A review of real estate and infrastructure investments by the Norwegian Government Pension Fund Global (GPFG)

Stijn Van Nieuwerburgh   Richard Stanton
NYU   U.C. Berkeley

Leo de Bever
Former CEO, Alberta Investment Management Corporation

December 1, 2015

Preface

The Norwegian Government Pension Fund Global represents a rare instance of a people deciding that the financial proceeds from extracting a country's non-renewable resources should be shared with future generations. We are honored to be asked to assist in determining whether investments in real estate and infrastructure can augment the Fund's future growth.

Contents

1 Recommendations

1. The academic literature we review concludes that both listed and unlisted real estate have the same characteristics in the long run. It finds no strong evidence for superior performance or lower risk for unlisted-asset investments; neither does it find evidence for long-term diversification benefits of adding unlisted assets to a listed-real-estate portfolio. Pragmatically, the unlisted asset space is far too large a segment of the real estate market (about 80%) to ignore for any large investor like the GPFG. We see no compelling reason for a dramatic shift away from the current listed-unlisted split of GPFG's real estate portfolio. We recommend that the Ministry of Finance, in its management mandate to Norges Bank, allow for both listed and unlisted real estate investments, as it does today.

2. Global listed real estate has fairly compensated investors for exposure to global stock-market and bond-market risks in the period 1994–2015. We find that correlations between real estate and returns on stocks and bonds have been rising over time, reducing the diversification benefits from real estate and resulting in a higher required rate of return. This evidence does not warrant a recommendation to the Ministry to open up the mandate for a higher share of the Fund's capital to be invested in real estate.

3. Current valuation and risk levels of real estate assets, especially for core assets in top-tier gateway cities, seem elevated compared to historical pricing. We recommend caution for new real estate investments.

4. Global listed infrastructure investments have fairly compensated investors for exposure to global stock-market and bond-market risks in the period 1999–2015. However, there is some evidence for abnormal returns on core infrastructure assets (in the period 2003–2015 with available data), which could justify an overweight position in the portfolio. Current valuation and risk levels of listed infrastructure reflect more moderate income growth expectations than in real estate.

5. The Ministry should open up for unlisted infrastructure investment in the management mandate to Norges Bank to take advantage of investment opportunities unavailable in the listed space.

6. We recommend that the Ministry use the Opportunity Cost Model (OCM) for its real

estate and infrastructure investments in the management mandate to Norges Bank. The OCM approach delegates investment authority to the asset manager, who is closer to the market. This approach should result in better adaptability to changing opportunities across and within asset classes. With the OCM, real estate and infrastructure investments (beyond those already included in the benchmark stock and bond indices) are only justified if their expected returns exceed those of the appropriate combinations of stocks and bonds. Rather than filling a target allocation to real estate (or infrastructure), the OCM shifts the focus from asset-class labels to the underlying risk exposures. While the application of the OCM to real estate and infrastructure investments is more challenging than for listed assets, we provide specific recommendations in the report for how best to address these challenges.

7. The OCM obviates the need for a separate real estate benchmark like the IPD index. We discuss several problems with appraisal-based indices like IPD in the report. The OCM provides an easy-to-understand and measure benchmark consisting of listed stock and bond indices.

8. The Ministry should set *maximum* portfolio weights for real estate and infrastructure of 10% each to limit total portfolio volatility. Norges Bank should *only* use this flexibility if return expectations warrant it (point 6). The current environment may not be the best to fill this allocation.

9. Norges Bank should not use tracking error to measure the active risk in its real estate and infrastructure investments. Rather, we favor the use of the OCM with maximum weights, combined with a set of reporting requirements spelled out in the report. By construction, the OCM generates a portfolio with the same amount of systematic risk as the Reference Portfolio of stocks and bonds.

10. To further improve transparency, the Ministry should require Norges Bank to report detailed costs for managing the real estate and infrastructure portfolios.

11. We recommend that the Ministry open up for unlisted clean energy infrastructure investments in the management mandate to Norges Bank. This would give Norges Bank the opportunity to explore such investments. These investments will constitute a majority of energy investments over the next 30 years.

12. We recommend that the Ministry open up for unlisted emerging-market infrastructure investments in the management mandate to Norges Bank. This would give Norges Bank the opportunity to explore such investments. Listed EM infrastructure investments have shown strong historical performance, after accounting for standard sources of risk. The enormous need for all types of infrastructure in developing countries and the shrinking role of traditional funding sources provides a compelling rationale for continued growth. The main challenge lies in managing several incremental sources of risk such as political risk, regulatory risk, and management & governance risk.

13. In the management mandate to Norges Bank, the Ministry has opened up for real estate investments in developing countries. Norges Bank's strategy has been to start by investing in mature markets. We recommend that Norges Bank, having built up this investment expertise, explore investments in select real estate projects in developing countries. Due to urbanization, a growing middle class, and a rebalancing towards a larger service sector, much of the world's future demand for real estate will be in developing countries.

14. The management mandate should give Norges Bank the opportunity to continue to build up its internal real estate investment team for both its listed and unlisted portfolios, as it does today. We recommend that Norges Bank also start building a team with expertise in listed and unlisted infrastructure. Institutional investors of a certain size have achieved lower costs, and higher gross returns, with internal than with external management, both in the unlisted and listed space. GPFG's size should allow it to fully exploit these economies of scale.

15. The Opportunity Cost Model delegates the implementation of the real estate and infrastructure investment decision to Norges Bank. Nevertheless, we would like to express our differing views on the ultimate composition of the infrastructure portfolio. Two of us (Van Nieuwerburgh and Stanton) recommend that Norges Bank hold the majority of its infrastructure portfolio in listed form. Reasons for this choice are:

    (a) Most private infrastructure assets are currently held in listed form.
    (b) There is a lack of data that convincingly shows evidence of outperformance on a risk-adjusted basis of unlisted investments. Listed infrastructure assets have performed reasonably well (point 4).
    (c) Unlisted asset investments are subject to myriad non-financial risks (political, reputational, environmental) and operational complexities.

(d) Building a substantial direct infrastructure portfolio requires many boots on the ground. Norges Bank's compensation practices and culture may not allow it to compete globally for senior talent.

One of us (de Bever) believes that Norges Bank should invest the majority of its infrastructure portfolio in unlisted form. Reasons for this choice are:

(a) The observed investment strategies of GPFG's peers are almost exclusively focused on unlisted assets. No peer fund seems to target a proportionate slice of either the real estate or the infrastructure market, listed or unlisted. All seem to concentrate on large, high-quality assets in real estate and stable-return infrastructure assets.

(b) Being able to deploy stable capital long term provides opportunities to create value in relatively inefficient markets beyond trading in zero-sum-game listed markets.

(c) The main governance issue is seen to be lack of alignment with external manager objectives. Direct investments provide better control over asset governance. Stale valuations are a minor annoyance for very long term investors.

(d) Any liquidity advantage of listed assets is limited, given the large size of the holding for GPFG peer institutions.

2 Introduction

In line with the mandate given to us by the Ministry of Finance, the goal of this report is to consider how GPFG should think about investing in real estate and possibly infrastructure going forward, including both an assessment of risk and return and a consideration of how Norges Bank's investments should be regulated and monitored by the Ministry of Finance.

We start with an overview of the real estate and infrastructure markets and their place in global capital markets (Section 3). This analysis also tells us the allocation to real estate and infrastructure of the "average" investor, a simple benchmark for GPFG. Real estate represents about 6% and (private) infrastructure about 2% of the world market portfolio. Listed assets make up only about 15% of all real estate investments, but about 85% of the infrastructure universe.

We next compare the risk and return of listed versus direct investments in Section 4, which addresses the following questions:

- What are the challenges with measuring direct real estate and infrastructure returns?
- After addressing these challenges, how do risk and return of direct asset investments compare with those for listed investments? Is there an illiquidity premium on directly held assets?
- How do large institutional investors choose their mix of direct and listed assets and how has their gross and net performance been? What is the role of external versus internal management?
- How do large institutional investors structure their real estate and infrastructure teams?

To decide whether the Fund should increase its real estate or infrastructure allocations, we evaluate the risk and return of real estate and infrastructure in Section 5, which explores the following specific questions:[1]

- What are the average returns and return volatilities on listed real estate and infrastructure? How do the returns correlate with stocks and bonds? We use a 1994–2015 sample, where we have global real estate data, and two shorter samples (1999–2015 and 2004–2015), where we add two global infrastructure indices.[2]
- How much of the return on real estate and infrastructure can be explained by the return on stocks and bonds? What is the expected return implied by such a two-factor model? Have real estate and infrastructure outperformed that expected return?

---

[1]Given the finding in Section 4 that the risk properties of direct real estate are well described by those of listed real estate, Section 5 focuses exclusively on listed real estate. Throughout, our focus is on equity, not debt interests.

[2]We also explore the U.S. in Appendix A, where we have high quality data going back to 1972, to provide a longer-term perspective.

- What other risk factors are real estate and infrastructure exposed to? How does this exposure affect expected returns?
- How much variation is there over time in the risk (volatility and correlation) of these asset classes?
- How does the return volatility change as the investment horizon changes?
- How do real estate returns correlate with inflation and economic growth?
- What should be the weight on real estate and infrastructure in a portfolio alongside stocks and bonds, using mean-variance optimization?
- Given the expected return on real estate and infrastructure implied by the model, what assumptions on future cash flow growth are implied by current valuation ratios? Are they reasonable?

Section 6 tackles the second part of the mandate: how the Ministry should regulate real estate and possible infrastructure investments in the management mandate to Norges Bank and how Norges Bank's performance should be reviewed. Given the drawbacks with the current real estate benchmark IPD index (Section 4), we propose that the Fund switch to the opportunity cost (OC) model recommended also by Ang, Brandt, and Denison (2014a). Section 6 describes how to deal with the complications specific to real estate and infrastructure when it comes to good governance, addressing the following specific issues:

- Why using tracking error as a risk-management tool may not be the best way to go for real estate and infrastructure investing.
- How to use the Opportunity Cost model to evaluate real estate and infrastructure investments, and how to address challenges with implementation.
- How to measure and manage risk in the real estate and infrastructure portfolios in the OC model framework.
- What reporting requirements the Ministry of Finance should mandate for Norges Bank's real estate and infrastructure investments.
- What return expectations the Ministry of Finance should have for Norges Bank's real estate and infrastructure investments and how it should evaluate Norges Bank's performance.
- How Norges Bank's senior management and real estate and infrastructure managers should be compensated.

Section 7 addresses the specific questions about renewable energy infrastructure and emerging-market infrastructure. In the case of renewable energy, the short time series of data necessitates a more qualitative, forward-looking perspective.

