereign vehicles (returns enter the governments fiscal budget) with high foreign asset exposure, nonstandard liabilities (could be as simple as wage growth or as complicated as maintaining external purchasing power) with long (generation spanning) time horizon.¹ In this paper, we focus on SWF sourced by oil

revenue as the currently most important (biggest) fraction of this class of new investors, as can be seen from Table 1.

Among the ten biggest SWF we find eight funds that are sourced from oil revenues. Given an estimated market size of about 3tn dollars at the beginning of 2008, the three biggest oil revenue funds account for 52 of total SWF assets. Given the long-term mediocre performance of spot oil (underground wealth), SWF have been created to perform an oil to equity

Table 1 The ten biggest SWF: size and source of funding

Sovereign	Assets	Inception	Source	Weight (%)
UAE	880	1976	Oil	29.33
Norway	390	1996	Oil	13.00
Singapore	350	1981	Misc	11.67
Saudi Arabia	290	1981	Oil	9.67
Kuwait	245	1953	Oil	8.17
China	200	2007	Misc	6.67
Lybien	55	1974	Oil	1.83
Katar	49	NA	Oil	1.63
Algerien	44	2000	Oil	1.47
USA (Alaska)	39	1976	Oil	1.30

All numbers are in billion dollars and based on public sources or our own estimates as of the end of 2007.

transformation to participate in global growth. The speed of this transformation will depend on the optimal patch of extraction, which in turn depends on the impact of increased supply on oil prices, extraction costs (technology) and oil price expectations. Given an estimated 40tn dollar value of underground oil, compared to 50tn dollars in global equities, SWF will have a major impact on global equity markets. It will also lead to a shift from traditional reserve currencies (dollar, yen) to emerging market currencies, where much of the global growth is expected.

For many oil exporting countries, crude oil or gas reserves are their single most important national asset. Any change in the value of reserves directly and materially affects these countries' wealth, and thus the wellbeing of their citizens. Having recognised this, a number of oil exporting countries have been depositing oil revenues in funds dedicated to future expenditure. Devising optimal investment policies for such oil revenue funds is the aim of this paper. We analyse optimal allocations among standard partitions of the investment universe, taking into account that aggregate wealth consists of financial assets and oil reserves. We show that, in the absence of short sale constraints, the investment decision can be separated into two steps: determining (standard) efficient portfolios, and determining a zero investment, hedge portfolio, which is unique for all efficient

portfolios. The positions in the hedge portfolio depend on the oil sensitivities and the covariance matrix of the assets under consideration. Empirically, we find considerable variation in oil sensitivities across assets for sufficiently fine partitions of the universe. At the global level, there is no significant variation for equity and only low variation for debt. At the country and industry level, however, we find statistically and economically significant variation in oil sensitivities. Consequently, the hedge potential is considerable, and volatility of aggregate wealth can be reduced substantially. If, for example, oil and financial assets each represent 50 per cent of aggregate wealth, the variance of aggregate wealth is nearly cut in half when using an industry stratification of the universe. This risk reduction implies 3.26 per cent in certainty equivalent return annually. In the presence of short sale constraints, however, the hedge potential is lower, and the average (across the efficient frontier) risk reduction is 6.69 per cent, implying a certainty equivalent return of 0.48 per cent. Overall, we show that the hedge potential is substantial, and produces certainty equivalent returns exceeding those of active portfolio management.

An example of an oil revenue fund is Norway's State Petroleum Fund. The policy goals of the fund, as stated in the Norske Finansdepartementet's (Norwegian Ministry of Finance) Summary,² is '[f]irst, [...] to act

as a buffer to smooth short-term variations in the oil revenues [in the Fiscal Budget, ... and second to] serve as a tool for coping with the financial challenges connected to an ageing population and the eventual decline in oil revenues, by transferring wealth to future generations'. The second objective is operationalised as '[...] invest[ing] the capital in such a way that the fund's international purchasing power is maximised, taking into account an acceptable level of risk'. This suggests that the benchmark of the fund is future consumption in the form of imports. The same reason also motivates the inclusion of equity, which is expected to enhance the performance of the fund.³ Concerning the definition of risk, it appears that the Finansdepartementet is mostly concerned with changes in the market value of the fund. We were not able to infer the Finansdepartementet's views on operationalising the first objective, smoothing oil revenues in the short term. We believe that both objectives, smoothing revenues and maximising long-term welfare, suggest the more extensive definition of risk we propose in this paper.

As of December 2001, Norway's Petroleum Fund had a market value of approximately USD 68bn.⁴ At the same time, the estimated remaining petroleum resources were 3.7 billion tons of oil and 6,300 billion tons of gas, translating into a gross market value of approximately USD 1,143bn. During 2001, 251 million tons of oil equivalents were sold and delivered, which we value at an average price of USD 24 per barrel, or roughly USD 37bn in total. For the same period, the Norwegian Fiscal Budget shows a net cashflow of USD 27bn from petroleum activities. Thus, the government appears to be capturing about 70 per cent of revenues from oil extraction.⁵ Therefore, we estimate the government's claim on the remaining petroleum resources as 70 per cent of the gross market value, or USD 800bn as of December 2001. Thus, the value of the financial portfolio relative to the

value of the claim on oil assets is currently about 8 per cent.

Other examples⁶ of portfolios funded by revenues from natural resources include the Alaska Permanent Reserve Fund (funds of USD 23bn), the State Oil Fund of Azerbaijan (USD 0.5bn), Chad's Revenue Management Fund, the National Fund of Kazakhstan (USD 1.2bn), Venezuela's Investment Fund for Macroeconomic Stabilisation (USD 3.7bn), the Alberta Heritage Savings Trust Fund (CAD 3.7 billion), and the Nunavut Trust (CAD 0.5 billion). Furthermore, certain central bank funds of oil exporting countries, such as Iran, Kuwait, Oman, and Saudi Arabia, are defacto oil revenue funds. In general, stated investment objectives are similar to those of the Norwegian fund, that is, a favourable long-term trade-off of return and risk of the financial portfolio. The risk in aggregate wealth stemming from price changes in natural reserves is typically ignored.

More generally, our paper is an example of how risk stemming from nonfinancial assets can be hedged, at least partially, through financial assets. In other words, we talk about asset allocation with nontradeable wealth. The key is exploiting the correlation between financial and nonfinancial assets to reduce the overall risk of the portfolio, compared to an allocation that considers only the correlation structure of the financial assets. Although the general idea is straightforward, empirical or practical implementations are rare. An exception is asset/liability management, in which interest rate exposure on one side of the balance sheet is offset by interest rate exposure on the other side. This paper applies a similar idea to a more general problem.

Theoretical framework

Setup

The investor, in the present case a national government, is endowed with two assets: an

oil asset, which is the investor's claim to national oil resources, and a portfolio of financial assets. The investment policy for the portfolio is under government control. The oil asset is not necessarily under the government's control, but a fraction of future revenues, through taxation, usually is.

Restricting the analysis to these types of assets is obviously an abstraction from reality. A more complete analysis would include additional real assets such as a country's nonoil capital stock and human capital. In general, such an extension is straightforward. For most countries that have an oil revenue fund, however, the value of nonoil assets is small relative to the value of oil reserves, or nonoil assets are highly oil-related.

For the present, we assume that the value of oil assets and the fund changes only due to price changes of the underlying assets. In other words, we do not explicitly consider depletion of oil reserves, increases in oil reserves due to exploration, or cash flows into or out of the monetary funds. We do, however, consider a range of ratios of the value of oil resources to the value of the financial assets. Therefore, we present a dynamic, albeit myopic, asset allocation policy as oil resources are depleted and financial assets increase.

The current amount of oil resources is x_0 , denominated in million barrels, and the value per barrel in U\$ is p_o . The interpretation of this value depends on a variety of factors. In general, it reflects the fraction of the price per barrel that the government deposits in the financial portfolio. The total value of the government's claim on oil resources is $x_0 p_o$. The relative change in oil prices over the time period considered in the analysis is denoted as r_o .

