Network for Studies on Pensions, Aging and Retirement

Bert Scholtens

Laura Spierdijk

Does Money Grow on Trees? The Diversification Properties of US Timberland Investments

Discussion Paper 03/2009 - 017

March 23, 2009

Does Money Grow on Trees?
The Diversification Properties of US Timberland Investments

Bert Scholtens[*] Laura Spierdijk[†]

March 23, 2009

[*]University of Groningen, Faculty of Economics and Business, Department of Finance and the Centre for International Banking, Insurance and Finance (CIBIF), P.O. Box 800, 9700 AV Groningen, The Netherlands. Phone +31 50 363 7064. E-mail L.J.R.Scholtens@rug.nl

[†]Corresponding author. University of Groningen, Faculty of Economics and Business, Department of Economics & Econometrics and CIBIF, P.O. Box 800, 9700 AV Groningen, The Netherlands. Phone: +31 50 363 5929. E-mail: L.Spierdijk@rug.nl. The authors would like to thank Jack Lutz and two anonymous referees for useful suggestions. The usual disclaimer applies.

Abstract

This paper quantifies the diversification potential of timberland investments in a mean-variance framework. The starting point is a broad set of benchmark assets represented by various indexes. Including publicly traded timberland investments from the US and Canada in the portfolio does not significantly increase mean-variance efficiency. At first sight, US private equity timberland seems to improve the mean-variance frontier, even if the portfolio already contains a forestry and paper equity index. Adding privately held timberland to the investment set increases the risk-adjusted excess return on the tangency portfolio with about 10 bp per quarter. However, after removing the appraisal smoothing bias from the raw timberland data, there is much less evidence that private equity timberland investments increase mean-variance efficiency. Even under mild assumptions regarding the appraisal smoothing bias, the inclusion of the unsmoothed Timberland Index increases the risk-adjusted excess return on the tangency portfolio with only 1 bp per quarter.

Keywords: timberland investments, mean-variance spanning, portfolio management

JEL Classification: G10, G12, L73

1 Introduction

Investments in timberland have become increasingly popular with institutional investors both in the Unites States and elsewhere in the world. According to the UGA Center for Forest Business, the global timberland market value in 2006 was about 400 billion dollar, of which 230 billion in the United States. Within the US, private landowners' timberland had a value of $ 160 billion, forest products companies owned 52 billion and institutional investors possessed 14 billion.[1] Timberland investment returns are driven by four main factors: biological growth, timber prices, land appreciation, and inflation (Healey et al. 2005). The attractiveness of timberland for institutional investors is often explained by its low correlation with more traditional assets (such as stocks and bonds), which would make it a suitable diversification instrument. See e.g. Redmond and Cubbage (1988) and Sun and Zhang (2001) who estimate CAPM models and find negative beta values for various timberland investments.

The CAPM framework is the conventional approach to assess the diversification properties of timberland investments from the investor's perspective. Studies based on CAPM focus on excess returns and on the risk level relative to the market portfolio. The negative beta generally found in the literature suggests that there is some potential for improving the risk and return characteristics of a portfolio by including timberland.

This paper adopts a different approach by explicitly quantifying how much the risk-adjusted excess return will increase when timberland is added to an institutional portfolio. For this purpose, we analyze the diversification potential of timberland investments in a mean-variance framework. We apply the mean-variance spanning and intersection tests of Huberman and Kandel (1987) and Kan and Zhou (2008) to assess to what extent the mean-variance frontier improves by including timberland in the portfolio. In particular, we investigate how much the risk-adjusted excess return on the tangency portfolio (the portfolio with highest possible risk-adjusted excess return) will increase when timberland is included. Existing studies on timberland performance generally use a simple proxy for the market portfolio such as the S&P500. The resulting negative beta's tell us that timberland investments have low correlations with such an index. By contrast, we consider different investment sets including US and global stock, bond, real estate, and commodity indexes. Moreover, we focus on both private equity and publicly traded timberland investments.

Our results show that adding publicly traded timberland investments from the US and Canada to the portfolio does not significantly improve the mean-variance frontier. At first sight, US private equity timberland seems to increase mean-variance efficiency, even if the investment set already contains a forestry and paper equity index. Adding timberland to the investment set increases the risk-adjusted excess return on the tangency portfolio with about 10 bp per quarter. However, after removing the appraisal smoothing bias from the raw timberland data, there is much less evidence that privately held timberland investments increase mean-variance efficiency. Even under mild assumptions regarding the appraisal smoothing bias, the inclusion of the unsmoothed Timberland Index increases the risk-adjusted excess return on the tangency portfolio with only 1 bp per quarter.

The setup of the remainder of this paper is as follows. Section 2 describes the data and provides some sample statistics. The mean-variance framework and the tests for spanning and intersection are explained in Section 3. Section 4 discusses the empirical results and Section 5 investigates the robustness of the analysis. Finally, Section 6 concludes.

2 The data

This section describes the data used for the empirical part of this paper and provides some sample statistics.

2.1 Timberland investments

Our goal is to assess how the inclusion of timberland investments in an institutional investor's investment set affects the mean-variance efficient portfolio of this investor. Hence, we have to set clear how we represent the institutional portfolio and the timberland investment. With respect to the institutional portfolio, we construct a set of benchmark assets, covering different investment classes in various countries. The asset classes are represented by one or more total return indexes. International diversification is ensured by including both US and global indexes. The assets have been chosen to reflect the elements of a well-diversified portfolio and include US and global (small and large cap) stocks, bonds, real estate, and commodity indexes, but not with the goal to mimick an existing portfolio. The benchmark assets are listed in the first column of Table 1, with a short index description and the data source in the second and third column, respectively.[2]

Insert Table 1 about here.

We distinguish between private equity investments in timberland for institutional investors (and wealthy individuals) and publicly traded timberland investments available to any investor. Regarding private equity investments in timberland, there is little choice with respect to available data. As far as we know, there are only two relevant indexes: the Timberland Performance Index and the National Council of Real Estate Investment Fiduciaries (NCREIF) Timberland Index. The former ended in 1999. The latter is a property-based index reporting returns for three regions of the United States: the South, Northeast and Pacific Northwest. The returns on the NCREIF Timberland Index are determined by the income and appreciation returns on the timberland managed by eight Timberland Investment Management Organizations (TIMO's) in the US.[3] The income returns arise from log and stumpage sales, whereas the appreciation returns results from timber and land appreciation. TIMO's are privately owned companies who basically work as timberland brokers for institutional investors. They try to earn cash and capital return for their investors. Also Timberland Real Estate Investment Trusts (REIT's) own and explore timberland properties. They are listed on a stock exchange, but are not considered corporations and therefore do not pay corporate income tax. They have to pay out 90% of their profit as dividend to investors. In contrast to TIMO's, REIT's do not contribute to the NCREIF Timberland Index. Apart from timberland REIT's, there are many publicly traded companies in the paper, lumber and wood production industry that own or explore timberland and can therefore be regarded as timberland investments.