Section 8 discusses how the Fund's peers (especially in Canada) have sought to maximize net return on risk, with a focus on:

- The investments beliefs that underpin their management culture.
- The challenges of being a long-term investor.
- Their governance structure and risk management.
- Their performance measurement and compensation.

This section reflects the experience of one of the members of this committee (de Bever) with investments in real estate and infrastructure.

3 Market sizes

3.1 GPFG

The Norwegian Government Pension Fund Global (GPFG) was set up in 1990 to manage petroleum revenues on behalf of the Norwegian people. It is now the world's largest sovereign-wealth fund, with assets under management of 7.019 trillion Norwegian kroner ($822.9 billion) on September 30, 2015. The Fund was first opened up to real-estate investment (both listed and unlisted, and including real estate derivatives), up to a limit of 5% of assets, in March 2010, and the first unlisted real estate investment was made in April 2011. As of September 30, 2015, real estate represented 3.0% of the Fund's portfolio, mostly focused on core markets in Europe and the U.S. The overall breakdown is shown in Table 1.

	Value (Billions)		Share
	(NOK)	($)	(%)
Equity	4,191	491.3	59.7
Fixed-income	2,620	307.2	37.3
Real estate	208	24.4	3.0
Total	7,019	822.9	100.0

Table 1: Market value of Government Pension Fund Global (GPFG) on September 30, 2015. Source: `http://www.nbim.no/en/the-fund/market-value/key-figures/`. Dollar values calculated using the September 30, 2015 exchange rate of 8.53 NOK/$ (source: Datastream).

Despite this growth, both the target allocation and invested amount lag well behind other institutional investors. For example, in their survey of 231 institutional investors in 28 countries, Funk, Weill, and Hodes (2014) report an average target allocation to real estate of 9.38% in 2014 (up 49 bps from 2013), with an intention to increase this to 9.62% during 2015. The average fraction actually invested is only slightly lower, at 8.49%. Preqin (2015a) reports average current allocation to infrastructure for investors has increased from 3.5% of AUM in 2011 to 4.3% in 2014 and 2015. Target allocations to the asset class have continued to grow in 2015 and now stand at 6.3% of AUM for those investors allocating to the asset class. Allocations are likely to continue to grow over the coming years, with 44% of investors planning to increase the amount of capital they invest in infrastructure.

3.2 The world market portfolio

In the Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1965), homogeneous investors obtain the optimal trade-off between risk and return by investing in a

combination of the "market portfolio"[3] and riskless bonds. Black (1989, 1990) and Black and Litterman (1992) extended this idea to a global context, showing that an investor in global markets would choose some combination of the global market portfolio plus riskless borrowing or lending, combined with some amount of currency hedging. While there has been a huge amount of research on the topic of optimal portfolio allocation since then,[4] the global market portfolio nevertheless provides one simple benchmark for the proportions to be invested in any given asset class.

The global market portfolio Estimating the size of the global market portfolio is obviously difficult, but there are a few attempts in the literature. Doeswijk, Lam, and Swinkels (2014), estimate the total market capitalization of the invested global multi-asset market portfolio in 2012 at $90.6 trillion, of which the largest components are $32.9 trillion in listed equities (36.3%), $26.7 trillion in government bonds (29.5%), and $16.8 trillion in investment-grade corporates (18.5%). Hewitt EnnisKnupp (2014) estimate the size of the world market portfolio as of June 30, 2013, at $98.66 trillion, of which $18.19 trillion is U.S. equities, $13.83 trillion is non-U.S. equities, $22.65 trillion is in non-U.S. bonds (developed), and $15.34 trillion is in U.S. bonds (investment grade).

Real estate Doeswijk et al. (2014) estimate real estate equity to make up $4.6 trillion of the total, and Hewitt EnnisKnupp (2014) estimate $5.46 trillion as of June 2013, split into private equity at $4.2 trillion and public equity at $1.26 trillion. Other estimates for the total size of the real estate equity market include Almond (2015) ($6.1 trillion in 2014, split into $5.1 trillion private equity and $1.0 trillion public equity) and Clacy-Jones and Teuben (2014) ($6.8 trillion at the end of 2013, but not split into public/private). Based on the numbers from Almond (2015), listed real estate equity represents about 1%, and private real estate equity about 5%, of the global market portfolio.

Infrastructure RARE (2013) estimates the total size of all global infrastructure assets in 2012 to be $20 trillion, with 75% of this amount government owned and 25% ($5 trillion) privately owned. Of this, listed equity represents about $2.5 trillion (roughly 2.5% of the global market portfolio) and unlisted fund equity a further $250 billion (0.25% of the global market portfolio).

The average investor Taken together, these estimates suggest that real estate represents roughly 6% of the global market portfolio, the vast majority of this unlisted, and infrastruc-

---

[3]This portfolio contains all assets available to investors, in amounts proportional to their market values.
[4]Not to mention the classic critique of the CAPM, Roll (1977).

ture roughly another 3%. By definition, these are the proportions invested by the "average" investor, and thus represent a simple first pass at target allocations for GPFG. Of course, GPFG is quite different from the average investor in several ways, which interact with the characteristics of different investment types to affect its optimal portfolio allocations. In the rest of this section, we shall study the real estate and infrastructure markets in more detail.

3.3 The global real estate market

Geography   A very useful study of the investable commercial real estate market is Clacy-Jones and Teuben (2014), who estimate the size of the global, investable real-estate market to be \$6.8 trillion as of the end of 2013. Figure 1, taken from this paper, shows the market sizes in all 33 countries covered by IPD. The top five individual markets are
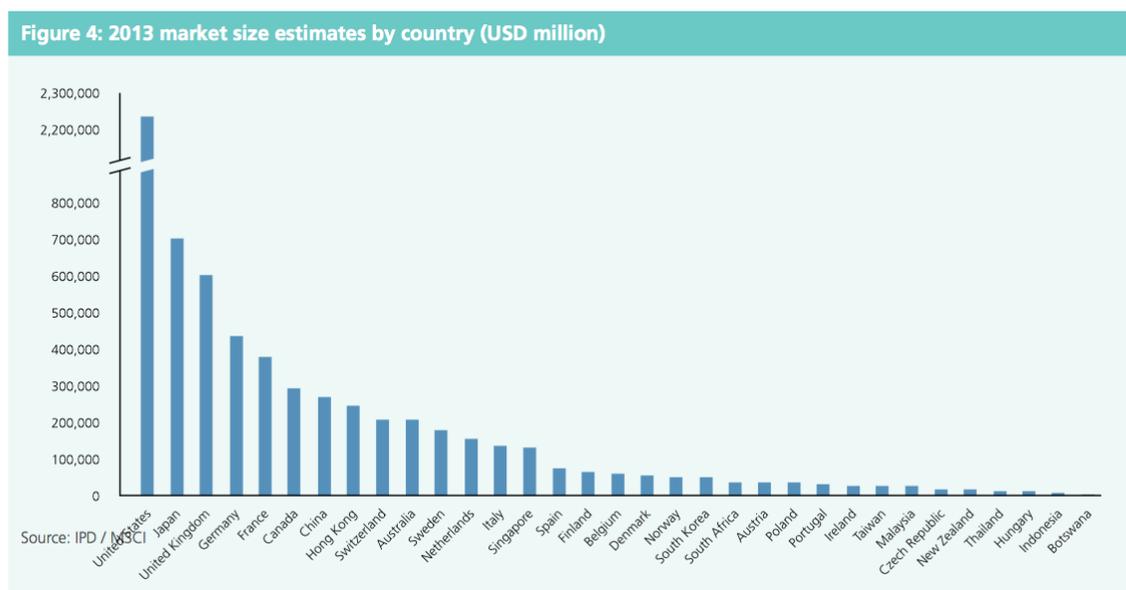
Figure 1: 2013 Real-estate market sizes by country (Figure 4 from Clacy-Jones and Teuben, 2014)

1. U.S. (\$2.2 trillion).
2. Japan (\$0.7 trillion).
3. U.K. (\$0.6 trillion).
4. Germany (\$436 billion).
5. France (\$379 billion).

While there has been some variation, the shares of the largest markets have remained fairly stable over the last decade. Figure 2 shows the geographical breakdown of pension funds' real

15

estate investments, taken from a survey by MSCI/IPD. While pension funds invest globally, they invest most heavily in domestic real estate, especially for funds in the U.K., Australia and the U.S. At the other end of the spectrum, funds in the Benelux countries and Asia invest a higher fraction of their assets internationally, split fairly evenly between the U.S., Europe and Asia. This large home bias seems to leave substantial gains from international real estate diversification on the table.
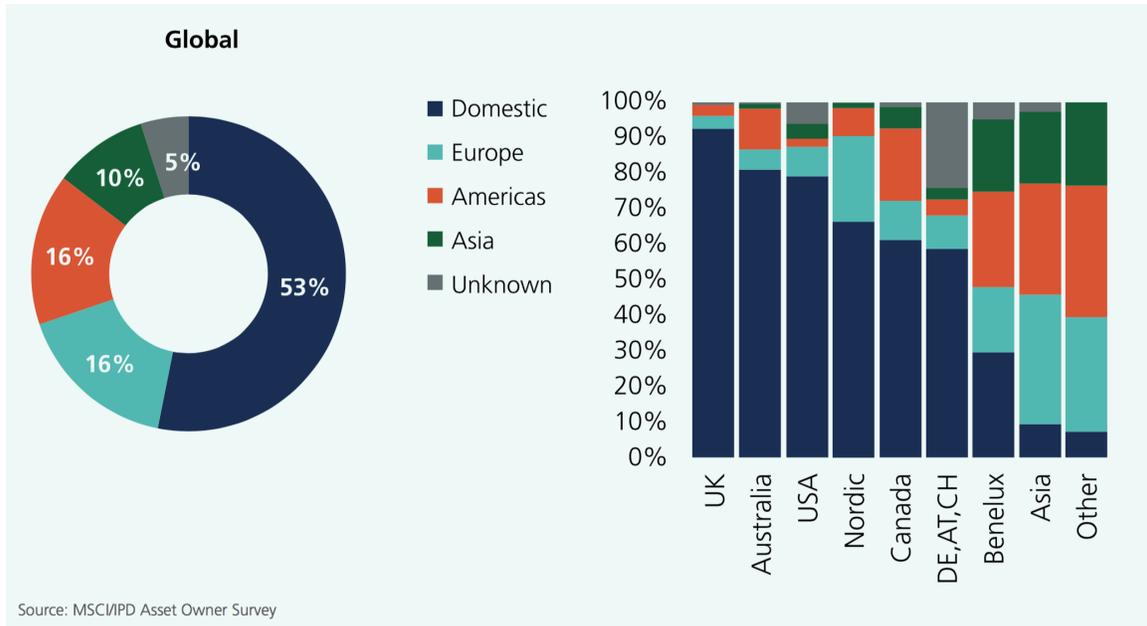
Figure 2: 2013 Real-estate market sizes by country (Figure 2 from Clacy-Jones and Teuben, 2014)

Sector Clacy-Jones and Teuben (2014) estimate that the global real estate investment market breaks down by sector as

- 34% office
- 29% retail
- 18% residential
- 9% industrial
- 10% other

These shares show substantial variation across countries, especially in the amount of investable residential real estate available, which ranges from only a few percent of the total for the UK to over half in Switzerland.