The current value of the monetary fund is v_f , denominated in U\$ million. The current relative portfolio weight on an individual asset or an asset class is w_i such that $\sum w_i = 1$. The return on a asset i over the period of the analysis is r_i . The return on the portfolio is r_p . Therefore, the change

in total wealth, oil reserves and financial assets is $r = \omega r_o + (1 - \omega) r_p$, where $\omega = x_0 p_o / (x_0 p_o + v_f)$, the value of oil reserves relative to aggregate wealth. For example, the Norwegian fund currently has a ratio of $\omega = 0.92$. For ease of exposition, we assume that returns are distributed multivariate normal. This implies that expected return and volatility completely characterise portfolio return distributions. Together with constant relative risk aversion, normality implies that the standard mean-variance framework for preferences is applicable. We assume that the function $u = E[r_i] - (1/2)\alpha \text{Var}[r_i]$ adequately describes the investor's preferences, where $E[\]$ and $\text{Var}[\]$ denote the expected return and the variance, and α is the degree of risk aversion.

Oil exposure

We measure financial assets' oil exposure by the oil sensitivity of the assets' return, $b_i = \text{Cov}(r_i, r_o) / \sigma_o^2$. We estimate b_i as the regression coefficient of historical asset returns on contemporaneous oil price changes and an intercept. We collect the assets' oil sensitivities in a vector $\mathbf{b}$. Accordingly, we define the oil exposure of a financial portfolio $\mathbf{w}$ as $\mathbf{w}^T \mathbf{b}$. The variance of relative changes in aggregate wealth is

$$\text{Var}(r) = \omega^2 \sigma_o^2 + 2\omega(1 - \omega) \sigma_o^2 \mathbf{w}^T \mathbf{b} + (1 - \omega)^2 \mathbf{w}^T \mathbf{\Sigma} \mathbf{w} \quad (1)$$

where $\mathbf{\Sigma}$ is the covariance matrix of the financial assets' returns. The equation shows that the risk of aggregate wealth depends on the oil exposure of the financial portfolio. Since the portfolio's oil exposure is the weighted average of the individual assets' oil exposure, the risk of aggregate wealth depends immediately on the oil exposure of the individual assets. Holding constant the volatility of the financial portfolio, investing in assets with small oil beta reduces overall volatility.

The traditional measure of financial portfolio risk is only the last term, ignoring the weighting factor $(1-\omega)$. Equation (1) shows that the traditional measure misestimates the risk of aggregate wealth, since it ignores the volatility of oil price and also the correlation between the prices of oil and financial assets. For any given financial portfolio, equation (1) allows us to compute the oil exposure of the portfolio by taking the average, weighted by the portfolio positions, of the individual assets' oil betas.

Efficient portfolios

According to the standard definition, an efficient portfolio minimises the portfolio variance for a given expected return. There are, in the present application, however, two possible choices for the portfolio: the financial portfolio and the combined portfolio including oil reserves. Standard portfolio construction, as typically applied in practice, requires financial portfolios to be efficient, that is, to solve the program $\mathbf{w}^T \Sigma \mathbf{w}$ subject to achieving a target expected return μ , $\bar{\mathbf{z}}^T \mathbf{w} = \mu$, and a budget constraint $\mathbf{1}^T \mathbf{w} = 1$. The solution, depending on the target return μ , we denote as $\mathbf{w}_L(\mu)$. From a more comprehensive point of view, only the combined portfolio need be efficient. We call the financial portfolio, $\mathbf{w}_L(\mu)$, which is efficient in isolation, 'locally efficient', while a financial portfolio $\mathbf{w}_G(\mu)$ is 'globally efficient' if it yields an efficient combined portfolio, that is, solves the program

$$\begin{aligned} \text{Var}(r) = & \omega^2 \sigma_o^2 + (1-\omega)^2 \mathbf{w}^T \Sigma \mathbf{w} \\ & + 2\omega(1-\omega) \sigma_o^2 \mathbf{w}^T \mathbf{b} \end{aligned} \quad (2)$$

subject to the constraints $\mathbf{1}^T \mathbf{w} = 1$ and $\bar{\mathbf{z}}^T \mathbf{w} = \mu$. These constraints do not ensure, however, that assets have generally positive weights, that is, we allow for unrestricted short positions in the financial portfolio. In our empirical analysis we relax this

assumption, but we confine the theoretical discussion to the analytically tractable case of unrestricted short positions. We stress that we define the programs such that the solutions, which are financial portfolios, have the same expected return μ . Thus, for a given ratio of oil resources to value of financial assets, the combined portfolios, based on either $\mathbf{w}_L(\mu)$ or $\mathbf{w}_G(\mu)$, have the same expected return, independent of the expected oil price change. The advantage is that we need not make any assumptions on the expected oil price change. The return on the combined portfolio is, however, in general, different from μ . As Appendix A shows, the globally efficient portfolio for target return μ is

$$\mathbf{w}_G(\mu) = \mathbf{w}_L(\mu) + \frac{\omega}{(1-\omega)} \mathbf{w}_H \quad (3)$$

where

$$\begin{aligned} \mathbf{w}_H = & \sigma_o^2 \left\{ -\Sigma^{-1} \mathbf{b} + \Sigma^{-1} \bar{\mathbf{z}} \frac{1}{\Delta} [(1^T \Sigma^{-1} \mathbf{1})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{b}) \right. \\ & - (\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})(1^T \Sigma^{-1} \mathbf{b})] \\ & + \Sigma^{-1} \mathbf{1} \frac{1}{\Delta} [(1^T \Sigma^{-1} \mathbf{b})(\bar{\mathbf{z}}^T \Sigma \bar{\mathbf{z}}) \\ & \left. - (\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{b})(1^T \Sigma^{-1} \bar{\mathbf{z}})] \right\} \end{aligned} \quad (4)$$

is a zero-net investment, zero-expected return hedge portfolio that does not depend on the expected target return μ , and $\Delta = (\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}})(1^T \Sigma^{-1} \mathbf{1}) - (1^T \Sigma^{-1} \bar{\mathbf{z}})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})$. Thus, given a locally efficient portfolio for any target expected return, we can easily construct the corresponding globally efficient portfolio using the hedge portfolio $\mathbf{w}_H$. In other words, we have established a useful fund separation property.

The composition of the hedge portfolio is easily interpreted. The first term in equation (4) calculates the optimal hedge, the second ensures that expected portfolio returns are zero, and the third term scales portfolio weights such that the portfolio requires no net investment. Finally, equation (4) provides an answer to the important question of whether a locally efficient portfolio is also

globally efficient. Inspection of equation (4) shows that necessary conditions are fairly complex. Sufficient conditions are, however, easy to state. The trivial sufficient condition is $\mathbf{b} = \mathbf{0}$, that is, no financial asset exhibits any oil sensitivity. A more interesting condition is $\Sigma^{-1}\mathbf{b} = \mathbf{0}$, which has the form of a system of linear equations in $\mathbf{b}$. It follows that there is, for a given (nonsingular) covariance matrix, exactly one vector of oil sensitivities that satisfies the condition. In other words, the condition is unlikely to be satisfied in practice.

Measuring the benefit of reducing oil exposure

To gauge the effect of taking into account oil exposure when constructing financial portfolios, we employ two measures: the reduction in the variance of aggregate wealth, and the certainty equivalent return associated with the variance reduction. The first measure is preference free, while the second requires assumptions regarding the investor's preferences. The variance reduction by taking into account financial assets' oil sensitivity is

$$\begin{aligned}\Delta Var(r|\mu) &= Var(r|\mathbf{w}_L(\mu)) - Var(r|\mathbf{w}_G(\mu)) \\ &= (1 - \omega)^2 [\mathbf{w}_L(\mu)^T \Sigma \mathbf{w}_L(\mu) \\ &\quad - \mathbf{w}_G(\mu)^T \Sigma \mathbf{w}_G(\mu)] \\ &\quad - 2\omega^2 \sigma_o^2 [\mathbf{w}_H^T \mathbf{b}]\end{aligned}\quad (5)$$

The first term is the difference between the variance of the locally efficient financial portfolio and the globally efficient financial portfolio, which is negative. The second term, after the minus sign, is the oil sensitivity of the hedge portfolio, which is also negative. Thus, the variance reduction is the result of trading off increasing volatility of the financial assets against decreasing oil sensitivity. The variance of well-diversified portfolios is typically not very sensitive to the exact portfolio composition. Thus, the first term in equation (5) is approximately zero,

and, since the second term does not depend on μ ,

$$\Delta Var(r) \approx -2\omega^2 \sigma_o^2 [\mathbf{w}_H^T \mathbf{b}] \quad (6)$$

The equation shows also that the hedge portfolio's oil sensitivity is a useful summary measure of the risk reduction potential for a particular set of assets. The smaller the hedge portfolio's oil sensitivity, the larger the reduction in overall, that is, financial and oil-related risk.