Privately and publicly held timberland may differ substantially in terms of risk and return profiles. Apart from the appraisal smoothing bias (which we discuss in Section 2.3), several other factors may contribute to these differences. Publicly traded timberland funds are leveraged, while privately owned timberland is not. Leverage generally increases average returns and volatility. Moreover, privately and publicly held timberland investments are likely to differ with respect to the asset composition. For example, timberland located in the South, Northeast and Pacific Northwest of the US is overrepresented in the NCREIF Timberland Index. Another factor of potential importance is liquidity. Due to the underlying investment vehicle and the illiquidity of timberland itself, private equity investments in timberland are relatively illiquid. It is usually not possible to sell a sizeable timberland property instantaneously. Publicly traded timberland investments are much more liquid, since investors can sell their shares at any time. Investors generally value the liquidity of an investment. However, privately held timberland is bought by institutional investors with a long-term investment horizon for whom illiquid assets are not a problem. It is therefore unclear whether the illiquid nature of privately owned timberland negatively affects its risk-return profile. Also efficiency may play a role. Publicly traded timberland may benefit from economies of scale and greater reporting transparency. Finally, the fee structure could play a role. The TIMO's included in the NCREIF Timberland Index basically act as timberland brokers. They charge their institutional clients a fee which is not corrected for in the reported timberland returns. For a more detailed analysis explaining performance differences between privately and publicly held assets, we refer to Riddiough et al. (2005) who compare privately and publicly owned commercial real estate.

2.2 Description of the data

The NCREIF Timberland Index index has its limitations (see Section 2.3), but we have few other possibilities to represent private equity timberland investments as an asset class. Although institutional investors will presumably also invest in timberland outside the US, little or no data is at hand for such investments. For this reason we focus on 'institutional' timberland in the US as represented by the quarterly NCREIF Timberland Index, which has been used in other studies as well (Sun and Zhang, 2001; Healey et al., 2005). We focus on the longest sample period for which we have returns on both the Timberland Index and our virtual institutional portfolio. The resulting time span runs from the first quarter of 1994 until the third quarter of 2007 and comprises 55 quarterly observations. As opposed to institutional timberland represented by the TIMO-based NCREIF Timberland Index, we also construct three (value-weighted) indexes of publicly traded timberland.[4] The first consists of the three publicly traded timberland REIT's (Rayonier, Plum Creek and Potlatch). Our second index is broader and covers the largest companies in the paper and paper products industry in the US and Canada. The third index consists of the major companies in the lumber and wood production industry in these countries.[5] For the three indexes of publicly held timberland the sample period runs from the fourth quarter of 1994 until the third quarter of 2007.[6]

2.3 The Timberland Index and the appraisal smoothing bias

Timberland properties are not traded frequently enough to construct a transaction-based index. For this reason the NCREIF Timberland Index is based on appraisal values. Timberland properties are appraised on the basis of recent transactions of comparable properties and these appraisal values are used to construct the Timberland Index. As a consequence, the Timberland Index may suffer from an appraisal smoothing bias and may contain certain inertia. A major concern is that the volatility of the observed index returns is too low compared to the true (unobserved) index. This would seriously distort a mean-variance analysis, resulting in a too optimistic picture of the diversification potential of timberland. We refer to Appendix A.1 for a more detailed exposition of the potential problems caused by the use of appraisal values. To avoid analyzing Timberland Index returns with too low a volatility compared to the true returns, we consider an 'unsmoothed' version of the index.[7] We use the unsmoothing approach introduced by Fisher et al. (1994); see Appendix A.2. The unsmoothed index returns have the same mean as the raw index returns, but a substantially higher volatility. As explained in more detail in Appendix A.2, the unsmoothing approach involves the problem that the true volatility of the unsmoothed index is unknown. Instead of choosing an arbitrary value for this volatility, we proceed in a different way. For a range of possible volatilities, we obtain the unsmoothed Timberland Index. Subsequently, we apply the tests for mean-variance efficiency (to be discussed in the next section) to the resulting series. To save space, we do not report the results for each of the unsmoothed series. Instead we focus on one version of the unsmoothed Timberland Index. We consider the highest volatility for which including the unsmoothed Timberland Index in the investment set significantly improves the mean-variance efficiency of the tangency portfolio. It turns out that the unsmoothed index significantly increases the mean-variance efficiency of the tangency portfolio for volatilities up to 5.3%. This is the version of the unsmoothed Timberland Index for which we will report results in the remainder of this paper. Since the mean of the unsmoothed series is not affected by the unsmoothing procedure, the unsmoothed Timberland Index will become less attractive in terms of mean-variance when its volatility is increased. Hence, the reported unsmoothed index is based on a relatively favorable choice of the volatility. We emphasize that we do not view the unsmoothed index included in Table 1 as the 'true' index. Instead, we consider it a tool for sensitivity analysis with respect to the volatility of the unsmoothed Timberland Index. We will turn back to this issue in Section 5.3.

2.4 Preliminary data analysis

The fourth and fifth column of Table 1 provide quarterly means and volatilities for all asset returns. The last two columns report the correlations between the returns on the benchmark assets and the returns on the raw NCREIF Timberland Index. The Dow Jones Canada Index generates the highest average quarterly returns during the sample period and the Dow Jones US Technology Index is the most volatile. Some indexes are mean-variance inefficient, in the sense that there exists at least one other index with a higher average return and lower volatility. For instance, the MSCI Far East Index is inefficient compared to the MSCI World. The means of the raw and unsmoothed Timberland Index returns are the same, but the volatility of the unsmoothed Timberland Index (5.3%) is almost twice as large as the volatility of the raw index (2.7%). The raw and the unsmoothed Timberland Index have about the same correlation with the other assets. Table 2 presents the full correlation matrix corresponding to the benchmark assets, the (raw and unsmoothed) NCREIF Timberland Index, and our three indexes of publicly traded timberland. The latter timberland indexes have a much higher volatility than the NCREIF Timberland Index. Moreover, their returns show weakly positive or even negative correlation with the return on the Timberland Index. Also the Dow Jones Forestry & Paper Indexes in our set of benchmark assets have a relatively high volatility and are negatively correlated with the Timberland Index.

Insert Table 2 about here.