Portfolio type Investment in real estate can be achieved via direct ownership, listed funds (including REITs), and unlisted funds. Figure 3 shows the breakdown of portfolio type by country, again showing great variation. For example, listed investments are a large share of the total in the U.S. and Asia, and much less important in Europe. Figure 4 (Figure 2
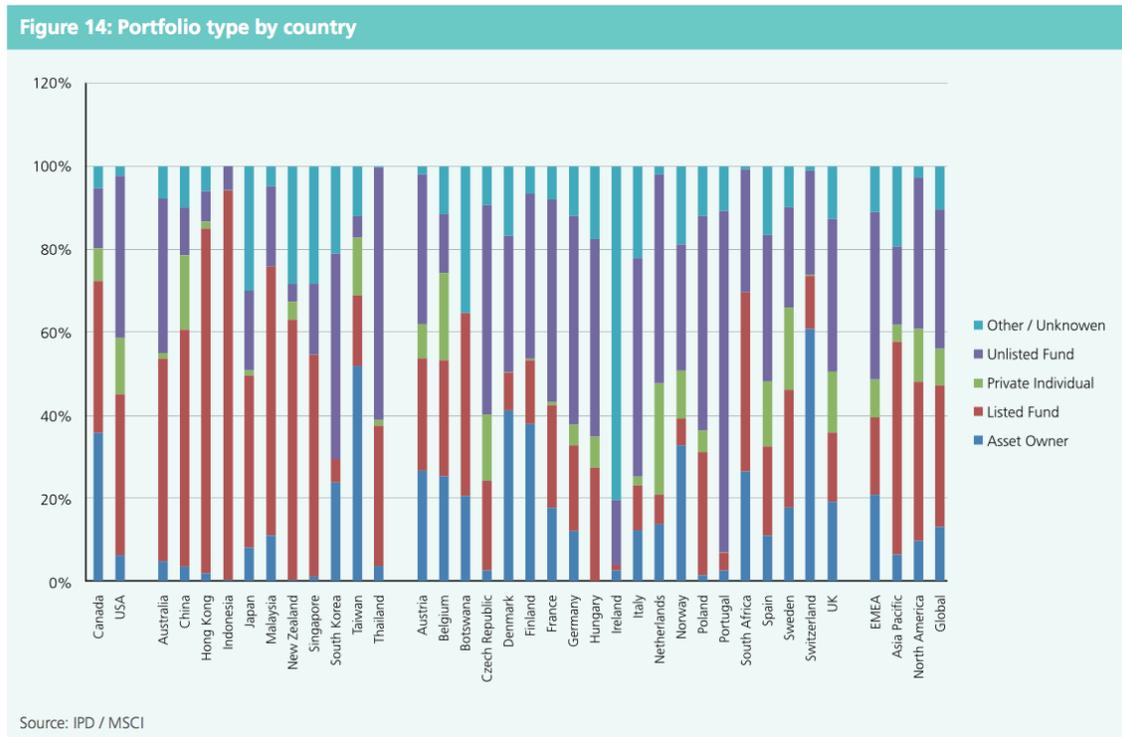
Figure 3: Portfolio type by country (Figure 14 from Clacy-Jones and Teuben, 2014)

from Andonov, Eichholtz, and Kok, 2015) shows real estate holdings by pension funds, split into REITs versus direct investment. Two things stand out. First, the overall proportion of investment in real estate, at 6–7%, is close to the proportion that real estate makes up of the global market portfolio, as discussed above. Second, while the proportion of investment in REITs has increased over time, REITs still make up only a small fraction (12–13%) of total real estate investment by pension funds.

3.4 The infrastructure market

The World Economic Forum (WEF, 2014) defines infrastructure as "the physical structures — roads, bridges, airports, electrical grids, schools, hospitals — that are essential for a society to function and an economy to operate," and notes that there is currently a global investment shortfall in infrastructure of at least \$1 trillion per year. A good overview of the
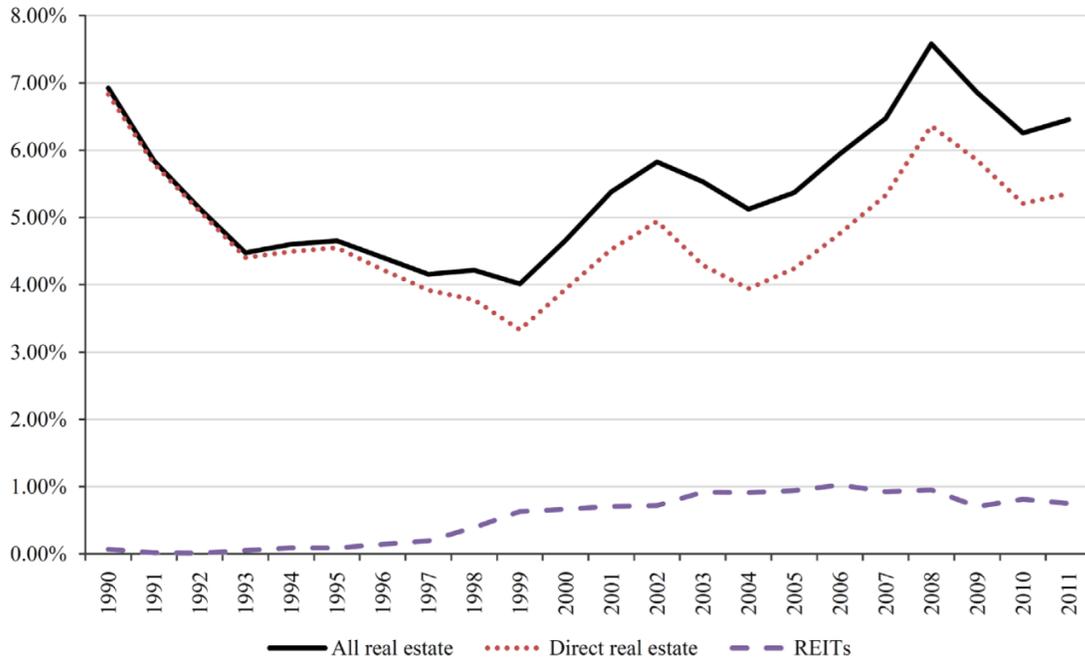
Figure 4: Real estate as a fraction of total pension fund assets (Figure 2 from Andonov et al., 2015)

market is IPD/MSCI (2014) who (using data from RARE (2013)) estimate the total size of all global infrastructure assets in 2012 to be $20 trillion, with 75% government owned and 25% ($5 trillion) privately owned. Of this $5 trillion, roughly $2.5 trillion is held via listed equity and $250 billion via unlisted funds.

Market sectors   The infrastructure market contains many different sectors. Figure 5 (from IPD/MSCI, 2014) presents a taxonomy of infrastructure by asset sector, with description and examples for each sector. Figure 6 (from RARE, 2013) shows the breakdown of the listed and unlisted investable universe in 2012 by sector. Infrastructure investments can be classified in various other ways too, including the degree of regulatory/contractual protection; asset stage (greeenfield/development vs. brownfield/mature); asset use (IPD/MSCI, 2014, contrast "social" investments, such as schools and hospitals, with "economic" investments such as toll roads and transmission towers).

		Definition	Example
SECTOR	Power Generation (Traditional)	Infrastructure that generates energy or power. Excludes infrastructure that falls under renewable energy.	• nuclear reactor • coal-fired power plant
	Power Generation (Renewable)	Infrastructure that sources energy from natural resources e.g. wind, sunlight, rain, tides, etc. Excludes non-renewable forms of power generation, transmission, & distribution.	• wind farms • hydro-power generation • waste-to-energy/biomass • solar farms
	Power Transmission & Distribution	Transfer, distribution and storage of power including oil and gas transportation, distribution and storage. Excludes infrastructure that falls under renewable energy.	• electricity transmission grid • gas pipelines
	Water	Water management infrastructure.	• sewage treatment plants • desalination plants • major irrigations systems • flood control systems • drinking water filtration
	Communication	Facilities consisting of the physical plants and/or equipments for disseminating information.	• broadcast transmission • mobile telephony towers
	Transport	Facilities consisting of the means and/or equipment necessary for the movement of passengers or freight from one point to another.	• toll roads / bridges / tunnels • mass transit systems (metro, bus) • airports & seaports • railways (freight & passenger)
	Public Facilities	Facilities consisting of social infrastructure such as education, health, social housing, convention centres, public car parks, student accommodation	• medical (hospitals) • educational (schools, universities) • correctional (prisons) • events (convention centres, arenas)

Source: IPD

Figure 5: Infrastructure market by sector (Exhibit 1 from IPD/MSCI, 2014)

Sector weightings of listed and unlisted infrastructure universe
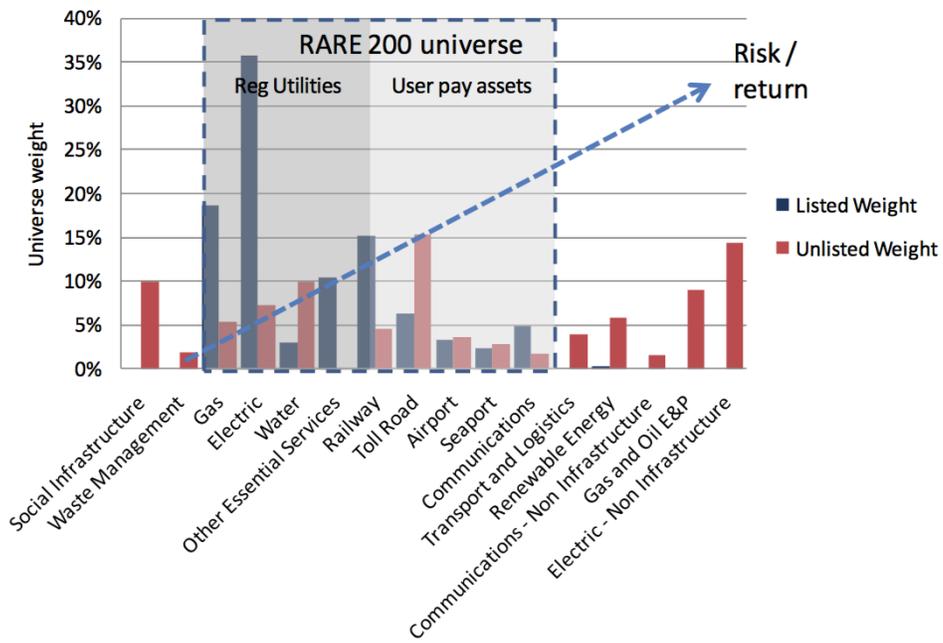
Source: RARE, Factset, Preqin12 years data to 2012

Figure 6: Sector weightings of listed vs. unlisted infrastructure investment (From p. 4 of RARE, 2013)

4   Investing in private RE and infrastructure

Besides investing in listed real estate (either REITs or shares in real-estate companies), another way to gain exposure to real estate is via private equity investment, which can take the form of direct investment, investment in unlisted funds, or investment in a fund-of-funds vehicle. In addition to deciding how much of a portfolio to invest in real estate, GPFG therefore also needs to consider how that allocation should be split between private and public real estate. Estimates differ (see Section 3), but while listed real estate holdings are significant, the "average" investor has 75–85% of its real estate holdings in private investments.