Since the economic importance of a variance reduction can be difficult to interpret, we present the gain in certainty equivalent return corresponding to the variance reduction. The gain in certainty equivalent return is the amount of riskless return the investor would be willing to give up in exchange for reducing risk by the specified amount, that is,

$$\begin{aligned}\Delta CER &= \frac{1}{2} \alpha \Delta Var(r) \\ &\approx -\frac{1}{2} \alpha \omega^2 \sigma_o^2 \mathbf{w}_H^T \mathbf{b}\end{aligned}\quad (7)$$

For the risk aversion parameter, we choose a moderate value of three.

Data

Here we give only a brief overview of the data and the assumption underlying the analysis. Further details are in Appendix B. As proxies for financial assets, we use standard, broad asset classes. Using widely diversified portfolios reduces estimation error in the covariance matrix of asset returns and in oil sensitivities. On a more practical level, restricting the investable universe to standard asset classes facilitates implementation within existing investment processes.

As examples, we work with four different investable universes, each comprising government debt and different stratifications of developed market equity. For global government debt, we use standard benchmarks for the three major economic blocks: EMU, Japan, and US. For equity,

we use the MSCI developed market indices in four different stratifications: major economic blocks (Europe, North America, and Pacific), countries, industry sectors, and industries. As an estimate of the covariance matrix, we use the historical covariance matrix. Expected returns on government debt are the yields corrected for currency effects. We estimate the expected returns for equity as the expected return implied by market capitalisation weights. All return series are in local currency, which is equivalent to assuming that currency risk is perfectly hedged.

To model oil price risk, we rely on historical data for Dated Brent Crude Oil available from Independent Commodity Information Services (ICIS). In the next section, we describe the estimator of assets' oil price sensitivities and discuss results. As we note in the theoretical development, we can dispense with estimating expected returns for oil.

Oil sensitivity

Table 2 contains evidence on the individual asset classes' oil sensitivity, which we estimate as the slope coefficients in univariate time-series regressions

$$r_{i,t} = a_i + b_i r_{o,t} + u_{i,t} \quad (8)$$

where a_i is the intercept and $u_{i,t}$ is the error term in the regression. Next to the estimate of the slope coefficient, we report standard errors. Slope coefficients that are significantly different from zero are typeset in bold face. From the theoretical analysis above, oil sensitivities that are large in absolute magnitude and widely dispersed indicate substantial hedging potential.

Scanning Table 1, panels A and B reveal that oil sensitivity for sovereign debt and regional equity is generally low, negative, and not very dispersed. At the global level, oil price increases affect financial assets' prices negatively, as expected. The magnitude of the effect is surprisingly small, given the

considerable importance that economists usually ascribe to oil price changes. Thus, allocating among these broad asset classes promises few advantages in hedging oil price risk. Turning to the country stratification in panel C, more interesting results emerge. For some countries, we find statistically and economically significant oil sensitivities. In particular, for Belgium, France, Ireland, and Spain oil sensitivities are significantly negative. For countries rich in natural resources, such as Australia, Canada, and Norway, we find positive, albeit statistically insignificant oil sensitivities. Overall, these results suggest that oil sensitivities are sufficiently dispersed as to have an impact on asset allocation. For the sector stratification in panel D, we find widely different oil sensitivities across industrial sectors. In particular, the energy and information technology sectors exhibit statistically and economically positive oil sensitivity, that is, these sectors' equity prices move in the same direction as oil prices to a considerable degree. In contrast, sectors such as consumer staples or utilities have negative oil price sensitivity. Of course, the industry stratification, which is just a finer partition of the sector stratification, exhibits similar patterns. For some industries, oil sensitivities are surprising. For example, the automobile industry has virtually zero oil sensitivity (0.0009), while the media sector has a substantial, positive oil sensitivity (0.0792).

In Table 3, we summarise the information on oil sensitivity for the different asset classes. For each broad class, we show the arithmetic average of oil sensitivities and the corresponding standard deviation across asset classes. The standard deviation indicates the cross-sectional dispersion of the estimates, measuring the hedge potential of a particular partition of the universe. If a partition leads to only tightly dispersed oil sensitivities, reallocation among the members of this partition cannot change the oil sensitivity of the portfolio by much. Conversely, if a partition is widely dispersed, changing the

Table 2 Oil price sensitivity of financial assets

	Oil sensitivity	Standard error	Correlation
<i>Panel A Debt</i>			
Europe	−0.0168	0.0070	−0.17
Japan	−0.0146	0.0078	−0.13
US	−0.0168	0.0082	−0.15
<i>Panel B Regional equity</i>			
Europe	−0.0409	0.0283	−0.10
North America	−0.0183	0.0271	−0.05
Pacific	−0.0071	0.0391	−0.01
<i>Panel C Country equity</i>			
Australia	0.0552	0.0413	0.10
Austria	−0.0594	0.0409	−0.10
Belgium	−0.0858	0.0309	−0.20
Canada	0.0194	0.0287	0.05
Denmark	0.0135	0.0329	0.03
Finland	−0.0046	0.0712	0.00
France	−0.0751	0.0362	−0.15
Germany	−0.0433	0.0378	−0.08
Hong Kong	−0.0265	0.0535	−0.04
Ireland	−0.0815	0.0407	−0.14
Italy	−0.0371	0.0443	−0.06
Japan	−0.0206	0.0371	−0.04
Netherlands	0.0055	0.0301	0.01
New Zealand	0.0277	0.0452	0.04
Norway	0.0818	0.0416	0.14
Portugal	−0.0868	0.0448	−0.13
Singapore	−0.0092	0.0474	−0.01
Spain	−0.1027	0.0409	−0.18
Sweden	−0.0538	0.0454	−0.09
Switzerland	−0.0459	0.0320	−0.10
United Kingdom	−0.0430	0.0290	−0.11
United States	−0.0209	0.0274	−0.06
<i>Panel D Sector equity</i>			
Energy	0.1440	0.0425	0.33
Materials	0.0211	0.0474	0.05
Industrials	0.0573	0.0409	0.15
Consumer discretionary	0.0287	0.0473	0.06
Consumer staples	−0.0613	0.0353	−0.18
Health care	−0.0489	0.0400	−0.13
Financials	−0.0326	0.0499	−0.07
Information technology	0.2192	0.0860	0.26
Telecommunications	0.0277	0.0539	0.05
Utilities	−0.0341	0.0292	−0.12
<i>Panel E Industry equity</i>			
Energy	0.1440	0.0425	0.33
Materials	0.0211	0.0474	0.05
Capital goods	0.0725	0.0449	0.17
Commercial services	0.0401	0.0425	0.10
Transportation	0.0012	0.0354	0.00
Automobiles	0.0009	0.0584	0.00
Consumer durables	0.1083	0.0531	0.21
Hotels and leisure	−0.0078	0.0475	−0.02
Media	0.0792	0.0555	0.15
Retail	−0.0566	0.0525	−0.11
Food & drug retail	−0.0454	0.0345	−0.14
Food, beverage and tobacco	−0.0726	0.0373	−0.20
Household and personal products	−0.0481	0.0513	−0.10
Health care	−0.0611	0.0487	−0.13
Pharmaceuticals	−0.0465	0.0418	−0.12
Banks	−0.0226	0.0539	−0.04
Diversified financials	0.0058	0.0564	0.01

Table 2 Continued

	Oil sensitivity	Standard error	Correlation
Insurance	-0.0945	0.0480	-0.20
Real estate	-0.0198	0.0576	-0.04
Software	0.2345	0.0953	0.25
Hardware	0.2123	0.0876	0.25
Telecommunications	0.0277	0.0539	0.05
Utilities	-0.0341	0.0292	-0.12

Oil sensitivity is the slope coefficient of individual time-series regressions of asset returns on oil price changes. Correlation is the simple correlation coefficient between asset returns and oil price changes. The sample consists of monthly return observations between January 1988 (except when return series from MSCI are shorter) to December 2001. Data for government debt are from Merrill Lynch, data for equities are from MSCI, and data on Brent.