As a first step, we estimate a simple CAPM model for the raw Timberland Index with the S&P500 as the market index and the quarterly return on a 1-month T-Bill as the risk-free rate. This results in a beta equal to 0.04 (p-value 0.43) and a risk-adjusted excess return (alpha) of 1.4% (with p-value 0.00). For the unsmoothed Timberland Index we find virtually the same results and also for the three indexes of publicly traded timberland we find beta's close to zero; see Table 3. This suggests that adding timberland has some potential for improving the risk and return characteristics of a portfolio. However, the key question is whether the beta is indeed low enough. The spanning framework presented in the next section can be used to answer this question in a formal way and to quantify the added value of timberland investments.

Insert Table 3 about here.

3 Testing for mean-variance spanning and intersection

Given a collection of benchmark assets, portfolio weights can be chosen in such a way that the resulting portfolio is mean-variance efficient. For a given variance, a mean-variance efficient portfolio has maximum expected return. Stated differently, it has the lowest variance at a given level of return for all possible portfolios. The weights associated with a mean-variance efficient portfolio depend on the degree of risk aversion of the investor.

The key question is whether the inclusion of another asset class in the set of benchmark assets improves the mean-variance efficient portfolio. More specifically, we ask: 'Does the inclusion of timberland to the investment set improve mean-variance efficiency?' This question will be answered using the framework of mean-variance spanning and intersection. DeRoon and Nijman (2001) present a survey of the various methods used to test whether the mean-variance frontier of a set of benchmark assets spans or intersects the frontier of a larger set of assets. Two situations can arise. First, if there is only a single value of the risk aversion parameter for which mean-variance investors cannot improve upon their mean-variance efficient portfolio by including the additional assets in their investment set, the mean-variance frontiers of the benchmark assets and the extended set of assets intersect. With intersection, the mean-variance frontier of the benchmark assets and the frontier of the benchmark assets and the additional asset have precisely one point in common. Second, if there is no value of the risk aversion parameter for which a mean-variance investor can improve her mean-variance efficient portfolio, the mean-variance frontiers of the benchmark and the extended set of assets coincide. This is called spanning. Spanning implies that the mean-variance frontier of the benchmark assets coincides with that of the benchmark assets and the additional asset. In that case, no mean-variance investor is better off by investing in the additional asset, regardless of the zero-beta rate.

We use the regression framework of Huberman and Kandel (1987) to test for spanning and intersection. We denote the returns on the K benchmark assets by the K-dimensional column vector R_t and the returns on the additional asset by the scalar r_t. We consider the regression of r_t on R_t, i.e.

$$r_t = \alpha + \beta R_t + \varepsilon_t, \tag{1}$$

where α is the intercept and β represents a K-dimensional row vector of coefficients. In terms of parameter restrictions the hypothesis of intersection is stated as

$$\alpha - \eta(1 - \sum_{k=1}^{K} \beta_k) = 0, \tag{2}$$

where η equals the zero-beta rate (which we assume to be known). The hypothesis of spanning implies that restriction (2) hold for all values of η and reduces to

$$\alpha = 0, \quad \sum_{k=1}^{K} \beta_k = 1. \tag{3}$$

As shown by Kan and Zhou (2008), testing $\alpha = 0$ boils down to testing whether the tangency portfolio (i.e. the portfolio with the highest possible Sharpe ratio) has zero weights in the N test assets. Furthermore, the restriction $\sum_{k=1}^{K} \beta_k = 1$ is a test of whether the global minimum-variance (GMV) portfolio has zero weights in the test assets. From the two-fund separation theorem (Tobin, 1958) we know that if there are two distinct minimum-variance portfolios that have zero weight in the N test assets, every portfolio on the minimum-variance frontier of the $N+K$ assets will also have zero weights in the N test assets. This explains why the two restrictions in Equation (3) are sufficient to test for mean-variance spanning. However, not all portfolios on the mean-variance frontier are equally relevant. According to the CAPM model, each investor will hold a combination of the risk-free asset and the tangency portfolio.[8] For this reason, the tangency portfolio is considered to be a more relevant portfolio than the GMV portfolio. Therefore, we can refine the original 'overall' spanning test by testing separately whether the inclusion of timberland improves the tangency portfolio or the GMV portfolio, or both. If the overall spanning hypothesis is mainly rejected due to a large improvement of the GMV portfolio, the added value of timberland is limited despite the formal rejection of spanning. Adding timberland to the investment set only pays off if it results in a substantial improvement of the tangency portfolio. In this context Kan and Zhou (2008) propose a 'step-down' approach, which consists of testing the restriction $\alpha = 0$ followed by $\sum_{k=1}^{K} \beta_k = 1$ conditional on $\alpha = 0$. The first step-down test assesses whether the tangency portfolio is improved by adding timberland to the investment set, the second whether the inclusion of timberland improves the GMV portfolio. Spanning is not rejected if both tests are not rejected.[9] To get a complete picture of the impact of timberland on the mean-variance frontier, we will apply both the overall spanning tests and the step-down approach. A Wald test will be used for this purpose (see e.g. Greene, 2007).

If the spanning tests show that the mean-variance efficient portfolio is improved by adding timberland to the portfolio, we would like to quantify the resulting increase in mean-variance efficiency. In other words, we want to assess the economic benefit of adding timberland to the set of benchmark assets. To this extent, we consider the Sharpe ratio (Sharpe, 1966). The Sharpe ratio of a portfolio with return R_t^p is defined as the expected portfolio excess return (relative to the risk-free rate R_t^f) divided by its standard deviation

$$\text{Sharpe}(R_t, R_t^f) = \frac{E(R_t^p) - R_t^f}{\sigma(R_t^p)}. \tag{4}$$

This reward-to-volatility ratio is a performance measure which can be used to compare different portfolios in terms of their risk-adjusted excess return. To assess the economic benefit of investing in timberland, we consider the change in the Sharpe ratio for the tangency portfolio caused by adding timberland to the investment set. For completeness, we also report the resulting change in the Sharpe ratio of the GMV portfolio.

4 Empirical results: spanning and intersection tests

Using the historical returns on the benchmark assets and the Timberland Index, we test for spanning and intersection and assess the economic relevance of adding timberland to the investment set.