Of course, to fully address the public/private split for GPFG, we need to compare the return characteristics of private versus public real estate, and also to take into account the specific features of the Fund that might make it more suitable than the average investor to invest in either public or private real estate.

The standard series used to study private and listed real estate returns are the NAREIT REIT index (listed) and the NCREIF appraisal-based index (private). Before drawing any definitive conclusions from comparisons of these indices, it is important to note that there are some significant data problems that affect this comparison. In particular,

1. While we can observe daily, transaction-based prices for REITs, privately held real estate trades infrequently. As a result, the NCREIF index is based on appraised prices, which tend to exhibit significant smoothing, serial correlation, and lags relative to REIT returns (see, for example, Geltner, 1991, 1993; Ross and Zisler, 1991).
2. While NCREIF returns are unlevered, REIT returns are calculated for levered equities.
3. The mix of property types may differ between the two indices.

A sizable literature has attempted to control for some or all of these issues. Geltner (1991, 1993) and Ross and Zisler (1991) first showed how to correct for smoothing of prices,[5] and their work has been extended by others, including Shepard, Liu, and Dai (2014), who use Bayesian methods to explicitly handle uncertainty and time-variation in Geltner's smoothing parameter, λ, as well as correcting for leverage and the differences between property types in the respective indices.

---

[5]Assuming the smoothed returns, s_t, follow the process,

$$s_t = (1 - \lambda)r_t + \lambda s_{t-1},$$

where r_t is the true return and λ is the *smoothing parameter* (assumed known), true returns can be obtained by inverting this expression to obtain

$$r_t = \frac{s_t - \lambda s_{t-1}}{1 - \lambda}.$$

In addition to these data issues, another big difference between direct and indirect investment in real estate is that REIT shares, like closed-end funds, often trade at prices quite different from their net asset value (NAV). This premium (or discount) reflects the capitalized benefits of the REIT over holding the assets directly (e.g., liquidity, managerial ability, taxes, funding advantages, lower management costs) minus costs (e.g., management fees, agency costs).[6] Figure 7 shows the average REIT premium from 1990–2015, obtained from Green Street's Web site on October 6, 2015. Over this period, REIT shares have traded as low as 40% below NAV and as high as 30% above NAV, so there is a fair degree of variability in the level of the premium (or discount) over time. There are a few important things to
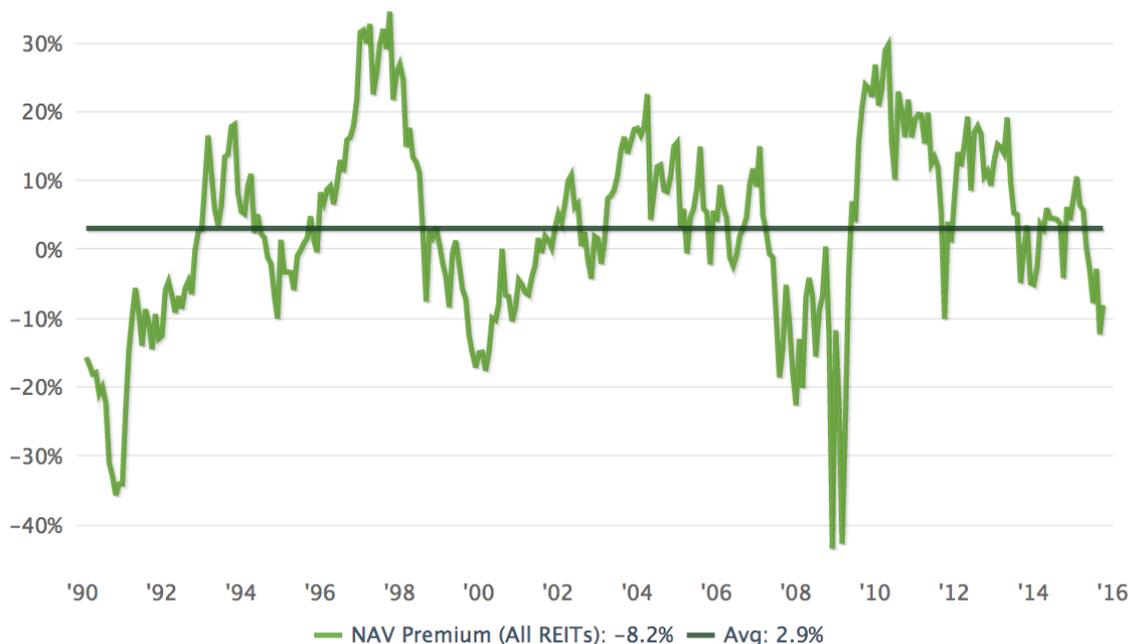
Figure 7: REIT premium over NAV, 1990–2015 (source: Green Street Advisors)

note. First, the mere existence of a premium (or a discount) does not necessarily imply anything about returns. Indeed, if the premium always stays constant, the return on the REIT is identically equal to that on the NAV. Second, just as discussed above, we do not actually know the true NAV for a REIT; the NAV's used to plot Figure 7 are estimates, based on a model of Green Street Advisors.

---

[6]For a discussion of these issues in the context of closed-end funds, see, for example Berk and Stanton (2007); Cherkes, Sagi, and Stanton (2009).

4.1 Returns on public versus private real estate

Figure 8 shows quarterly REIT returns (NAREIT) from 1978Q1–2015Q1, together with the appraisal-based NCREIF Property Index (NPI) over the same period. Figure 9 shows
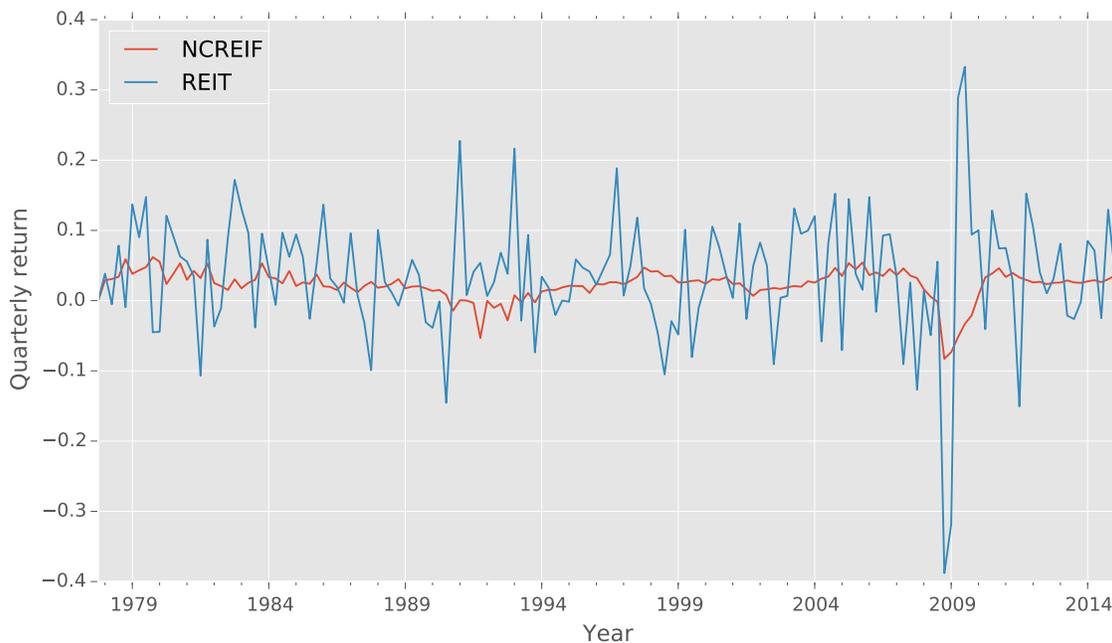
Figure 8: Quarterly returns on REITs versus NCREIF Property Index (NPI), 1978Q1–2015Q1.

cumulative returns on the same series. It can be seen from the figures that total returns on REITs have been substantially higher over this period. One dollar in 1978Q4 would have grown to \$27.03 by 2015Q1 if invested in NCREIF properties, compared with \$95.76 in REITs, a value over 3.5 times as high. Returns on REITs also appear substantially more volatile; the annualized volatility of quarterly REIT returns is 17.7%, compared with only 4.3% for NCREIF returns.