Table 3 Summary statistics

Stratification	Oil sensitivity		F Statistic:		Correlation	
	Average	Standard deviation	Beta = 0	p value (%)	Average	Standard deviation
Debt	-0.0161	0.0013	4.72	11.17	-0.15	0.02
Regional equity	-0.0221	0.0172	1.47	43.37	-0.06	0.05
Country equity	-0.0272	0.0479	2.17	1.49	-0.05	0.09
Sector equity	0.0321	0.0895	3.37	1.80	0.04	0.17
Industry equity	0.0191	0.0875	2.72	0.26	0.02	0.15

The table reports the average oil sensitivity and correlation within each stratification of the investment universe, and the corresponding cross sectional standard deviation within each stratification. The *F* statistic is computed under the null hypothesis that all oil sensitivities are zero within a particular stratification.

allocation has a large impact on portfolio oil sensitivity.

Table 3 shows that the dispersion of oil sensitivities is, at 0.0013, very low for debt. Oil sensitivities for the regional stratification are more widely dispersed at 0.0172. Partitioning by individual country increases the dispersion to 0.0479. The dispersion of the partition based on industrial sectors is even greater at 0.0895. The industry stratification is slightly less dispersed at 0.0875. Table 2 also contains average oil sensitivities for the various asset classes and partitions. As a reference, we calculate the oil sensitivity of the market portfolio, that is, the value-weighted average of the asset classes' oil sensitivities. Using the market values of the corresponding indices, we find an oil sensitivity of -0.0251 for the value-weighted

portfolio of debt and equity. The oil sensitivity of the Norwegian Petroleum Fund is -0.0238. The last two columns show that the dispersion is statistically significant. For each asset class, we calculate the *F* statistic for the null hypothesis that the oil sensitivity of every asset is zero. The last column reports the corresponding *p* values. We cannot reject the hypothesis that the different classes of debt or regional equity have no oil sensitivity at conventional levels of statistical significance. In contrast, the oil sensitivities of the country, sector, and industry portfolios are significantly different from zero.

Hedge portfolios

In the absence of short-sale restrictions, we can easily determine globally efficient

Table 4 Hedge portfolio

	Stratification				
	Regional	Country	Sector	Industry	
$w_H^T b$	−0.0119	−0.2840	−0.5004	−0.6162	
Aggregate short position	84%	341%	666%	745%	
Weight on debt	3%	−6%	−43%	−36%	
	Weights in hedge portfolio (%)				
	Oil sensitivity	Regional	Country	Sector	Industry
Debt					
Europe	−0.0168	−33	−26	−161	−166
Japan	−0.0146	−4	−9	135	179
US	−0.0168	40	28	−17	−50
Regional equity					
Europe	−0.0409	44			
North America	−0.0183	−32			
Pacific	−0.0071	−16			
Country equity					
Australia	0.0552		−30		
Austria	−0.0594		12		
Belgium	−0.0858		69		
Canada	0.0194		−3		
Denmark	0.0135		−42		
Finland	−0.0046		−1		
France	−0.0751		23		
Germany	−0.0433		7		
Hong Kong	−0.0265		−1		
Ireland	−0.0815		44		
Italy	−0.0371		−23		
Japan	−0.0206		−10		
Netherlands	0.0055		−104		
New Zealand	0.0277		−11		
Norway	0.0818		−61		
Portugal	−0.0868		20		
Singapore	−0.0092		6		
Spain	−0.1027		33		
Sweden	−0.0538		19		
Switzerland	−0.0459		33		
United Kingdom	−0.043		46		
United States	−0.0209		−20		
Sector equity					
Energy	0.144			−174	
Materials	0.0211			105	
Industrials	0.0573			−280	
Consumer discretionary	0.0287			110	
Consumer staples	−0.0613			0	
Health care	−0.0489			47	
Financials	−0.0326			129	
Information technology	0.2192			−17	
Telecommunications	0.0277			−17	
Utilities	−0.0341			140	

Table 4 Continued

<i>Industry equity</i>		
Energy	0.1440	–185
Materials	0.0211	75
Capital goods	0.0725	–128
Commercial services	0.0401	–37
Transportation	0.0012	–43
Automobiles	0.0009	33
Consumer durables	0.1083	–22
Hotels and leisure	–0.0078	–7
Media	0.0792	–43
Retail	–0.0566	89
Food and drug retail	–0.0454	–34
Food, beverage and tobacco	–0.0726	6
Household and personal products	–0.0481	–25
Health care	–0.0611	21
Pharmaceuticals	–0.0465	14
Banks	–0.0226	23
Diversified financials	0.0058	–5
Insurance	–0.0945	104
Real estate	–0.0198	36
Software	0.2345	2
Hardware	0.2123	–1
Telecommunications	0.0277	1
Utilities	–0.0341	162

The table contains the composition of the hedge portfolio and selected summary statistics, where $w_{i,b}$ is the oil sensitivity of the hedge portfolio. Aggregate short positions are the sum of all negative weights in the hedge portfolio. Weight on debt is the aggregate weight on government debt.

portfolios, that is, portfolios minimising overall risk for a given target expected return. As we show in the first section, for a given universe and any locally efficient financial portfolio $\mathbf{w}_L(\mu)$ there is a unique hedge portfolio $\mathbf{w}_H$. A globally efficient portfolio $\mathbf{w}_G(\mu)$ is the locally efficient portfolio $\mathbf{w}_L(\mu)$ combined with the hedge portfolio $\mathbf{w}_H$, weighted according to the value of oil assets relative to financial assets.

In Table 4, we present the hedge portfolios for each investment universe. The first row contains the oil sensitivity of the hedge portfolio. The oil sensitivity of the hedge portfolio ranges from -0.0119 for the regional debt and equity stratification to -0.6162 for the industry equity universe. The oil sensitivity of the hedge portfolio for the sector stratification is, at -0.5004 , slightly lower in absolute magnitude. The hedge portfolio's oil sensitivity for the country stratification is at the intermediate level of -0.2840 . Patterns of the hedge

portfolio's oil sensitivity are consistent with the dispersion of individual assets' oil sensitivities of a particular stratification in Table 3: a wider dispersion of individual oil sensitivities translates into lower oil sensitivity of the hedge portfolio.

Table 4 also contains the weights of the individual assets in the hedge portfolio. We focus our discussion first on the simplest universe of regional debt and equity. After giving the interpretation of the numbers and describing general features for this example, we will highlight salient features for other stratifications. For regional debt and equity, the weights are -33 per cent on European debt, -4 per cent on Japanese, 40 per cent on US debt, 44 per cent on European equity, -32 per cent on North American equity, and -16 per cent on Pacific equity. It is easy to check whether this is a zero net investment portfolio. Overall, the hedge portfolio is tilted slightly towards debt, with an aggregate weight on debt of 3 per cent.⁷

The aggregate weight on equity is -3 per cent, since the hedge portfolio is zero net investment. Intuitively, one might expect a direct relation between individual assets' oil sensitivities and the weights in the hedge portfolio; holdings of assets with positive oil sensitivity, increasing overall risk, should be reduced, and holdings of assets with negative sensitivity should be increased. Interestingly, this is not necessarily the case. For example, despite having exactly the same oil sensitivity (-0.0168), displayed in the first column of Table 3, the weights on European debt (-33 per cent) and US debt (40 per cent) have the opposite sign. Similarly, the magnitude of the short position in Japanese debt (-4 per cent) is smaller than the short position in European debt (-33 per cent), although Japanese debt has the lower oil sensitivity (-0.0146). The reason for these apparently counterintuitive weights is the existence of cross-effects from the covariance matrix. Fundamentally, a financial asset in the hedge portfolio has two roles: to act as a hedge against oil risk and to hedge the financial risk stemming from another financial asset. An asset with a counterintuitive weight acts in the second capacity, hedging risk stemming from an asset that acts as a hedge against oil. This dual role of the financial assets is also highlighted in the dramatic increase, in absolute magnitude, of the hedge portfolio's oil sensitivity as the partition of the investment universe becomes finer.