4.1 Testing for spanning and intersection

We first assess how the NCREIF Timberland Index performs in relation to each of the asset classes under consideration. We estimate the regression model of Equation (1) and subsequently run the overall spanning tests and apply the step-down approach. We first take the raw Timberland Index as the test asset. Table 4 shows that all spanning and step-down tests are rejected at a 1% significance level. As a second step, we include all assets listed in Table 1 in the benchmark portfolio. Subsequently, we test for spanning and intersection with respect to adding timberland to the full set of benchmark assets.[10] Initially, we do not include the Dow Jones Forestry & Paper Index. Hence, we first regress the Timberland Index on the returns of 13 benchmark assets. The Wald test rejects the hypothesis of spanning at a 1% significance level. Assuming that the zero-beta rate equals

3.9% per year on the basis of historical yearly returns on a 1-month T-Bill during the period 1994 – 2007, we reject the intersection hypothesis at a 1% significance level. Hence, we conclude that adding timberland to the investment set results in a significant improvement in mean-variance efficiency. Moreover, the step-down tests for the tangency and GMV portfolios are rejected at the 1% level. This leads to the conclusion that adding timberland to the investment set does not only improve the mean-variance efficiency of the GMV portfolio, but also that of the tangency portfolio. We repeat the former analysis, but add the Dow Jones Forestry & Paper Index to the benchmark assets. This allows us to assess whether investing in timberland improves the mean-variance frontier if the set of benchmark assets already contains forestry-related investments. Since the regional and global Forestry & Paper indexes are highly correlated, we only include one of them at a time. Again, the hypotheses of spanning are rejected at a 1% significance level; see Table 4. Thus, even when the investment set already contains stocks from the forestry and paper sector, adding the Timberland Index still improves mean-variance efficiency of the portfolio. In particular, both the tangency and the GMV portfolio significantly benefit from adding timberland to the investment set. Even when we additionally add our three indexes of publicly traded timberland to the set of benchmark assets, spanning is still rejected. For the unsmoothed Timberland Index the picture is less favorable. With the complete set of benchmark assets, the overall spanning tests are rejected at the 10% significance level. The same holds for the mean-variance efficiency tests for the tangency portfolio. However, adding timberland to the investment set does not significantly improve the mean-variance efficiency of the GMV portfolio. For the smaller sets of benchmark assets all spanning tests are rejected at the 5% level.

Insert Table 4 about here.

Finally, we turn to the three indexes of publicly traded timberland. Table 5 reports the results of the spanning and step-down tests with publicly traded timberland as a test asset and the same benchmark assets as before. Adding the timberland REIT Index to the investment set only improves the mean-variance efficiency of the tangency portfolio (at a 5% significance level) if the set of benchmark assets consists of the S&P500 or bonds. Regardless of the set of benchmark assets, the other two indexes of publicly traded timberland do not significantly increase the mean-variance efficiency of the tangency portfolio. Although the mean-variance efficiency of the GMV portfolio often benefits from adding timberland, this portfolio is considered less relevant for investors than the tangency portfolio. In sum, publicly traded timberland does not add value from a portfolio perspective. Possible explanations for the differences in diversification potential between private and publicly held timberland investments have been addressed in Section 2.2.

Insert Table 5 about here.

4.2 Economic gains of investing in timberland

We assume a risk-free rate of 3.9% per year and calculate the change in Sharpe ratios for the tangency and GMV portfolios caused by adding the raw Timberland Index to the various investment sets. The results in the upper pane of Table 4 show that the increase in the Sharpe ratio is often much lower for the tangency than for the GMV portfolio. For example, for the full set of benchmark assets the increase in the Sharpe ratio of the GMV portfolio is slightly more than 30 bp, whereas the rise is less than 10 bp per quarter for the tangency portfolio. We propose two additional measures to assess the economic impact of adding timberland to the investment set. First, we consider the weights of the Timberland Index in the tangency portfolio, which reflect the relative importance of the timberland in this portfolio. Next, we compare the mean-variance frontier based solely on the benchmark assets to the frontier including timberland in addition to the benchmark assets. The (long) position in timberland is second largest among all positions, emphasizing its importance. Furthermore, the mean-variance frontier based on the benchmark assets and the raw Timberland Index differs substantially from the frontier derived from the benchmark assets only.

We repeat the above analysis using the unsmoothed Timberland Index, for which the results are reported in the lower pane of Table 4. Regardless of the benchmark portfolio, the Timberland Index significantly improves the mean-variance efficiency of the tangency portfolio at a 10% significance level. However, the economic significance of the increase in Sharpe ratios is substantially smaller than in case of the raw Timberland Index. For example, for the complete set of benchmark assets, including the global Dow Jones Forestry & Paper Index, the rise in the Sharpe ratio of the tangency portfolio is only 1 bp. The corresponding 95% confidence interval around this increase equals $[0.02, 13]$ bp.[11] Thus, although the economic relevance of the rise in the Sharpe ratio of the tangency portfolio seems small, it is statistically significant. The limited economic impact of adding the unsmoothed Timberland Index to the investment set is confirmed by its relatively small weight in the tangency portfolio. Moreover, adding timberland to the investment set does hardly affect the mean-variance frontier.

Our unsmoothing procedure is based on relatively mild assumptions regarding the nature of the appraisal smoothing bias (see Section 2.3), resulting in an unsmoothed index with a relatively low volatility. Even under these favorable assumptions, the economic gains of investing in private equity timberland are limited. By contrast, our analysis based on the raw data suggests that the increase in mean-variance efficiency from adding private equity timberland to the investment set is substantial. Strikingly, the appraisal smoothing has been ignored in previous studies using the NCREIF Timberland Index (see e.g. Sun and Zhang, 2001; Healey et al., 2005). The relatively minor diversification potential of the unsmoothed index emphasizes that any results based on the raw index should be interpreted with caution, as they possibly overstate the mean-variance properties of institutional timberland. Hence, our results also serve as a cautionary warning for future studies using the NCREIF Timberland Index.

5 Robustness of the results

In this section we assess the robustness of several assumptions underlying our analysis.

5.1 Risk-free rate

The way in which the Sharpe ratio is affected by adding timberland to the investment set depends on the risk-free rate, which we assumed to be 3.9% per year on the basis of historical yearly returns on a 1-month T-Bill during the period 1994 – 2007. As a robustness check, we calculate the increase in Sharpe ratios for risk-free rates of 5.7% and 2.1%. These figures reflect the average of 3.9% over the period 1994 – 2007 plus and minus one standard deviation, respectively. In the first part of Table 6 we report some results for the full set of benchmark assets, including the global Dow Jones Forestry & Paper Index. We first consider the raw Timberland Index. The increase in the Sharpe of the tangency portfolio equals 0.2 bp with a 5.7% risk-free rate and 21 bp with a 2.1% risk-free rate. We emphasize that the 0.2 bp increase is still significant, with the 95% confidence interval equal to $[0.02, 11]$ bp.[12] For the unsmoothed Timberland Index the rise in the Sharpe ratios is even smaller, but again statistically significant. Hence, even for a historically low risk-free rate of 2.1%, inclusion of this index in the investment set leads to a significant increase in mean-variance efficiency. Nevertheless, for the tangency portfolio the economic significance of the increase in the Sharpe ratio is very small.

Insert Table 6 about here.