NCREIF versus IPD returns One drawback of using the NCREIF return series is that it is a U.S.-only measure. For a more international comparison, we could in principle use the IPD Global series, but this is only available annually and only since 2001. Figure 10 shows annual NCREIF returns versus IPD Global returns from 2001–2013, and it can be seen that the two series exhibit very similar behavior. As a result, our statistical analysis (like the vast majority of the academic literature) will use NCREIF data.
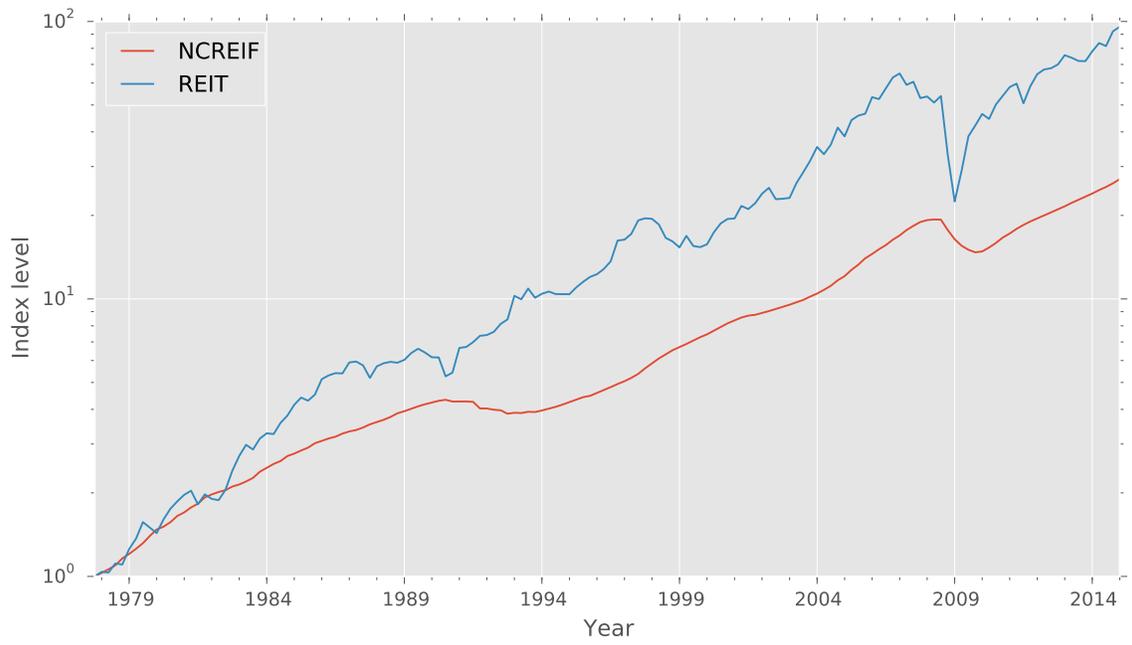
Figure 9: Cumulative returns on REITs versus NCREIF Property Index (NPI), 1978Q1–2015Q1.

Figure 10: NCREIF versus IPD annual returns, 2001–2013

24

Expected returns   Expected returns comparisons are complicated by the data issues described above, especially property-type mix and leverage. However, even after correcting for these issues, the majority of the literature concludes that expected returns on listed real estate exceed those on unlisted real estate. Riddiough, Moriarty, and Yeatman (2005) "...propose an alternative approach that involves adjusting the characteristics of assets constituting an index or portfolio to match the asset characteristics of a reference index or portfolio. This approach is applied to commercial real estate, where we create an index of REIT returns to compare to the NCREIF index. To enhance comparability, return indices are adjusted for partial-year financial data, leverage, asset mix and fees. Adjusted results over a 1980–1998 sample period show general convergence between the indices, although an annual return difference of over three percentage points remains in favor of public market asset ownership." Tsai (2007) and Pagliari, Scherer, and Monopoli (2005) find similar results, as summarized by Tsai (2007): "In contrast to the author's 2.66% difference in mean returns between public and private markets from 1987 to 2005, Riddiough et al. (2005) reported a 3.08% difference over 1980 to 1998 and Pagliari et al. (2005) found a 3.00% difference during 1981 to 2001." Most recently, using more recent data, Ling and Naranjo (2015) report that "Unconditionally, we find that passive portfolios of unlevered core real estate investment trusts (REITs) outperformed their private market benchmark by 49 basis points (annualized) over the 1994–2012 sample period." On important caveat in interpreting all of these results is that they do not deal with risk-adjusted returns. It is therefore possible that the different expected returns found by these authors merely reflect different degrees of risk.

Volatility   As mentioned above, the annualized volatility of quarterly NAREIT index is 17.7%, compared with only 4.3% for the NCREIF index. However, this comparison is affected both by the leverage of REITs and by the smoothing of prices in the NCREIF series. Figure 11 shows annualized volatilities for returns calculated over different horizons. Both volatilities increase with horizon up to 28-quarter returns, then REIT volatility starts to decline while the volatility of NCREIF returns continues to increase all the way to a horizon of 40 months, as the effect of smoothing becomes less and less significant. At this horizon the estimated volatilities are much closer: 19.2% for NCREIF versus 25.1% for REITs.[7] Shepard et al. (2014) find, after adjusting for smoothing, leverage and portfolio composition, that the total volatility of public and private real estate is very close.   Other authors find similar results. For example, Pedersen, He, Tiwari, and Hoffmann (2012) report that, after correcting for smoothing, the volatility of private real estate is about 2.5 times higher than implied by unadjusted volatility estimates.

[7]Of course, the number of independent observations is rather small at this horizon.
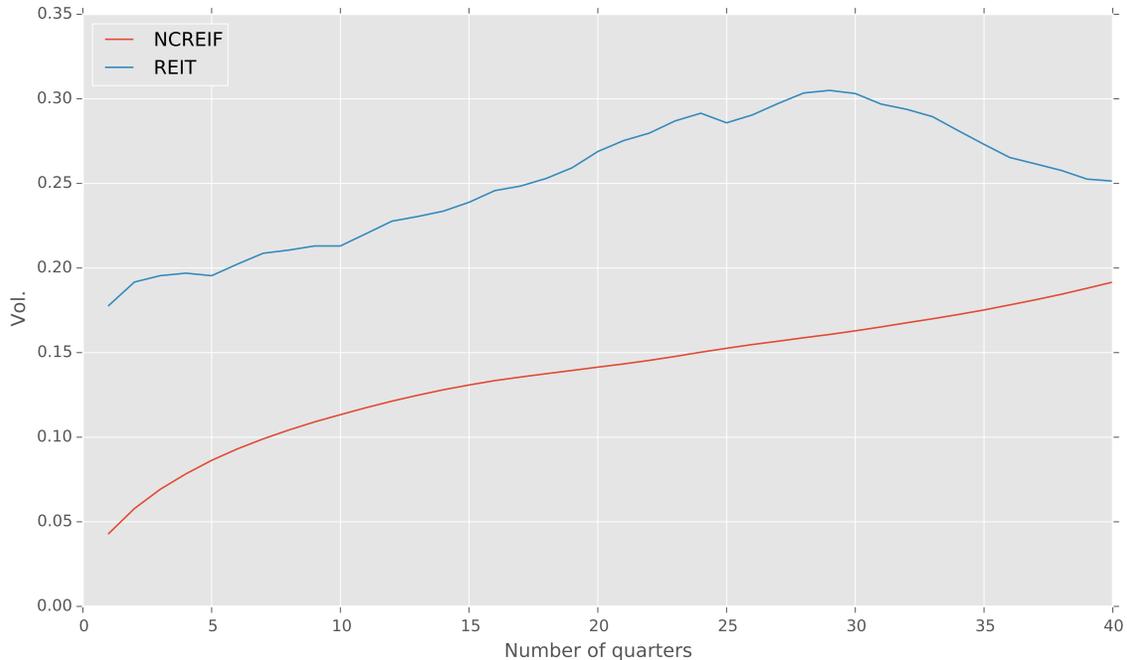
Figure 11: Volatility of returns on REITs versus NCREIF Property Index (NPI), 1978Q1–2015Q1.

Correlation   The sample correlation between quarterly NAREIT and NCREIF returns from 1978–2015 is only 0.152. Taken at face value, this seems to show that private and public real estate investment are very different, and that there may be substantial scope for diversification gains by investing in both. However, as with volatility, these calculations are affected by the data issues discussed above, especially smoothing. A rather different picture starts to emerge when we look at Figure 12, which shows correlations between NCREIF and NAREIT returns calculated over different horizons. The correlation increases almost monotonically with horizon from 0.152 at a quarterly horizon to a maximum of 0.653 for 33-quarter returns.[8]   The marked difference between short and long-run correlation has been noted by others (see, for example, Giliberto, 1990; Geltner and Kluger, 1998), and is consistent with the data issues noted above — in particular, if there is smoothing in one of the series, this will affect short-horizon correlations more than long-horizon. Shepard et al. (2014), after controlling for smoothing and leverage, find that estimates of the correlation between public and private real estate returns over a quarterly horizon increase to around 80% for many countries, including the U.S. and U.K., suggesting that, once measurement errors are controlled for, public and private real estate are close substitutes (see Case, 2015, for a heated presentation of this point).   Pedersen et al. (2012) obtain similar results,

---

[8]Of course, the number of independent observations is rather small at this horizon.

26
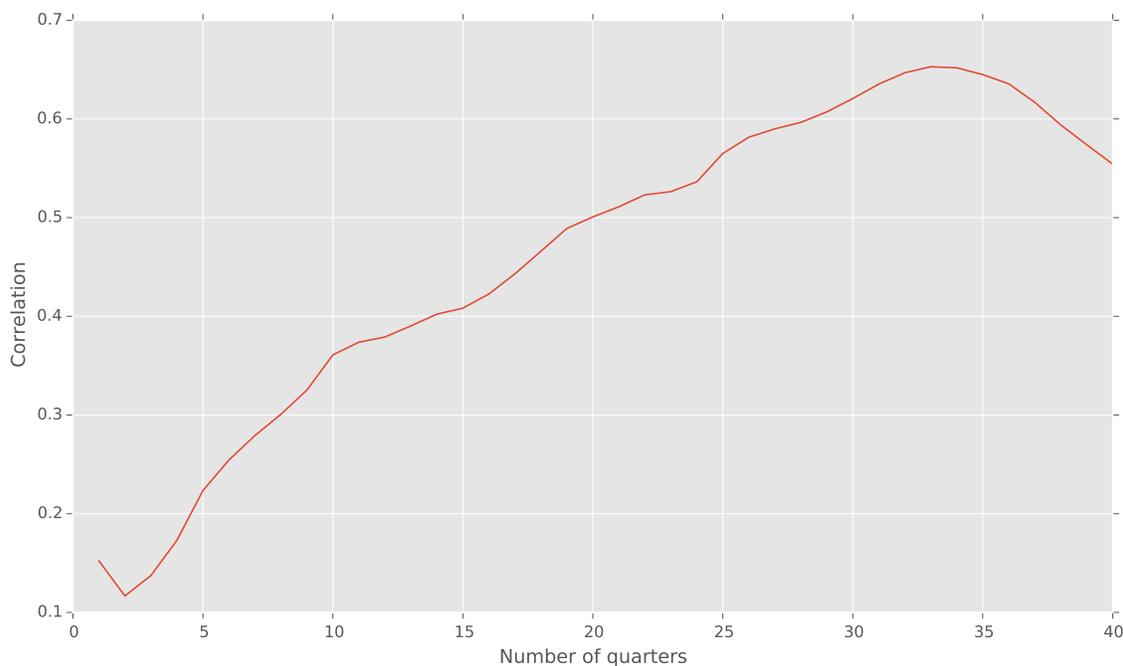
Figure 12: Correlation between returns on REITs versus NCREIF Property Index (NPI), 1978Q1–2015Q1.

finding that after correcting for smoothing, "the implied return correlations between public REITs and private real estate investments fell in the range of 60% to 80% from January 1989 through June 2011."