We turn now to the structure of the hedge portfolio for the other stratifications of the universe. For the country and industry stratification, the hedge portfolio is tilted towards equity, with debt weights ranging from -6 to -43 per cent. The country stratification leads to a hedge portfolio with moderate weights, the smallest of -104 per cent on the Netherlands and the largest of 69 per cent on Belgium. The sector equity universe yields more extreme portfolio weights, ranging from -280 per cent on industrials to 140 per cent on utilities. The industry stratification also delivers large

weights, albeit less extreme. In general, we find hedge portfolio weights in line with oil sensitivities. Countries such as Belgium or Switzerland that have a negative exposure to oil prices receive positive weights in the hedge portfolio. On the other end of the spectrum, countries such as the Netherlands or Norway, in which natural resources companies make up a substantial fraction of market capitalisation, receive negative weights in the hedge portfolio. As we point out, however, the relation between oil sensitivity and hedge portfolio weight is not necessarily monotone. Italy, for example, has a negative oil sensitivity, yet receives a negative weight in the hedge portfolio. Similar patterns emerge for the sector and industry stratifications. Not surprisingly, energy producers receive negative weights in the hedge portfolio, while energy users, such as utilities, have positive weights. In Table 4, we also report the aggregate short positions as a measure of the hedge portfolio's leverage. For the combination of debt and regional equity, the aggregate short positions is 84 per cent. For the industry stratification, the aggregate short position is 745 per cent. In other words, the hedge portfolios are highly geared at ratios between, roughly, one and seven.

Globally efficient portfolios

Combining the hedge portfolio and the locally efficient financial portfolio delivers the globally efficient portfolio, that is, a financial portfolio that minimises the risk of total wealth (oil and financial assets). In this section, we present and discuss the reduction in variance and the additional certainty equivalent return that switching from locally efficient to globally efficient allocations can attain. These statistics quantify the advantage of taking into account the correlation between oil and financial assets in the asset allocation process.

We determine the locally efficient frontier, ranging from 4.4 per cent expected

Figure 1 Locally and globally efficient frontiers. Country stratification and $\omega = 50$ per cent

return (the yield of debt) to 7.5 per cent expected return (the highest expected return for the coarsest partition) in increments of 10 basis points, and combine it with the hedge portfolio, weighted by the appropriate ω , delivering the globally efficient portfolio. For both the locally and the globally efficient portfolio we calculate the variance of the combined portfolio (oil and financial assets). As an example, Figure 1 contains both frontiers for $\omega = 50$ per cent and the country stratification. On the vertical axis, we plot the expected return of the financial portfolio, and on the horizontal axis, we plot the volatility of the combined (oil and financial assets) portfolio. By construction, the globally efficient frontier plots strictly to the left (has a strictly lower volatility) of the locally efficient frontier. What is surprising, however, is the magnitude of the difference, which is roughly 3 per cent in this particular example. The graph also shows that choosing an allocation along the frontier is far less important than choosing the relevant frontier, that is, whether to hedge or not. Allocating along the locally efficient frontier, volatility of aggregate wealth ranges roughly from 20.5 to 21.5 per cent, compared to approximately 3 per cent volatility reduction moving to the globally efficient frontier.

From these two frontiers we proceed to calculate the relevant statistics. As an example, Figure 2 contains the gain in

certainty equivalent return for the country stratification of the universe and a value of $\omega = 50$ per cent. The gain in certainty equivalent return, around 1.7 per cent, barely varies along the efficient frontier. This is true in general, not only in this particular example. Therefore, we restrict ourselves to reporting averages of variance reductions and added certainty equivalent returns. Figure 3 contains the graph of the variance reduction scaled by the total variance of the locally efficient portfolio. Given that the variance reduction itself is stable across the frontier, not surprisingly the relative (or percentage) variance reduction varies with the variance of the portfolio. For consistency as well as brevity, we also report only the average across the frontier.

In Table 5, we report statistics separately for varying relative values of oil to financial assets ($\omega = 25, 50, 75, 90$ per cent) and for different stratifications of the investment universe. In discussing the results, we focus on the case $\omega = 50$ per cent, which is in the second block of columns. The reduction in variance ranges from 0.0005 for the regional stratification (debt and equity for the three major economic blocs) to 0.0217 for the industry stratification. Correspondingly, the relative variance reduction (in per cent of the variance of the locally efficient portfolio) ranges from 0.25 per cent for the regional stratification to 48.47 per cent for the

Figure 2 Gain in certainty equivalent returns across the frontier. Country stratification and $\omega = 50$ per cent

Figure 3 Relative reduction in variance. Across the frontier and $\omega = 50$ per cent

industry stratification. The country stratification yields an intermediate variance reduction of 26.17 per cent. These figures show that the potential for reducing risk in the overall portfolio through appropriate allocation in the financial portfolio is considerable. The added certainty equivalent returns, measuring the economic importance of the variance reduction, emphasise the same fact. The increase in certainty equivalent return (evaluated at a risk aversion coefficient of three) for switching from locally to globally efficient portfolios ranges from 8 basis points for the regional stratification to 3.26 per cent pa for the industry stratification. For the country stratification, the added certainty equivalent return is 1.70 per cent. Overall, proper asset allocation yields economically significant utility gains that are far beyond the figures

that active portfolio management can generate.

Scanning across the column blocks of Table 5 shows that the potential for variance reduction is larger when oil constitutes a large proportion of aggregate wealth. For example, the added certainty equivalent return for the country stratification is 5.52 per cent at $\omega = 90$ per cent and 1.7 per cent at $\omega = 50$ per cent. This might be surprising given that a small financial portfolio can conceivably hedge only small amounts of oil risk. We allow, however, unconstrained short positions in financial assets. Thus, the potential benefit from hedging depends only on the contribution of oil to the variance of the combined portfolio. The financial portfolio can be levered to any degree to offset the oil risk. Empirically, the leverage in the globally efficient financial portfolio can

Table 5 Locally versus globally efficient portfolios

Universe	$\omega = 25\%$			$\omega = 50\%$		
	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)
Regional	0.0001	0.11	0.02	0.0005	0.25	0.08
Country	0.0028	21.14	0.43	0.0114	26.17	1.70
Sector	0.0044	31.74	0.66	0.0176	39.35	2.64
Industry	0.0054	39.06	0.81	0.0217	48.47	3.26
Universe	$\omega = 75(\%)$			$\omega = 90(\%)$		
	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)
Regional	0.0012	0.38	0.17	0.0017	0.45	0.25
Country	0.0255	26.63	3.83	0.0368	26.55	5.52
Sector	0.0396	40.81	5.95	0.0420	30.18	6.30
Industry	0.0489	50.30	7.33	0.0704	50.58	10.56

The table contains summary measures on the economic advantages from hedging oil risk, where $\Delta \text{Var}(r)$ is the difference between the variance of the locally efficient portfolio and the globally efficient portfolios and $\% \Delta \text{Var}(r)$ is $\Delta \text{Var}(r)$ divided by the variance of the locally efficient portfolio, while ΔCER denotes the gain in certainty equivalent return from switching from the locally efficient portfolio to the globally efficient portfolio. w is the relative value of oil reserves to aggregate wealth.

be substantial. For example, the aggregate short positions for $w = 50$ per cent and the country universe range from roughly 670 per cent to about 815 per cent across the efficient frontier, with an average of 740 per cent.

Globally efficient portfolios under short-sale constraints

Short-sale constraints imply that the allocation problem is a convex program with linear inequality constraints. From the previous section — the financial portfolio is highly levered — we know that these constraints are binding. Consequently, the potential for reducing the variance of total (oil and financial) wealth is strictly less than in the absence of short-sale constraints. We repeat the analysis from the previous section when short sales are not allowed. In Figures 4 and 5, we present both the globally and the locally efficient frontier for financial portfolios with positive weights for $\omega = 50$ per cent and the country specification. Necessarily, the globally efficient portfolio plots strictly to the left of the locally efficient portfolio, except at the minimum return

portfolio and the maximum return portfolio, which coincide. While the average difference in volatility between the two frontiers is, at 0.49 per cent pa, much smaller than in the unconstrained case, it is still economically significant. Figure 5 shows the gain in certainty equivalent return for switching from the locally to the globally efficient frontier. The certainty equivalent return increases along the efficient frontier from approximately zero for the minimum return portfolio to its maximum of 0.43 per cent for the financial portfolio with 7 per cent expected return, and drops afterwards. Intuitively, large gains of diversification can be realised by shifting from debt, which has low oil sensitivity, into suitable equity. Generating very high expected returns, however, requires disregarding oil sensitivity, hence the hump-shaped graph in the presence of short-sale restrictions.