5.2 Wald test

Kan and Zhou (2008) show that the asymptotic Chi-square distribution of the Wald statistic is not always accurate for small samples with a substantial amount of benchmark assets. Therefore, we use a bootstrap approach to obtain the small-sample distribution of the Wald statistic. To deal with any heteroskedasticity in the error term of the regression model of Equation (1), we apply the wild bootstrap (Mammen, 1993). For the full set of benchmark assets including the global Dow Jones Forestry & Paper Index, the second part of Table 6 reports the bootstrapped p-values for the overall spanning and step-down tests. The asymptotic p-values are very close to the bootstrapped ones and tend to be conservative. We establish similar robustness results for the spanning and step-down tests applied to the smaller investment sets. Therefore, the use of the asymptotic Wald test does not seem to be a problem in our case.

5.3 Unsmoothing procedure

For the complete set of benchmark assets, the inclusion of the unsmoothed Timberland Index increases the Sharpe ratio of the tangency portfolio by only 1 bp. The assumed volatility of the unsmoothed index is relatively low compared to the volatility of publicly traded timberland as reported in Table 1. Even for this low volatility, the increase in the Sharpe ratio is small (although statistically significant). For higher volatilities of the unsmoothed Timberland Index the results are even worse, see the third part of Table 6. Assuming a volatility of 10% for the unsmoothed Timberland Index (which comes close to the volatility of publicly traded timberland investments listed in Table 1), the spanning tests are not rejected. With fewer benchmark assets the increase in the Sharpe ratio is more substantial, but obviously the full set is more representative for an institutional investor's well-diversified portfolio.

5.4 Other concerns

We make some final reservations regarding our analysis. First, US private equity timberland investments are represented by the NCREIF US Timberland Index. This index covers only three regions in the US and might not be representative for the US as a whole. Second, timberland is a non-traded, illiquid asset. The quarterly historical returns analyzed in this paper can only be realized by investors with a long-term horizon, such as pension funds or other institutional investors. Hence, there may be a discrepancy between the data frequency (and the resulting Sharpe ratios) and the investment horizon. As shown by Levy (1972), Sharpe ratios computed at frequent intervals are not always the right instrument for making long-term investment decisions. As a consequence, a quarterly 40 bp increase in the maximum Sharpe ratio does not necessarily imply a 80 bp increase on a yearly basis.[13] Unfortunately, the return history of the NCREIF Timberland Index is too limited to do the entire analysis at, say, the yearly level. Moreover, the relevant investment horizon will depend on e.g. the preferences of the institutional investor. Without exact knowledge about these preferences it is not possible to arrive at an appropriate horizon. Furthermore, the mean-variance framework applied in this paper relies on the assumption that investment decisions of institutional investors are solely made on the basis of the mean-variance properties of assets. In reality, other asset characteristics will play a role as well, such as the fact that timberland is often claimed to be an inflation hedge (Washburn and Binkley 1993; Healey et al. 2005). Finally, our analysis is based on historical data and ex post Sharpe ratios, which do not necessarily have predictive power for future performance.

6 Conclusions

This paper analyzes the diversification potential of timberland investments in a mean-variance framework. We focus on both US private equity timberland and publicly traded timberland investments in the US and Canada, where the former is represented by the NCREIF Timberland Index and the latter by three equity indexes. Our approach contributes to the existing literature by explicitly quantifying to what extent timberland investments improve the mean-variance efficiency of a well-diversified portfolio including US and global (small and large cap) stock, bonds, real estate, and commodity indexes.

We find that publicly traded timberland investments from the US and Canada do not significantly improve the mean-variance frontier. At first sight, adding US private equity timberland to the investment set seems to increase mean-variance efficiency, even if the portfolio already contains a forestry and paper equity index. Including timberland in the investment set increases the risk-adjusted excess return on the tangency portfolio with about 10 bp per quarter. However, after removing the appraisal smoothing bias from the raw Timberland Index, there is much less evidence that privately held timberland investments increase mean-variance efficiency. Even under mild assumptions regarding the appraisal smoothing bias, the inclusion of the unsmoothed Timberland Index increases the risk-adjusted excess return on the tangency portfolio with only 1 bp per quarter. Hence, money does not seem to grow on trees.

Appendix A Unsmoothing the Timberland Index

This appendix addresses the problems caused by the use of appraisal values for the construction of the NCREIF Timberland Index.

A.1 Appraisal problems

Most timberland properties in the index are appraised only once a year, usually in the fourth quarter and – to a lesser extent – in the second quarter. Properties whose value is not updated in a given quarter are included in the index with the same value as in the previous quarter (Lutz, 1999). Consequently, temporal aggregation of individual properties could result in an artificially high quarterly index return variance due to the individual appreciation returns with value zero. Moreover, the seasonality of reappraisals may induce spurious positive fourth-order autocorrelation in the index returns. Second, appraisers tend to 'smooth' or 'partially adjust' property values over time, which leads to appraisal smoothing. As a consequence, the volatility of index returns could be biased towards zero and the index may lag changes in underlying property values. Similar appraisal problems have been discussed in the context of other appraisal-based property indexes such as the Russell-NCREIF (Fisher et al., 1994). Biases caused by appraisal smoothing, seasonality in reappraisals and temporal aggregation may seriously distort any mean-variance analysis, since return volatility may seem lower than it actually is. However, we emphasize that the appraisal problems do not necessarily occur. Edelstein and Quan (2006) argue that appraisal smoothing at the individual level does not always lead to an appraisal smoothing bias at the index level, since individual property appraisal biases may offset in the aggregate. Lai and Wang (1998) show that the use of appraisal-based data can in fact result in a *higher* variance than that of the true index returns.

A.2 Unsmoothing approach

Fisher et al. (1994) propose a method to 'unsmooth' the quarterly Russell-NCREIF Index. This property index suffers from the same potential appraisal problems as the NCREIF

Timberland Index. Therefore, we apply the methodology of Fisher et al. (1994) to remove the autocorrelation from the index returns and to increase its variance (while keeping the mean unchanged). First, it is assumed that the *observed* index return r_t^* is a weighted average of current and past *true* index returns r_{t-k}, i.e.