Conclusions of the academic literature   The finding that the correlation between listed and unlisted real estate, once data issues are corrected for, approaches one, has led many authors to conclude that listed and unlisted real estate are substitutes for one another, especially over long horizons. Oikarinen, Hoesli, and Serrano (2011) study the long-run relationship between securitized (NAREIT) and direct (NCREIF) real estate total return indices using a cointegration framework (see, for example, Engle and Granger, 1987; Johansen, 1988), and conclude that, "[d]ue to the tight long-run interdependence, the longer the investment horizon is, the greater the degree of substitutability between REITs and direct real estate in a mixed-asset portfolio. In other words, the correlation between NAREIT and NCREIF returns approaches one as the investment horizon lengthens . . . [S]ince the two real estate indices are cointegrated with one another and not with the stock market, REITs are likely to bring similar long-term diversification benefits to a stock portfolio as direct real estate." Ang, Nabar, and Wald (2013) conclude that "Over the full real estate cycle, the effects of these different innovation exposures largely disappear, and both public and private vehicles

27

exhibit similar characteristics." Other papers to conclude that public and private real estate investments can be regarded as close substitutes over long horizons include Kutlu (2010); Bond and Chang (2013); Boudry, Coulson, Kallberg, and Liu (2012); Hoesli and Oikarinen (2012); Stefek and Suryanarayanan (2012); Yunus, Hansz, and Kennedy (2012).

Despite the large volume of research that finds public and private real estate to be roughly equivalent over the long term, many papers (including many already mentioned) note that private and public real estate can behave quite differently over the short and medium term. For example, Ang et al. (2013) note that "[o]ur finding that both REITs and private real estate investments have different, idiosyncratic components further suggests there may be a short- and medium-term diversification benefit to holding both in an institutional portfolio;" Hoesli and Oikarinen (2013) comment that "...in the short run the observed REIT and direct real estate returns can substantially deviate from each other due to factors such as data complications, market frictions, and slow adjustment to changes in the fundamentals in the private market ..." Oikarinen et al. (2011) conclude that "[i]n the short-run, the diversification benefits of REITs and direct real estate may differ substantially, however." There are several differences between direct and indirect real estate returns over shorter horizons, notably the fact that, due primarily to greater liquidity and informational efficiency, the securitized real estate market tends to lead the direct real estate market, even after controlling for leverage and appraisal smoothing, which can make private real estate returns more predictable in the short run than those on listed real estate (see, for example, Gyourko and Keim, 1992; Myer and Webb, 1993; Li, Mooradian, and Yang, 2009; Oikarinen et al., 2011; Geltner and Kluger, 1998; Pagliari et al., 2005).

4.2   Real estate performance of institutional investors

Size

One significant difference between GPFG and other investors is size. In line with the proportion that real estate represents of the global market portfolio, Andonov, Kok, and Eichholtz (2013) note that from 1990–2009, pension funds' real estate holdings increased substantially to more than US$320 billion, of which 75% was in direct real estate. This weighting towards private real estate is even heavier for large funds; Pagliari et al. (2005) report that about 90% of real estate investment by the top 25 pension funds is private. Dyck and Pomorski (2011) note similar patterns for other alternative investments, noting that "[l]arger plans devote significantly more assets to alternatives, where costs are high and where there is substantial variation in costs across plans." Andonov, Bauer, and Cremers (2012), note that pension funds increased their exposure to real estate, private equity, hedge funds, infrastructure, and

commodities from 9% in 1990 to 16% in 2010, while university endowment funds increased their allocation to alternative assets from 7% in 1989 to 19% in 2005 (Brown, Garlappi, and Tiu, 2010).

In the mutual fund literature, diseconomies of scale at the fund level are a common feature of both the theoretical and empirical literatures (see, for example, Berk and Green, 2004; Chen, Hong, Huang, and Kubik, 2004). However, recent research has found the opposite to hold for pension funds (see, for example, Dyck and Pomorski, 2011; Andonov et al., 2015) and university endowments (see, for example, Brown et al., 2010; Lerner, Schoar, and Wang, 2008): larger funds earn higher (net) returns than smaller funds, due to lower fees and also (in some cases) higher gross returns. In particular, Dyck and Pomorski (2011) analyze the performance of 842 pension plans between 1990 and 2008, and find that the largest funds in their sample out-perform smaller funds by an average of 43–50 basis points in net abnormal returns (net returns relative to a plan-specific benchmark) per year overall; for the real-estate component of their portfolios, the difference is 4% per year.

Costs   Dyck and Pomorski (2011) find that between 1/3 and 1/2 of the overall difference in net returns between large and small plans reflects lower costs, primarily due to the greater use of internal management by larger plans,[9] which leads to substantial cost savings. In their sample, they find that internal management costs overall are about one third of external management costs,[10] but for the (mostly fixed) costs of internal management to make sense, a fund has to be of sufficient size. Dyck and Pomorski (2011) note, in particular, that "[t]he greatest organizational challenge in setting up an internal group is likely that of assembling, rewarding, and monitoring an internal private equity or real estate team. Up to the 4th size decile there is no internal active management of private equity, and it only becomes sizable in the 9th and 10th size deciles." These results are supported by Andonov et al. (2012) and Andonov et al. (2015), who conclude that "doubling the size of a real estate mandate reduces the annual costs by 10 bps ... A fund that invests internally has 29 bps lower investment costs than a fund that invests through external managers." Other authors have found a similar relationship between size and costs in other asset classes (see, for example, Bauer, Cremers, and Frehen, 2010, on equities).

Gross returns   Dyck and Pomorski (2011) also find large differences in *gross* returns be-

---

[9]Large plans manage 13 times more of their active assets internally (2.7% in the 1st quintile versus 35.4% in the 5th quintile).

[10]From p. 21: "External active managers are 2.6 to 3.5 times as expensive to use as internal managers for equities, 1.7 to 5.1 times more expensive in fixed income, and 3 to 6.9 times more expensive in alternatives, with broadly speaking more substantial savings from internal management for larger plans."

tween large and small plans, particularly for private equity and real estate. Specifically, they estimate that going from the 1st to the 5th size quintile would result in an increase in gross returns of 2.66% for real estate. Andonov et al. (2015) (using data from the same source) also find that larger pension funds earn significantly higher gross returns on their real estate investments. However, Fisher and Hartzell (2013), in a study of several hundred real-estate private-equity funds, find no relationship between a fund's IRR and its size, holding the fund vintage constant.

Possible explanations As possible explanations for these differences in gross returns, Dyck and Pomorski (2011) suggest (p. 30) that "[t]his is most easily interpretable as a superior ability to screen and monitor private equity and real estate funds or being provided superior access to the best opportunities. Larger plans may also have more clout with policy makers, which may help funds in regulatory arbitrage or, say, in winning contracts...Real estate funds, like private equity funds, often offer co-investment opportunities where skill is important." On pages 7 and 8, they also note that "[t]here might also be scale economies on the return side if larger plans are given special access to attractive deals, are able to attract and retain more skillful managers, or are treated differently from other investors and granted special co-investment opportunities or contractual protections...Taking advantage of the potential scale economies requires the ability to hire and retain the right staff, the ability to provide the right incentives, and effective oversight structures that are focused on risk and performance rather than political and other factors." Similarly, Andonov et al. (2015) conclude that "pension funds investing internally in real estate outperform those funds that delegate the investment management. Moreover, investing in real estate through fund-of-funds results in substantial underperformance (around 300 bps per year) ...Larger funds seem to have better skills, which enable them to select better properties when investing internally, and to select better investment managers when investing externally. When investing through external managers, larger funds are likely to get preferential treatment, have greater monitoring capacity, and may have access to better investment opportunities at lower cost." In support of these suggestions, Da Rin and Phalippou (2014) confirm that size correlates strongly with co-investment invitations and other special deals for an investor.

Caveats While these suggested explanations all make intuitive sense, it is important to note that the results quoted above on excess returns do not control for the type of real estate invested in (and their corresponding risk/return characteristics), so it is possible that they stem from funds of different sizes investing in different types of real estate, which can

have very different return characteristics.[11],[12]

Heterogeneity in Skill

Even if the average investor in unlisted real estate investments does not outperform an appropriately risk-adjusted return benchmark, it is still possible that a small set of highly skilled managers does outperform. Studying a large cross-section of U.S. real-estate private-equity funds, Fisher and Hartzell (2013) and Ang, Chen, Goetzmann, and Phalippou (2014b) find that there are substantial differences in performance across funds, even holding fixed the vintage and investment style of the funds. Such heterogeneity arises naturally from the heterogeneity in real estate projects such funds undertake, and from the fact that each fund only undertakes a handful of such projects (with limited diversification). While a group of funds outperforms a public benchmark (for example the U.S. NAREIT index) in each year, the key question is whether it is the same set of funds that consistently outperforms after risk-adjustment. We are not aware of a study that answers this question.

Outside the U.S., Cremers and Lizieri (2015) also find significant heterogeneity in the returns among funds. They look at data on 256 UK real-estate funds between 2002 and 2011 from IPD. They find that more active portfolios (defined as portfolios whose weights differ more from those of the index) outperform less active portfolios, and that this is not driven by fund size, beta or volatility. There is some persistence in a fund's measure of activeness. However, the authors do not study how this persistence in active management translates into persistence in returns.

Illiquidity

Another major difference between GPFG and many other investors is GPFG's longer horizon, which could, in principle allow the Fund to capture any additional returns associated with holding illiquid investments for a long period. Clearly, private real estate investments are much less liquid than public. A REIT is a publicly traded stock, so buying or selling can be done very rapidly and with minimal cost. It is important to note, however, that most REITs are small stocks, and do not have a large trading volume. Greater reporting requirements

---

[11]For example, Figure 5 in Preqin (2015b) shows returns on four different real-estate strategies from 2008–2104 ("overall", "value added", "opportunistic" and "debt"). Even ignoring debt (the highest-performing of the strategies), the difference in the final investment value between the value-added and opportunistic strategies was about 20% of the initial amount invested.