Figure 6 shows the globally efficient allocations for the case $w = 50$ per cent across the frontier. In Figure 7, we report the locally efficient allocations. Comparing the two graphs, equity of globally efficient allocations is heavily concentrated in such

Figure 4 Locally and globally efficient frontiers. Country stratification and $\omega = 50$ per cent under short-sale constraints

Figure 5 Gain in certainty equivalent returns across the frontier. Country stratification and $\omega = 50$ per cent under short-sale constraints

countries as Belgium, France, Ireland, Portugal, and Spain, that is, countries with negative oil sensitivities. Heavily capitalised markets such as the United States or the United Kingdom, which dominate locally efficient allocations, receive moderate weights only for relatively high expected returns. Concentration in a small number of markets with low oil sensitivities is also typical for globally efficient allocations in other cases. This relatively poor diversification within the financial portfolio is not a defect of our approach, but a consequence of the hedging property of the financial portfolio. Only a concentrated (relative to the locally efficient portfolio) allocation exhibits the negative oil sensitivity necessary for hedging large oil risks.

In Table 6, we report the variance reduction and the associated added certainty equivalent in the presence of short-sale constraints. In panel A, we report the average statistics. As before, we focus on the case $\omega = 50$ per cent, which is in the second column block. In general, the patterns match those in Table 4: partitioning the investment universe to finer degrees from regions to industries increases the potential to reduce oil risk. The magnitude of the risk reduction is, however, necessarily different. For the country stratification, the variance reduction is 0.0020 (versus 0.0114 in the unconstrained case) and the added certainty equivalent return is 0.31 per cent (versus 0.57 per cent). Thus, the advantage of switching from locally to globally efficient portfolios is considerably less if short sales are not

Figure 6 Globally efficient portfolios. Country stratification. $\omega = 50$ per cent under short-sale constraints

Figure 7 Locally efficient portfolios. Country stratification and $\omega = 50$ per cent under short-sale constraints

possible. Similar comparisons for the other partitions of the investment universe show that imposing short-sale constraints roughly cuts in half the certainty equivalent return. The advantage, however, is still economically significant. Exploiting the wider dispersion

of oil sensitivities of the industry stratification yields an average gain in certainty equivalent returns of 0.48 per cent.

Scanning along the rows of Table 6, panel A shows that when short sales are prohibited the certainty equivalent return is hump

Table 6 Locally versus globally efficient portfolios under short sale constraints

Panel A Average across efficient portfolios						
Universe	$\omega = 25\%$			$\omega = 50\%$		
	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)
Regional	0.0001	0.83	0.02	0.0003	0.80	0.02
Country	0.0009	6.19	0.13	0.0020	4.63	0.31
Sector	0.0010	5.54	0.14	0.0021	4.37	0.31
Industry	0.0016	9.06	0.24	0.0032	6.69	0.48
Universe	$\omega = 75\%$			$\omega = 90\%$		
	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)
Regional	0.0001	0.83	0.02	0.0002	0.18	0.04
Country	0.0022	2.29	0.33	0.0012	0.87	0.18
Sector	0.0020	2.05	0.30	0.0011	0.75	0.16
Industry	0.0029	2.96	0.44	0.0015	1.09	0.23
Panel B Efficient portfolio for expected return of 6%						
Universe	$\omega = 25\%$			$\omega = 50\%$		
	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)
Regional	0.0001	0.96	0.02	0.0004	1.01	0.06
Country	0.0011	8.50	0.17	0.0026	6.15	0.40
Sector	0.0014	9.69	0.21	0.0030	6.78	0.46
Industry	0.0025	17.25	0.37	0.0046	10.26	0.69
Universe	$\omega = 75\%$			$\omega = 90\%$		
	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)	$\Delta \text{Var}(r)$ (%)	$\Delta \text{Var}(r)$ (%)	ΔCER (%)
Regional	0.0005	0.52	0.07	0.0003	0.19	0.04
Country	0.0025	2.63	0.38	0.0013	0.97	0.20
Sector	0.0034	3.45	0.50	0.0018	1.32	0.27
Industry	0.0042	4.30	0.63	0.0023	1.62	0.34

The table contains the same statistics as Table 5 for portfolios constrained to have nonnegative weights.

shaped in w , that is, the value of oil increases relative to total wealth. For example, the certainty equivalent return increases from 0.13 per cent for $w = 25$ per cent to around 0.3 per cent for $w = 50$ per cent and $w = 75$ per cent, and drops again to 0.18 per cent for $w = 90$ per cent in the case of the country stratification of the universe. In contrast, if short sales are allowed, the certainty equivalent return increases approximately linearly in squared w 's. If financial portfolios cannot be levered, the potential for reducing oil risk is lowest when either oil or financial assets are a small part of aggregate wealth. If aggregate wealth is mostly in oil, the financial

portfolio is simply too small to have a substantial impact on the volatility of the aggregate portfolio. Conversely, if aggregate wealth is mostly in financial assets, there is not much gained from reducing oil risk, which does not contribute much to total risk.

Figure 5 shows that, contrary to the unconstrained case, the variance reduction and the added certainty equivalent return are not stable across the frontier, but increasing and concave. Therefore, we also present, in panel B of Table 6, the variance reduction and additional certainty equivalent return for an expected return of 6 per cent on the

financial portfolio. For example, the certainty equivalent return for the country stratification and $\omega = 50$ per cent is 0.40 per cent for the financial portfolio, with 6 per cent expected return. In comparison, the average across all portfolios (from panel A) is only 0.31 per cent. Browsing the table, this is true for all ω 's and all stratifications of the universe. Since the portfolio roughly corresponds to the median, this result is a consequence of Jensen's inequality.

Conclusion

We show that taking into account the risk stemming from oil, as an example of a nontradable asset, can have substantial consequences for the risk of aggregate wealth and the efficient allocation of financial assets. Standard financial assets, depending on the sensitivity to oil price risk, partially act as hedge instruments. For relatively coarse partitions along countries or industries of a global investment universe, we find significant differences in oil sensitivity. For these partitions, we achieve reductions in variance of aggregate wealth between 20 and 50 per cent, depending on the partition and the relative importance of the oil asset, compared to standard efficient portfolios. These risk figures translate into gains in certainty equivalent returns of 43 basis points to 10.56 per cent pa. If short sales of financial assets are not permitted, the average reduction in variance is between 1 and 10 per cent. The corresponding average gain in certainty equivalent returns is 13 basis points to 48 basis points.

Notes

1. A long time horizon does not imply low risk aversion. This is one of the most common fallacies made in asset management and usually rests with the focus on quantile-based risk management.
2. [Http://www.odin.dep.no/fin/engelsk/p10001617/p10002780/indexbna.html](http://www.odin.dep.no/fin/engelsk/p10001617/p10002780/indexbna.html). Further information regarding the aims and policies of the Fund is in the Annual Reports, Kjaer (2001), and Norges Bank (2002).

3. Kjaer (2001) reports that Norges Bank, when advising on whether to invest in equities, provided standard deviation of returns and shortfall probabilities as the relevant risk measures.
4. Data on Government Revenues are from the English Summary of Norway's National Budget for 2003. Data on remaining resources are from the Norwegian Petroleum Directorate's 'The Petroleum Resources on the Norwegian Continental Shelf as at 31st December, 2001'.
5. This overstates the government's claim to revenues from production, since licences generate cashflows. The licencing fee is then a capitalised claim to future oil revenues.
6. Fasano (2000) and Melby (2002) contain overviews.
7. This is peculiar to that universe. The other, finer partitions of equity yield hedge portfolios tilted towards equity.
8. A proof is in Ingersoll (1987), p. 84, whose notation we adopt.
9. We have pointed out above that it is not necessary to establish expected oil price changes since they do not affect efficient portfolios.