$$r_t^* = w(L)r_{t-1} + w_0 r_t, \tag{A.1}$$

with w_0 a scalar satisfying $0 < w_0 < 1$ and $w(L)$ a polynomial in terms of the lag operator. Alternatively, we can write

$$r_t^* = \psi(L)r_{t-1}^* + w_0 r_t. \tag{A.2}$$

To deal with appraisal smoothing and seasonality in reappraisals, we take $\psi(L) = \psi_1 + \psi_4 L^3$ for the lag polynomial $\psi(L)$. This results in

$$r_t^* = \psi_0 + \psi_1 r_{t-1}^* + \psi_4 r_{t-4}^* + \eta_t, \tag{A.3}$$

with $\eta = w_0 r_t$. We can now write Equation (A.3) in terms of true returns r_t, assuming $\mathbb{E}r_t = \mu$. This yields

$$r_t = \mu + (r_t^* - \psi_0 - (\psi_1 + \psi_4 L^3)r_{t-1}^*)/w_0. \tag{A.4}$$

We can estimate the coefficients ψ_0, ψ_1 and ψ_4 from Equation (A.3) by assuming that true returns (r_t) are serially uncorrelated (in which case $\eta_t = w_0 r_t$ is white noise). The value of w_0 can only be derived by imposing an additional restriction, e.g. that the volatility of the true returns is κ%. This yields

$$w_0 = \frac{\sigma(\eta_t)}{\kappa}. \tag{A.5}$$

According to Table 1, the volatility of publicly traded timberland is around 7–12%, whereas the volatility of the raw NCREIF Timberland Index is only about 3%. Therefore, the true volatility of the unsmoothed timberland index is likely to be in the range of 3–12%. We therefore consider a volatility range of 3-12% and obtain the unsmoothed Timberland Index for all (rounded values of the) volatilities in this range using the approach suggested by Fisher et al. (1994). We take the complete set of benchmark assets and determine the *largest* volatility for which the step-down test with respect to the tangency portfolio is not rejected at a 10% significance level. Since the mean of the unsmoothed series is not affected by the unsmoothing procedure, the unsmoothed asset will become less attractive when its volatility is increased. Our results show that the unsmoothed Timberland Index significantly increases the mean-variance efficiency of the tangency portfolio for volatilities up to 5.3%. Therefore, we use this version of the unsmoothed index in our analysis. OLS estimation of Equation (A.3) applied to the NCREIF Timberland Index results in the following estimates (p-values based on White's heteroskedasticity robust covariance matrix in parentheses): $\psi_0 = 1.53$ (0.005), $\psi_1 = 0.04$ (0.505), $\psi_4 = 0.30$ (0.100), and $w_0 = 0.48$.

References

Belsley, D., E. Kuh, and R. Welsch. 1980. *Regression Diagnostics*. Wiley.

DeRoon, F.A. and Th. E. Nijman. 2001. Testing for mean-variance spanning: A survey. *Journal of Empirical Finance* 8: 111-155.

Edelstein, R.H. and D.C. Quan. 2006. How does appraisal smoothing bias real estate returns measurement. *Journal of Real Estate Finance and Economics* 32: 41-60.

Fisher, J.D., D.M. Geltner, and R.B. Webb. 1994. Value indices of commercial real estate: A comparison of index construction methods. *Journal of Real Estate Finance and Economics* 9: 137-164.

Greene, W.E. 2007. *Econometric Analysis*, 6th edition. Prentice Hall.

Healey, T., T. Corriero, and R. Rosenov. 2005. Timber as an institutional investment. *Journal of Alternative Investments*, Winter 2005: 60-74.

Huberman, G. and S. Kandel. 1987. Mean variance spanning. *Journal of Finance* 42: 873-888.

Kan, R. and G. Zhou. 2008. Tests of mean-variance spanning. OLIN Working Paper No. 99-05; Rotman School of Management Working Paper.
Available at SSRN: http://ssrn.com/abstract=231522.

Lai, T-Y. and K. Wang. 1998. Appraisal smoothing: The other side of the story. *Real Estate Economics* 26: 511-535.

Levy, H. 1972. Portfolio performance and the investment horizon. *Management Science* 18: 645-653.

Lutz, J. 1999. Measuring timberland performance. Timberland Report 1(2). James W. Sewall Company. URL: http://www.jws.com/pdfs/timberlandreport/v1n4.pdf.

Mammen, E. 1993. Bootstrap and wild bootstrap for high - dimensional linear models. *Annals of Statistics* 21: 255 - 285.

Redmond, C.H. and F.W. Cubbage. 1988. Portfolio risk and returns from timber asset investments. *Land Economics* 64: 325-337.

Riddiough, T.J., M. Moriarty and P.J. Yeatman. 2005. Privately versus publicly held asset investment performance. *Real Estate Economics* 33: 121-146.

Sharpe, W.F. 1966. Mutual fund performance. *Journal of Business* 39: 119-138.

Sun, C. and D. Zhang. 2001. Assessing the financial performance of forestry-related investment vehicles: capital asset pricing model vs. arbitrage pricing theory. *American Journal of Agricultural Economics* 83: 617-628.

Tobin, J. 1958. Liquidity preference as behavior towards risk. *Review of Economic Studies* 25: 65:86.

Washburn, C.L. and C.S. Binkley. 1993. Do forest assets hedge inflation? *Land Economics* 69: 215-224.

index	description of index	data source	mean (%)	volatility (%)	correlation with Timberland Index
S&P		Thomson Datastream			
S&P500	US equity index		2.60	7.78	0.14
MSCI		Thomson Datastream			
World	global equity index		2.23	7.50	0.19
Far East	regional equity index		0.45	10.22	0.10
Emerging Markets	regional equity index		2.08	12.88	-0.02
Lehman		Thomson Datastream			
Global Aggregate	global bond index		-0.18	1.73	-0.15
FTSE		www.ftse.com			
EPRA/NAREIT Global Real Estate	global real estate equity index		2.80	7.57	-0.02
EPRA/NAREIT US Real Estate	US real estate equity index		2.87	6.88	0.07
Dow Jones		www.djindexes.com			
AIG Commodity	global commodity index		2.37	6.24	-0.02
Forestry & Paper (global)	global forestry & paper stocks index		1.02	9.86	-0.12
Global Small Cap	global small cap equity index		2.32	8.95	-0.01
Latin America	regional equity index		2.04	16.27	0.01
Canada	regional equity index		3.40	10.44	0.10
Forestry & Paper (America)	regional forestry & paper stocks index		1.27	11.19	-0.11
Composite	US equity index		2.62	7.30	0.15
US Technology	US technology stocks index		2.94	16.47	0.07
NCREIF		www.ncreif.com			
Timberland			2.39	2.74	1.00
Unsmoothed Timberland			2.39	5.3	0.95
Other indexes (value-weighted)		Thomson Datastream			
Timberland REIT	index of 3 timberland REIT's (Rayonier, Plum Creek and Potlatch)		2.63	7.51	-0.02
Paper & Paper Products	Paper & Paper Products Industry index (US and Canada)		1.67	9.81	0.02
Lumber & Wood Production	Lumber & Wood Production Industry index (US and Canada)		2.48	9.81	-0.11