[12]It is also important to note that Dyck and Pomorski (2011) has not been published, and has thus not been through the academic review process. In addition, in private conversation, we have been alerted to the possibility that the cost numbers reported in the CEM data may be substantially understated, which would have a significant impact on these results.

may also cause a greater focus on short-term performance for REITs than for private real-estate investments. Direct real estate investment, by contrast, is relatively illiquid, with large costs associated with each transaction. As a very long-term investor, GPFG is in a good position to capture any additional returns associated with such investments.[13] However, the existence of a liquidity premium in real estate is controversial. In particular, as reported above, after adjusting for leverage, smoothing and asset mix, most of the academic literature concludes that the return on listed real estate actually exceeds that on private real estate.

Conclusions

In conclusion, listed assets offer better transparency and liquidity, and more straightforward risk management. Unlisted investments may offer greater potential for value creation and do not suffer from managerial short-termism that may affect listed investment. Direct control over the asset may increase the potential to control the operational performance of the asset and reduces the number of intermediaries who may have conflicting incentives. The academic literature concludes that both types of real estate have the same characteristics, on average, in the long-run. There is anecdotal evidence that the largest investors may display skill in project selection and/or enjoy access to better unlisted investment opportunities. However, the academic evidence on this topic is too scant and has too many methodological issues related to measurement of real estate returns and lack of risk adjustment to warrant firm conclusions.

4.3   Infrastructure

The choice between direct and indirect infrastructure investment is also an important one, though a topic on which there has been much less academic and non-academic research, in large part a reflection of the relative lack of data. As we saw above, the relative sizes of the listed and unlisted infrastructure markets are approximately the opposite of those for listed/unlisted real estate, with the listed market approximately ten times as large as the unlisted. Pension funds do, however, seem to hold a disproportionately large share of unlisted infrastructure investments. OECD (2014a) reports large pension funds holding 1.6% of their assets in unlisted infrastructure, compared with the roughly 0.25% this sector makes up of the global market portfolio (see Section 3).   Pension funds in the OECD survey (see Table 8) exhibit wide variation both in the fraction of total assets invested in any form of infrastructure (ranging from zero for about half of the funds to 21.4% for the Portuguese

---

[13]For a discussion of the effects of liquidity on asset returns, see, for example, Amihud and Mendelson (1986); Amihud, Mendelson, and Pedersen (2005).

Banco BPI Pension Fund (2), which had another 21.4% of assets invested in infrastructure debt) and in the allocation between listed and unlisted infrastructure equity (for example, the Ontario Municipal Employees Retirement System (OMERS) reports holding 14.9% of its asset in the form of unlisted infrastructure equity and 0.0% as listed equity, while the fourth Swedish National pension fund (AP4) holds 7.0% of its assets in listed infrastructure equity and 0.0% in unlisted).

Unfortunately, the lack of good infrastructure data means that there have been no serious comparisons of correlations, expected returns and variances, or attempts to correct for smoothing, asset type, and leverage, as discussed in the real-estate section above. One of the few data sources on unlisted infrastructure is the MSCI Global Infrastructure Asset Index, though this only has data going back to 2009 (the index was only made public in 2014). De Francesco, Doole, Hobbs, McElreath, and Sharma (2015) compare this index with the (listed) MSCI World Infrastructure Index, and find that in the five years since its inception (2009–2014), the unlisted index has outperformed the listed index with an annualized return over the last five years of 14.0%, compared with 9.6% for the listed index. However, they note that the sectoral composition of the two indices is quite different (as we saw earlier in Figure 6). When they create a listed proxy for the unlisted index by picking a portfolio of listed stocks to match the sectoral composition of the index, they find that much (though not all) of the difference in returns over the period disappears, as shown in Figure 13. They comment that bias in valuations may also partly account for the difference, but that it is too soon to be able to address this issues with any degree of precision.

The discussion on costs of external versus internal management of infrastructure portfolios echoes that for real estate. In a 2012 report, CPPIB estimates that "the fees and expenses for external management of an infrastructure portfolio the size of ours would cost approximately 10 times more than our all-in internal costs." Depending on the geographic footprint of the infrastructure portfolio, setting up an effective organization with many boots on the ground may require a substantial expansion in staff. On the other side of this argument, a large internal staff represents a costly-to-reverse investment, which reduces the flexibility of the organization.[14]

Additional challenges of direct infrastructure investment   Direct infrastructure investment subjects the Fund to a number of sources of risk that are not present with listed investments. In addition to construction and development risk, operational, demand, and market risk, financial and interest rate risk, the Fund would face: political risks (e.g., changes in government or infrastructure policies, shifting popular sentiment towards privatized na-

---

[14]We thank Ludo Phallipou for this argument.

tional services), regulatory risks (e.g., changing energy regulations, no viable PPP legal framework), and management and governance risk (e.g., corruption or expropriation). These risks go a long way towards explaining why only about one percent of institutional investors' assets have been allocated to direct infrastructure investments globally (OECD, 2014a; Inderst, 2013). Another challenge is to find investable projects of sufficient quality and scale. At least some of this risk can be insured. The World Bank, for example, sells sovereign risk insurance, which at least allows one to insure against a part of the risk. Likewise, at least some of the currency risk can be hedged. Local governments can offer debt guarantees or participate in direct equity financing to offset some of the risks that are hard to insure. Finally, co-investment with local partners or investment through a regional infrastructure fund or multilateral development bank (like the International Finance Corporation) may offer additional protection.

Source: MSCI
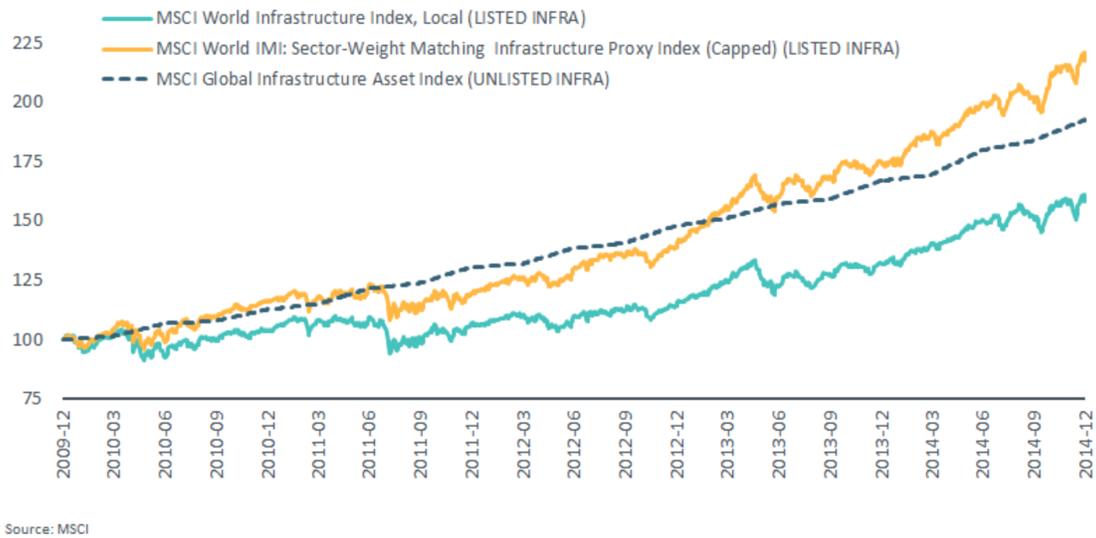
Figure 13: Listed versus unlisted infrastructure returns, Dec. 2009–Dec. 2014 (Figure from page 11 of De Francesco et al., 2015)

5 Risk, return, portfolio allocation, and valuation levels

The previous section has argued that both risk and return on publicly and privately held real estate and infrastructure assets are similar to one another. This similarity applies after both investments are compared on an unlevered basis, after controlling for asset type and geographic location, after clearing up measurement issues with direct asset performance, and when judged by standard statistical methods. It is a similarity that holds more strongly in the long-run, not in the least because there is an opportunity to take private companies public and vice versa. The similarity does not mean that there are no assets or investors that ever outperform by investing in private assets. It also does not mean that there are no gains from diversification in private assets. Especially in real estate, where direct assets comprise 85% of the market place, it is hard to imagine that all investment opportunities are equally available in public and private markets. This argument is less strong or even absent for infrastructure, where most of the private market currently consists of public assets. And it does not mean that there are no diversification benefits from combining direct and public assets in a portfolio in the short- to medium run. However, the evidence supports using publicly listed space as a good proxy, a *pars pro toto*, for the overall real estate and infrastructure markets.

In this section we investigate the first question in the mandate, regarding the target allocation to real estate and infrastructure. Specifically,

- Should the fund change its target allocation of 5% to real estate?
- Should it increase its allocation to infrastructure beyond what is already included in the benchmark equity index?

To answer these questions, in this section we study historical risk and return measures of real estate and infrastructure alongside those of stocks and bonds. The allocations to real estate and infrastructure asset classes that the analysis prescribes can be interpreted as the over-allocation beyond what is included in the equity (and bond) portfolio(s).

Because of the Fund's global investment mandate, we study global real estate and infrastructure returns. We use the longest available history to harness the best possible statistical power. Nevertheless, the available return histories are relatively short, at least by comparison with the overall stock and bond markets. This is a caveat that applies to any study of risk and return in real estate and infrastructure, whether its focus is on public or private markets. To partially mitigate this drawback, the appendix studies the United States, the largest real estate market in the world, over a longer history. We report return summary

statistics (means, volatility, and correlations), factor-model regressions, and a mean-variance portfolio analysis. The final part of this section is an analysis of the current valuation levels of real estate and infrastructure assets.

5.1 Global real estate analysis

We start by studying real estate and turn to infrastructure in the second part of our analysis.

5.1.1 Data

Global real estate data   For real estate, we use the MSCI World Core Real Estate index. As noted by MSCI (2014): "The MSCI Core Real Estate Indexes, based on the MSCI ACWI Investable Market Indexes (IMI) (the 'Parent Index'), are designed to reflect the performance of stocks in the Parent Index engaged in the ownership, development and management of specific core property type real estate. Specifically, these indexes exclude companies that do not own properties. For example, companies active in real estate services and real estate financing are not included in the MSCI Core Real Estate Indexes." REITs that are not involved in core real estate property (such as timber or infrastructure REITs) are not included; neither are mortgage REITs. The index focuses on mid- and large-cap stocks from 23 countries. The return series is available at monthly frequency from December 1994 onwards.[15] We refer to this index with the abbreviation CREI. All returns in this study are expressed in U.S. dollars. This makes the analysis more straightforward but may induce some additional correlation between the various asset classes.