References

- Fasano, U. (2000) 'Review of the Experience with Stabilisation and Savings Funds in Selected Countries', Working Paper, IMF.
- Ingersoll Jr, J. E. (1987) *Theory of Financial Decision Making*, Rowman and Littlefield, Lanham MD.
- Kjaer, K. N. (2001) A national strategy for investing resource wealth, Speech at the BSI Gamma Foundation Conference on Global Asset Management Long Term Asset Management.
- Melby, E. D. K. (2002) 'A Global Overview of Oil Funds', Presentation for the IGAF Symposium.
- Norges Bank (2002) An appraisal of the regional weighting of the Petroleum Fund, Letter to the Ministry of Finance on 11th April, 2002.

Appendix A Globally efficient portfolios

The standard program for calculating efficient portfolios of purely financial assets, which we call locally efficient portfolios, minimises $\mathbf{w}^T \Sigma \mathbf{w}$ subject to the constraints that the portfolio achieves a target expected return $\mathbf{w}^T \bar{\mathbf{z}} = \mu$ and the budget constraint $\mathbf{w}^T \mathbf{1} = 1$. Using familiar convex programming techniques, the optimal portfolio weights for target expected return μ are⁸

$$\mathbf{w}_L(\mu) = \Sigma^{-1} \mathbf{1} \frac{1}{\Delta} [\mu(\mathbf{1}^T \Sigma^{-1} \mathbf{1}) - (\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})] + \Sigma^{-1} \mathbf{1} \frac{1}{\Delta} [\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}} - \mu(\mathbf{1}^T \Sigma^{-1} \bar{\mathbf{z}})]$$

where $\Delta = (\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}}) - (\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})^2$.

The corresponding problem of calculating efficient portfolios, including the oil asset of fixed supply, which we call globally efficient portfolios, minimises $Var(r) = \omega^2 \sigma_o^2 + (1-\omega)^2 \mathbf{w}^T \Sigma \mathbf{w} + 2\omega(1-\omega)\sigma_o^2 \mathbf{w}^T \mathbf{b}$ subject to the constraints $\mathbf{1}^T \mathbf{w} = 1$ and $\bar{\mathbf{z}}^T \mathbf{w} = \mu$. The Lagrangian for the problem is

$$L = (1-\omega) \mathbf{w}^T \Sigma \mathbf{w} + 2\omega \sigma_o^2 \mathbf{w}^T \mathbf{b} + \lambda(1 - \mathbf{1}^T \mathbf{w}) + \gamma(\mu - \bar{\mathbf{z}}^T \mathbf{w})$$

The first order conditions are

$$2(1-\omega) \Sigma \mathbf{w}_G(\mu) + 2\omega \sigma_o^2 \mathbf{b} - \lambda \mathbf{1} - \gamma \bar{\mathbf{z}} = 0$$

and, of course, the constraints. Solving for the portfolio weights

$$\mathbf{w}_G(\mu) = \frac{1}{2(1-\omega)} \Sigma^{-1} (\lambda \mathbf{1} + \gamma \bar{\mathbf{z}} - 2\omega \sigma_o^2 \mathbf{b})$$

Substituting into the constraint equations, solving for the constraints, and resubstituting these into the optimal portfolio weight

$$\begin{aligned} \mathbf{w}_G(\mu) &= \frac{1}{2(1-\omega)} \left\{ \frac{2(1-\omega)}{\Delta} [\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}} - \mu \mathbf{1}^T \Sigma^{-1} \bar{\mathbf{z}}] \right. \\ &\quad + \frac{2\omega \sigma_o^2}{\Delta} [(\mathbf{1}^T \Sigma^{-1} \mathbf{b})(\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}}) \\ &\quad \left. - (\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{b})] \right\} \Sigma^{-1} \mathbf{1} \\ &\quad + \frac{1}{2(1-\omega)} \left\{ \frac{2(1-\omega)}{\Delta} [\mu \mathbf{1}^T \Sigma^{-1} \mathbf{1} - \bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1}] \right. \\ &\quad + \frac{2\omega \sigma_o^2}{\Delta} [(\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{b}) \\ &\quad \left. - (\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})(\mathbf{1}^T \Sigma^{-1} \mathbf{b})] \right\} \Sigma^{-1} \bar{\mathbf{z}} \\ &\quad - \frac{2\omega \sigma_o^2}{2(1-\omega)} \Sigma^{-1} \mathbf{b} \end{aligned}$$

where $\Delta = (\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{1}) - (\mathbf{1}^T \Sigma^{-1} \bar{\mathbf{z}})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})$. Rearranging

$$\begin{aligned} \mathbf{w}_G(\mu) &= \frac{1}{\Delta} [\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}} - \mu \mathbf{1}^T \Sigma^{-1} \bar{\mathbf{z}}] \Sigma^{-1} \mathbf{1} \\ &\quad + \frac{1}{\Delta} [\mu \mathbf{1}^T \Sigma^{-1} \mathbf{1} - \bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1}] \Sigma^{-1} \bar{\mathbf{z}} \\ &\quad + \frac{\omega \sigma_o^2}{(1-\omega)} \left\{ \frac{1}{\Delta} [(\mathbf{1}^T \Sigma^{-1} \mathbf{b})(\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}}) \right. \\ &\quad - (\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{b})] \Sigma^{-1} \mathbf{1} \\ &\quad + \frac{1}{\Delta} [(\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{b}) \\ &\quad \left. - (\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})(\mathbf{1}^T \Sigma^{-1} \mathbf{b})] \Sigma^{-1} \bar{\mathbf{z}} - \Sigma^{-1} \mathbf{b} \right\} \end{aligned}$$

But the first two terms are just $\mathbf{w}_L(\mu)$, the locally efficient portfolio for target expected return μ . Thus,

$$\begin{aligned} \mathbf{w}_G(\mu) &= \mathbf{w}_L(\mu) + \frac{\omega}{(1-\omega)} \\ &\quad \times \sigma_o^2 \left\{ \frac{1}{\Delta} [(\mathbf{1}^T \Sigma^{-1} \mathbf{b})(\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}}) \right. \\ &\quad - (\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{b})] \Sigma^{-1} \mathbf{1} \\ &\quad + \frac{1}{\Delta} [(\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{b}) \\ &\quad \left. - (\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})(\mathbf{1}^T \Sigma^{-1} \mathbf{b})] \Sigma^{-1} \bar{\mathbf{z}} - \Sigma^{-1} \mathbf{b} \right\} \end{aligned}$$

We further claim that the hedge portfolio $\mathbf{w}_H$ is a zero-net investment portfolio, that is, $0 = \mathbf{w}_H^T \mathbf{1}$. The net position of the hedge portfolio is

$$\begin{aligned} \mathbf{1}^T \mathbf{w}_H &= \sigma_o^2 \mathbf{1}^T \left\{ \frac{1}{\Delta} [(\mathbf{1}^T \Sigma^{-1} \mathbf{b})(\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}}) \right. \\ &\quad - (\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{b})] \Sigma^{-1} \mathbf{1} \\ &\quad + \frac{1}{\Delta} [(\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{b}) \\ &\quad \left. - (\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{1})(\mathbf{1}^T \Sigma^{-1} \mathbf{b})] \Sigma^{-1} \bar{\mathbf{z}} - \Sigma^{-1} \mathbf{b} \right\} \end{aligned}$$

Expanding,

$$\begin{aligned} \mathbf{1}^T \mathbf{w}_H &= \frac{\sigma_o^2}{\Delta} \{ (\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\mathbf{1}^T \Sigma^{-1} \mathbf{b}) \\ &\quad - (\mathbf{1}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{b}) \\ &\quad + (\mathbf{1}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\bar{\mathbf{z}}^T \Sigma^{-1} \mathbf{b}) \\ &\quad - (\mathbf{1}^T \Sigma^{-1} \bar{\mathbf{z}})^2 (\mathbf{1}^T \Sigma^{-1} \mathbf{b}) \\ &\quad - (\bar{\mathbf{z}}^T \Sigma^{-1} \bar{\mathbf{z}})(\mathbf{1}^T \Sigma^{-1} \mathbf{1})(\mathbf{1}^T \Sigma^{-1} \mathbf{b}) \\ &\quad + (\mathbf{1}^T \Sigma^{-1} \bar{\mathbf{z}})^2 (\mathbf{1}^T \Sigma^{-1} \mathbf{b}) \} \\ &= \frac{\sigma_o^2}{\Delta} 0 \end{aligned}$$

Finally, we claim that $\mathbf{w}_H$ has zero expected return, that is, $\bar{\mathbf{z}}^T \mathbf{w}_H = 0$. From the optimisation program,

$$\begin{aligned}\mu &= \bar{\mathbf{z}}^T \mathbf{w}_G(\mu) = \bar{\mathbf{z}}^T \mathbf{w}_L(\mu) + \frac{\omega}{(1-\omega)} \bar{\mathbf{z}}^T \mathbf{w}_H \\ &= \mu + \frac{\omega}{(1-\omega)} \bar{\mathbf{z}}^T \mathbf{w}_H\end{aligned}$$

Thus, $0 = \bar{\mathbf{z}}^T \mathbf{w}_H$ since $\omega/(1-\omega) > 0$.