Table 1: Benchmark indexes and timberland index

This table lists the benchmark indexes and the timberland indexes, together with a short description (second column), data source (third column), mean quarterly return (fourth column), quarterly volatility (fifth column) and quarterly correlation with the Timberland Index (sixth column) during the period 1994-2007. All indexes under consideration are total return indexes corrected for dividend payments and stock splits.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1:Timberland	1.00																			
2:Unsmoothed Timberland	0.95	1.00																		
3:Timber REIT	-0.02	-0.01	1.00																	
4:LWPI	0.02	0.01	0.80	1.00																
5:PPPI	-0.11	-0.13	0.52	0.67	1.00															
6:DJ AIG Commodity	-0.03	0.02	0.17	0.14	0.12	1.00														
7:Global Real Estate	-0.02	-0.04	0.20	0.16	0.33	0.15	1.00													
8:US Real Estate	0.07	0.09	0.40	0.27	0.36	0.06	0.77	1.00												
9:MSCI Emerging Markets	-0.01	0.00	0.11	0.14	0.24	0.14	0.68	0.35	1.00											
10:MSCI World	0.19	0.16	0.04	0.07	0.21	-0.07	0.60	0.31	0.73	1.00										
11:MSCI Far East	0.09	0.09	0.15	0.17	0.20	0.13	0.60	0.15	0.71	0.68	1.00									
12:Lehman Global Aggregate	-0.17	-0.25	0.03	-0.01	0.01	-0.20	0.06	0.16	-0.02	-0.02	-0.22	1.00								
13:DJ Forestry & Paper (Am.)	-0.10	-0.12	0.11	0.14	0.45	-0.01	0.65	0.38	0.63	0.68	0.50	0.02	1.00							
14:DJ Forestry & Paper (Global)	-0.12	-0.14	0.17	0.19	0.45	0.03	0.68	0.39	0.66	0.69	0.61	-0.02	0.96	1.00						
15:DJ Latin Am	0.02	0.04	0.15	0.12	0.19	0.15	0.63	0.40	0.91	0.72	0.55	0.04	0.61	0.60	1.00					
16:DJ Canada	0.12	0.13	0.14	0.14	0.27	0.18	0.62	0.35	0.78	0.85	0.62	-0.09	0.57	0.61	0.77	1.00				
17:DJ USA Technology	0.08	0.11	-0.04	-0.05	0.13	-0.13	0.39	0.13	0.61	0.83	0.53	-0.04	0.48	0.46	0.58	0.74	1.00			
18:DJ Global Small Cap	-0.01	-0.02	0.20	0.19	0.34	0.06	0.73	0.45	0.82	0.89	0.72	-0.02	0.70	0.75	0.78	0.87	0.74	1.00		
19:DJ Composite	0.15	0.11	0.12	0.15	0.30	-0.02	0.59	0.43	0.61	0.87	0.46	0.10	0.73	0.68	0.65	0.73	0.68	0.79	1.00	
20:S&P 500	0.14	0.11	-0.06	-0.08	0.11	-0.14	0.50	0.27	0.62	0.88	0.49	0.10	0.55	0.52	0.63	0.74	0.91	0.78	0.83	1.00

Table 2: Correlation matrix for quarterly returns on benchmark and timberland indexes (PPPI stands for Paper & Paper Products Index and LWPI for Lumber & Wood Production Index)

	α	p-value	β	p-value
Timberland Index	1.35	0.00	0.04	0.43
Unsmoothed Timberland Index	1.34	0.00	0.04	0.47
Timberland REIT Index	1.74	0.16	-0.05	0.00
Paper & Paper Products Index	0.50	0.67	0.11	0.46
Lumber & Wood Production Index	1.66	0.29	-0.09	0.00

Table 3: CAPM estimation results

Regression model:
$r_t = \alpha + \beta R_t + \varepsilon_t$

Sample period: 1994-2007
Frequency: quarterly

benchmark assets	***p*-value overall spanning**	***p*-value step-down TP**	***p*-value step-down GMVP**	**Δ Sharpe TP (bp)**	**Δ Sharpe GMVP (bp)**
Raw Timberland Index					
S&P500	0.000	0.000	0.000	33	33
Bonds	0.000	0.000	0.000	13	41
Equity (- forestry)	0.000	0.000	0.000	29	37
Equity (+ forestry America)	0.000	0.000	0.000	27	36
Equity (+ forestry Global)	0.000	0.000	0.000	27	37
Real estate	0.000	0.000	0.000	30	30
Commodity	0.000	0.000	0.000	35	35
All (- forestry)	0.000	0.000	0.001	10	32
All (+ forestry America)	0.000	0.000	0.001	9	32
All (+ forestry Global)	0.000	0.000	0.002	8	33
Unsmoothed Timberland Index					
S&P 500	0.000	0.001	0.000	11	11
Bonds	0.000	0.001	0.004	3	14
Equity (- forestry)	0.000	0.012	0.000	8	12
Equity (+ forestry America)	0.000	0.015	0.000	6	10
Equity (+ forestry Global)	0.000	0.019	0.000	6	11
Real estate	0.000	0.007	0.000	10	10
Commodity	0.000	0.010	0.000	13	13
All (- forestry)	0.097	0.093	0.123	2	7
All (+ forestry America)	0.092	0.097	0.134	1	8
All (+ forestry Global)	0.091	0.093	0.147	1	8

Table 4: Results of spanning tests with different benchmark assets (NCREIF Timberland Index)

For various investment sets, this table provides p-values for the overall panning and step-down tests. It also reports the increase in the Sharpe ratio caused by adding the NCREIF Timberland Index to the investment set. We consider the increase in the Sharpe ratio of the tangency portfolio (TP) and the global minimum variance portfolio (GMVP). Three situations are considered when the investment set consists of equity indexes: (1) the Dow Jones Forestry & Paper Index is not included in the set of benchmark assets, (2) the Dow Jones Forestry & Paper Index (Americas) is included, and (3) the global Dow Jones Forestry & Paper Index is included. All p-values are based on White's heteroskedasticity robust covariance matrix.