Appendix B Data

Oil price data

We employ two different series of monthly oil price data. The first series is the US Bureau of Labour Statistics' (BLS) Producer Price Index, Series WPU, Crude Petroleum — Domestic, in US\$, which is available on a monthly basis since 1947. The second series is Dated Brent Current Month FoB from ICIS' Oil Report, distributed by Thomson Financial Datastream and available since 1983. The BLS series has distinct disadvantages stemming from how these data are collected, but has the advantage of a long history. The time-series properties of the BLS series serve mainly as justification for exclusively using the Brent series in the remainder of the analysis.

BLS oil prices

The price index is calculated from prices sampled from a wide range of US producers as of the Tuesday of the week, in which the 13th calendar day of the month falls. Thus,

depending on the particular month under consideration, prices are sampled between the 9th and the 15th of the month. This implies that data on oil prices and financial assets are not sampled contemporaneously.

While the BLS data are available on a monthly basis from 1947 onwards, oil price dynamics change dramatically with the oil crisis in 1973. This is apparent from Figure B1, which shows the monthly index level since 1947, and Table B1, which reports descriptive statistics. Evidently, until 1973 oil prices are extremely stable. Prices increase steadily between 1973 and 1981, after which prices decrease steadily until 1985. Only after 1986 the oil price series exhibits the typical patterns of price series, high volatility and low predictability. The figures do not reveal any change in appearance after 1986 and there do not appear to be any structural breaks in the time series. In particular, volatility appears to be similar over the last 15 years.

Overall, there appear to be three historical periods during which different statistical processes describe BLS oil prices: a long first period between 1947 and 1973, when oil prices are very stable and increase barely noticeably, and a second period between 1973 and 1985 during which oil prices rise and volatility reaches modest levels. In the most recent period, no price trend is discernible and volatility is high. Table B1, which contains the sample average of the logarithmic changes and volatilities, both on an annual basis, documents these patterns as

Figure B1 BLS oil price index. Time period 1947–2001

Table B1 Descriptive statistics for oil price changes

	1947–2001	1947–1972	1973–1985	1986–2001
<i>Mean</i>				
Brent				–1.79%
BLS	3.33%	2.93%	13.05%	–3.93%
<i>Volatility</i>				
Brent				41.5%
BLS	20.57%	4.87%	11.48%	36.1%
<i>Autocorrelation</i>				
Brent				–0.03
BLS	0.18	0.36	0.41	0.15
Correlation of Brent(t) with				
BLS(t–1)	BLS(t)		BLS(t+1)	BLS(t+2)
–0.07	0.43		0.61	0.13

Figure B2 Brent and BLS oil price index. Time period 1983–2001

well. Over the entire period, BLS oil prices increase on average by 3.33 per cent pa, 2.93 per cent (13.05 per cent) (3.93 per cent) during the first (second) (third) period. Volatility is 20.57 per cent over the whole period, and 4.87, 11.48, and 36.07 per cent in the three subperiods. Table B1 also reports the serial correlation in the logarithmic changes. Depending on the time period, the time series correlation varies between 0.41 and 0.15, reflecting considerable persistence in the price changes. Serial correlation in the BLS price series is not unusual. However, both the correlation being positive and the magnitude are surprising..

Dated Brent prices

The DS Brent Crude prices are based on averages of major spot transactions with delivery up to one month, which ICIS samples daily. The monthly series is constructed from the prices reported for the last day of the month. Thus, the timing in the two oil price series is slightly different. Another difference arises from the fact that the two series are based on products of different quality. While the BLS price refers to a generic product ‘crude oil’, the ICIS series refers to a specific commodity of a very standardised nature with respect to quality, point of delivery, etc.

Figure B2 contains the ICIS Brent Price index level, using the same base date as the BLS Oil Price index and the BLS index for the same period. Obviously, the two series are very similar but not identical. Close inspection shows that the BLS index trails the ICIS index slightly. Since BLS prices are, on average, as of the 13th of the month, while ICIS prices are as of the last day of the month, this is not surprising. The ICIS index reflects changes in the second half of the month in the current observations, while the BLS index incorporates the same price changes only in the observation for the following month.

Table B1 also contains descriptive statistics for the ICIS Brent prices. On average, the price declines by 1.79 per cent pa between 1986 and 2001. The volatility is 41.53 per cent pa over the same period. The serial correlation is slightly negative at 0.03, which is consistent with a mild bid ask bounce in a liquid dealer market. The absolute magnitude of the series is negligible for all practical purposes, and is unlikely to cause statistical problems in the analyses below. Overall, the BLS and ICIS series appear, except for the time lag in the sampling procedure, very similar. Thus, it seems appropriate to treat them as close substitutes for empirical purposes.

Financial assets

As financial assets, we employ relatively broad market indices. First, the limited amount of available time-series data with which to estimate the correlation structure imposes restrictions on the number of assets that can be analysed. Secondly, and relatedly, increasingly fine partitioning of the investable universe increases the estimation error for the individual asset. Within an asset allocation exercise, this leads to maximisation of estimation error instead of well-diversified portfolios. Thirdly, allocating among market indices facilitates implementation of the investment policy within a traditional asset

management process. In addition, we require eligible assets to be sufficiently liquid, ensuring that the investment policy is implementable at reasonable transaction costs. For this reason, we restrict the analysis for debt issues to the major economic regions, and the analysis for equity to the MSCI World constituents.

We distinguish financial assets along three dimensions: asset class (ie equity versus sovereign debt), geography (economic regions or countries), and industry (economic sectors or industries). Other types of financial assets, such as corporate debt, derivatives, etc are considered replicable by the standard asset classes. Thus, we do not consider them explicitly.

We also focus on currency-hedged returns. Because there often are no hedged indices with a sufficient number of time-series observations, we use returns in local currency as a proxy for hedged returns. This is equivalent to assuming that perfect currency hedges are available.

Expected returns

To determine efficient portfolios we need to make assumptions on the expected returns on debt and excess returns on equity.⁹ For debt, we use the yield to maturity as quoted in the market. We assume that there is no credit risk in sovereign debt, and therefore, expected returns are equal for all issuers. In other words, differences in market yields for similar maturities are offset by the corresponding differential in forward currency rates. Thus, yields denominated in common currency are the same for the same maturity. Since the indices are composed of instruments of many different maturities, we compute forward rates based on duration. Using these forward rates, we translate local currency market yields into common, US dollar market yields. Thus, expected returns in the common currency are equal for all issuers except for differences in duration.

For equities, we use returns implied by the holdings of the value-weighted global portfolio proxied by the MSCI World and the historical covariance matrix. In other words, we assume that the CAPM describes expected returns on equities. In addition, we assume that the risk premium is 3 per cent pa of course, we can easily accommodate alternative assumptions on expected returns.

We estimate the covariance matrix from available historical data on the asset classes in the usual way. Since the estimates are standard, we do not report them.

Sovereign debt

Since interest rates on government debt are generally highly correlated across developed countries, we distinguish only three major

economic blocks: Europe, Japan, and North America. For European and Japanese sovereign debt, we employ the Salomon Smith Barney World Government subindices as proxies, and the Merrill Lynch Treasury Master in the case of the US.

Equity

Regional equity distinguishes between the three major economic blocks: North America, Europe, and Pacific. Country equity splits global equity into the developed markets using the MSCI definitions. Sector and Industry equity separates equity according to industrial sectors and industries following the Global Industry Classification Standards (GICS) classification used by MSCI.