benchmark asset	***p*-value overall spanning**	***p*-value step-down TP**	***p*-value step-down GMVP**
Timberland REIT Index			
S&P500	0.000	0.038	0.000
Bonds	0.006	0.014	0.174
Equity (- forestry)	0.000	0.097	0.000
Equity (+ forestry America)	0.001	0.137	0.000
Equity (+ forestry Global)	0.000	0.137	0.000
Real estate	0.000	0.210	0.000
Commodity	0.000	0.053	0.000
All (- forestry)	0.201	0.080	0.742
All (+ forestry America)	0.215	0.086	0.747
All (+ forestry Global)	0.204	0.081	0.741
Paper & Paper Products Index			
S&P500	0.000	0.254	0.000
Bonds	0.105	0.148	0.099
Equity (- forestry)	0.000	0.454	0.000
Equity (+ forestry America)	0.000	0.378	0.000
Equity (+ forestry Global)	0.000	0.369	0.000
Real estate	0.001	0.721	0.000
Commodity	0.000	0.220	0.000
All (- forestry)	0.590	0.581	0.385
All (+ forestry America)	0.456	0.395	0.379
All (+ forestry Global)	0.434	0.411	0.340
Lumber & Wood Production Index			
S&P 500	0.000	0.097	0.000
Bonds	0.055	0.077	0.224
Equity (- forestry)	0.001	0.205	0.002
Equity (+ forestry America)	0.005	0.287	0.003
Equity (+ forestry Global)	0.006	0.291	0.005
Real estate	0.003	0.344	0.001
Commodity	0.003	0.223	0.001
All (-forestry)	0.542	0.270	0.984
All (+forestry America)	0.568	0.290	0.966
All (+forestry Global)	0.568	0.289	0.963

Table 5: Results of spanning tests with different benchmark assets (publicly traded timberland)

For various investment sets, this table provides p-values for the overall spanning and step-down tests. The test assets are the three (value-weighted) indexes of publicly traded timberland. All p-values are based on White's heteroskedasticity robust covariance matrix.

1. impact of risk-free rate				
	raw index (bp)	**unsmoothed index (bp)**		
Δ Sharpe TP (risk-free rate 3.9%)	8	1		
Δ Sharpe TP (risk-free rate 5.7%)	0.2	0.1		
Δ Sharpe TP (risk-free rate 2.1%)	21	4		
Δ Sharpe GMVP (risk-free rate 3.9%)	33	8		
Δ Sharpe GMVP (risk-free rate 5.7%)	27	6		
Δ Sharpe GMVP (risk-free rate 2.1%)	39	9		
2. impact of asymptotic Wald test				
	raw index		**unsmoothed index**	
	p-value (asymp.)	**p-value (bootstrap)**	**p-value (asymp.)**	**p-value (bootstrap)**
overall spanning test	0.000	0.000	0.091	0.078
step-down TP	0.000	0.000	0.093	0.073
step-down GMVP	0.002	0.002	0.147	0.108
3. impact of unsmoothing approach				
mean (%)	2.39			
volatility (%)	10.00			
	p-value			
overall spanning test	0.710			
step-down TP	0.601			
step-down GMVP	0.520			

Table 6: Results of sensitivity analysis

For the complete set of benchmark assets (including the global Dow Jones Forestry & Paper Index), this table displays the outcomes of several sensitivity tests. The first sensitivity test relates to the risk-free rate. The first part of this table reports the increase in the Sharpe ratios of the tangency and GMV portfolios for several values of the risk-free rate (for both the raw and the unsmoothed Timberland Index). The second part of this table displays the usual asymptotic p-values for the overall spanning and step-down tests, as well as bootstrapped p-values for the same tests. Finally, the third part of the table reports the sample mean and volatility of the alternative unsmoothed Timberland Index, as well as the corresponding p-values for the overall and step-down spanning tests.

Notes

[1]See www.nga.org.

[2]Apart from these indexes, we also considered several other indexes. Eventually, we did not include them in the set of benchmark assets as they turned out collinear with one or more other indexes. Since collinearity might be problematic, we do not include these indexes in the investment set. However, we emphasize that our results are robust to the choice of the set of benchmark assets. The indexes we omitted because of collinearity with other benchmark assets are: MSCI Europe, Dow Jones US, Dow Jones US Small Cap, Dow Jones Euro Small Cap, Dow Jones Asia/Pacific, Lehman Global Treasury, Lehman Investment Grade, Lehman US Treasury, Lehman US Corporate Investment Grade, Lehman US Government Aggregate, and Lehman US Aggregate Index.

[3]These eight TIMO's are The Campbell Group, Global Forest Partners, Forest Investment Associates, Hancock Timber Resource Group, Molpus Woodlands, Resource Management Service, RMK Timberland, and Timberland Investment Resources. Prudential Timber was one of the founding contributors. Another past contributor is Forest Systems.

[4]The equally-weighted indexes yield very similar results. For this reason, we do not report them. However, they are available upon request.

[5]Our Paper & Paper Products Index contains International Paper, Avery Dennison, Sonoco Products, Meadwestvaco, Rock-Tenn, Domtar, Tembec, Potlatch, Mercer International, and Norbord. Our Lumber & Wood Production Index contains Leucadia, Masco, Universal Forest Products, Louisiana Pacific, Weyerhaeuser, Pope Resources, Cross Timbers Royalty Trust, West Fraser Timber, Cascades, Sino-Forest, International Forest Products, and Canfor. We downloaded all total return data from Thomson Datastream. At the end of 2007 (which is the end of our sample period) the Dow Jones Forestry & Paper Index (Americas) contains the stocks AbitibiBowater, Aracruz Celulose, Canfor, Cascades, Catalyst Paper, Domtar, Empresas, International Paper, Neenah Paper, Sino-Forest, Suzano Papel e Celulose, TimberWest Forest, Wausau Paper, West Fraser Paper, Weyerhaeuser. The global Dow Jones Forestry & Paper Index contains quite a few additional companies in the forestry and paper industry located outside America. The full list is available on http://www.djindexes.com/.

[6]We notice that Rayonier, Plum Creek and Potlatch became REIT's after 1994.

[7]Alternatively, we could apply the approach proposed by Riddiough et al. (2005), which consists of replicating the NCREIF Timberland Index by means of publicly traded timberland REIT's (while correcting for e.g. leverage). However, it seems difficult to apply this method to the Timberland Index, since there are currently only three timberland REIT's (who all adopted the REIT structure relatively recently). Therefore, we opt for a different approach and apply the unsmoothing method introduced by Fisher et al. (1994).

[8]We notice that, according to the CAPM model, the tangency portfolio will be a value-weighted mix of

all assets in the world.

[9]Suppose that the significance level of the first step-down test is α_1 and that of the second α_2. Then the significance of the joint step-down test is equal to $1 - (1 - \alpha_1)(1 - \alpha_2)$.

[10]To make sure that the benchmark assets in our analysis are not collinear, we follow the procedure proposed by Belsley, Kuh, and Welsch (1980) and inspect the condition indices and variance decomposition proportions corresponding to the matrix of benchmark returns. We do not find any evidence for multicollinearity.

[11]We obtained this confidence interval by means of a wild bootstrap (Mammen, 1993). We use the same bootstrap procedure to derive the finite-sample distribution of the Wald test in Section 5.2.

[12]Again we obtained this confidence interval by means of a wild bootstrap.

[13]This number is obtained as $80 = 40 \times 4/\sqrt{(}4)$.