Global stock and bond data   We use the global market portfolio from Fama and French (2012) as our global stock market benchmark. From the same source, we also obtain data on a global size, value, and momentum factor.[16]   We use the Barclays Global Aggregate

---

[15]As of May 2015, the CREI index included 390 companies in many different subsectors and geographies. The sub-sectors are Real Estate Development, Real Estate Operating Companies, Residential REITs, Retail REITs, Office REITs, Industrial REITs, Diversified REITs, Specialized REITs such as self-storage REITs, and Diversified Real Estate Activities. In terms of country composition that day, 20 countries are represented with the United States representing 49.9% on that day, Japan 13.9%, Hong Kong 9.7%, the U.K. 6%, and Australia 5.5%. Europe ex-U.K. represents 9.3% of the basket.

[16]The global factor include 23 countries: Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Hong Kong, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Switzerland, Sweden, United Kingdom, United States. The global market factor has a correlation of 99.3% with the MSCI All Country World Index equity index. It has an average return of 8.64%, which is 40 basis points higher than the 8.24% on the ACWI. We consider the Fama-French index a better proxy of the world equity market portfolio than the ACWI, which is skewed towards larger stocks. We note that this choice makes it harder for real estate or infrastructure indices to beat the equity market portfolio.

Bond Index as our global bond benchmark. This bond index contains both government and corporate bonds.

5.1.2   Return summary statistics

Table 2 lists summary statistics for the returns over three sample periods: December 1994–June 2015 (247 months), January 1999–June 2015 (198 months), and December 2003–June 2015 (139 months). These different samples are useful to gauge the stability of the returns over longer periods. They will also be useful to study the infrastructure returns below. The various infrastructure indices we consider only have data from 1999 or from 2003 onwards.

Full sample   Starting with the full sample in Panel A, we find that the global core real estate index (CREI) had an average annualized return of 9.8%. This is higher than that of global stocks whose return averaged 8.7%. The higher return came with higher volatility: 17.6% annualized return standard deviation for CREI versus 15.2% for stocks. For comparison, over the same period, U.S. REITs (last column) had even higher average returns of 12.8%, but with still higher volatility of 20.0%.

A performance measure that combines return and risk is the Sharpe ratio. It measures the return on a risky asset over and above a safe risk-free asset return and divides that excess return by the standard deviation of the risky asset return. The Sharpe ratio measures the return per unit of risk. Using the U.S. T-bill rate as the yield on a global safe asset, the Sharpe ratio on global real estate is 0.41. This compares to 0.41 for global stocks, 0.51 for global bonds, and 0.51 for U.S. REITs. Thus, global real estate has had the same performance as stocks on a volatility-adjusted basis, but has under-performed bonds. U.S. REITs have done better than real estate companies outside of the U.S., especially once one considers that the CREI index invests about 50% in U.S. REITs.

Volatility is a symmetric measure of risk that ignores the possibility that there may be more downside than upside risk. Therefore, the table also reports the skewness of the returns. A negative skewness means that large negative returns are more likely than large positive returns. A symmetric return distribution has a skewness of zero. The skewness of global real estate is −1.02, somewhat higher than the −0.83 for global stocks and −0.81 for U.S. REITs. Real estate returns have indeed suffered from large downside risk, especially in the financial crisis. Without the financial crisis period (August 2007–March 2009), the skewness of global real estate returns is essentially zero (−0.04) and that of U.S. REIT returns turns large and positive (+0.85).

What matters for the risk of a well-diversified portfolio is not the volatility of an asset but rather its covariance with the other assets in the portfolio. The covariance of two assets

37

is the product of their correlation and their individual volatilities. With that in mind, the second part of Panel A of Table 2 reports correlations between global real estate, stocks, and bonds, as well as between global real estate and U.S. real estate. First, we find a correlation of the CREI with U.S. REIT returns of 87.4%. This correlation is quite modest considering that U.S. REITs make up about 50% of CREI.[17] This suggests that there are substantial gains from international diversification in real estate. These gains from diversification are also visible in the lower return volatility of the CREI compared with the U.S. REIT index.

Global real estate has a 76% correlation with global stocks and a 39% correlation with global bonds over the full sample. Again these numbers suggest that there may be substantial gains from diversification that can be achieved when combining stocks and bonds with real estate.

Shorter samples   Panel B of Table 2 looks at the more recent January 1999–June 2015 period. This period includes the immediate aftermath of the Asian Financial Crisis and the technology boom and bust. Over this sample, global stock returns only averaged 6.9% over this sample, with a volatility of 15.7%. The Sharpe ratio was only 0.32, compared with 0.41 for the full sample (which includes only 49 more months). In sharp contrast, global real estate stocks suffered very little during this period and posted a strong performance, with average returns of 10.8%. Their Sharpe ratio was 0.48, compared with 0.41 over the full sample. The late 1990s and early 2000s are a great example of a period where real estate stocks offered diversification benefits relative to a stocks-only portfolio. The CREI has a correlation of 77% with global stocks and 43% with global bonds.

Panel C of Table 2 looks at the most recent December 2003–June 2015 period. Real estate posted strong returns in the past 11.5 years as well. The Sharpe ratio is 0.47 in this period. We see a high correlation of global real estate with stocks of 85% and with bonds of 49% over this period. The correlation between global and U.S. real estate stocks increases to 94%. Both the volatility and the negative skewness are larger over this period than over the full sample. The uniformly higher correlations and the higher volatility and skewness in this sample are attributable to the prominence of the global financial crisis.Unlike, the technology boom and bust, the financial crisis started to a large extent with a collapse in U.S. real estate markets with knock-on effects onto banks, and global equity markets.

---

[17]One can calculate that the data imply a correlation between U.S. REITs and the non-U.S. component of CREI of 52.4%. The non-U.S. return component has a volatility of 20.4%.

Table 2: **Global Return Summary Statistics**

Means, standard deviations, and Sharpe ratios are annualized. Skewness is the skewness of monthly returns. The different panels consider different samples.

	CREI	Stocks	Bonds	U.S. REITs	CII	II	EMII
Panel A: Dec 1994–June 2015							
Mean	9.79	8.73	5.39	12.78	–	–	–
Standard Deviation	17.63	15.18	5.52	20.00	–	–	–
Sharpe Ratio	0.41	0.41	0.51	0.51	–	–	–
Skewness	−1.02	−0.83	0.04	−0.81	–	–	–
Correlations							
CREI	1.00	0.76	0.39	0.87	–	–	–
Stocks	0.76	1.00	0.26	0.58	–	–	–
Bonds	0.39	0.26	1.00	0.28	–	–	–
U.S. REITs	0.87	0.58	0.28	1.00	–	–	–
Panel B: Jan 1999–June 2015							
Mean	10.80	6.93	4.28	12.68	–	3.96	10.39
Standard Deviation	18.35	15.72	5.73	21.55	–	14.36	20.71
Sharpe Ratio	0.48	0.32	0.41	0.50	–	0.14	0.41
Skewness	−1.07	−0.73	0.05	−0.81	–	−0.51	−0.34
Correlations							
CREI	1.00	0.77	0.43	0.92	–	0.58	0.63
Stocks	0.77	1.00	0.29	0.60	–	0.82	0.83
Bonds	0.43	0.29	1.00	0.31	–	0.33	0.26
U.S. REITs	0.92	0.60	0.31	1.00	–	0.44	0.45
CII	–	–	–	–	–	–	–
II	0.58	0.82	0.33	0.44	–	1.00	0.75
EMII	0.63	0.83	0.26	0.45	–	0.75	1.00
Panel C: Dec 2003–June 2015							
Mean	10.99	9.13	3.99	12.24	12.05	8.99	12.68
Standard Deviation	20.49	15.54	5.69	24.50	13.43	12.66	18.41
Sharpe Ratio	0.47	0.50	0.47	0.45	0.80	0.61	0.62
Skewness	−1.09	−0.94	−0.10	−0.78	−0.99	−0.80	−0.80
Correlations							
CREI	1.00	0.85	0.49	0.94	0.84	0.76	0.74
Stocks	0.85	1.00	0.44	0.71	0.90	0.85	0.83
Bonds	0.49	0.44	1.00	0.39	0.57	0.54	0.48
U.S. REITs	0.94	0.71	0.39	1.00	0.70	0.61	0.57
CII	0.84	0.90	0.57	0.70	1.00	0.93	0.81
II	0.76	0.85	0.54	0.61	0.93	1.00	0.78
EMII	0.74	0.83	0.48	0.57	0.81	0.78	1.00

5.1.3 Investment horizon

It is important to assess how the risk-return tradeoff changes with the investment horizon. Longer-term risk and return measures are arguably more appropriate for the long-term investment horizon of the GPFG and for long-term investments like real estate and infrastructure.

A key assumption when measuring longer-term performance concerns reinvestment. We assume that all proceeds from an investment are reinvested in the same investment during the investment horizon. We form cumulative returns over investment horizons ranging from 1 to 40 months. By the time we are up to 40 months, we only have 6 independent observations. We think it is not useful to even longer horizons since we do not have sufficient data to make meaningful statistical inference. The left panel of Figure 14 shows mean returns on global real estate, stocks, and bonds as a function of the investment horizon. There is a small increase in the annual return as the investment horizon increases. The variability in the lines, however, shows that this increase is swamped by increasing estimation uncertainty around that mean. In other words, per annum returns are not statistically different for one month or 40-month investment horizons.

The middle panel of Figure 14 uses the annualized average return in excess of a similarly annualized risk-free average return and divides it by the properly annualized standard deviation to arrive at an annualized Sharpe ratio. Because the volatility of real estate and stock returns *increases* with the investment horizon, their Sharpe ratio declines. There is no statistical evidence for the hypothesis that long-horizon Sharpe ratios are higher than short-horizon ones. The Sharpe ratio on bonds is higher than that of both real estate and stocks at all horizons, but is again flat in the investment horizon.

Finally, the right panel shows the correlation between real estate and stocks and between real estate and bonds, measured over 1- through 40-month horizons. The correlation between bonds and real estate is unstable and shows no clear pattern. However, the correlation between real estate and stocks does seem to decrease in investment horizon, from around 0.8 at one-month frequency to around 0.6 at 40-month frequency. This suggests increased diversification potential between stocks and real estate. Again, given the short sample of data (246 months), we only have 6 independent observations on 40-month returns, making it difficult to make conclusive statements.

5.1.4 Time-varying risk

The previous analysis calculated volatilities and correlations over a sample of 21.5 years of data as well as over two subsamples. To investigate the issue of time-variation in volatilities
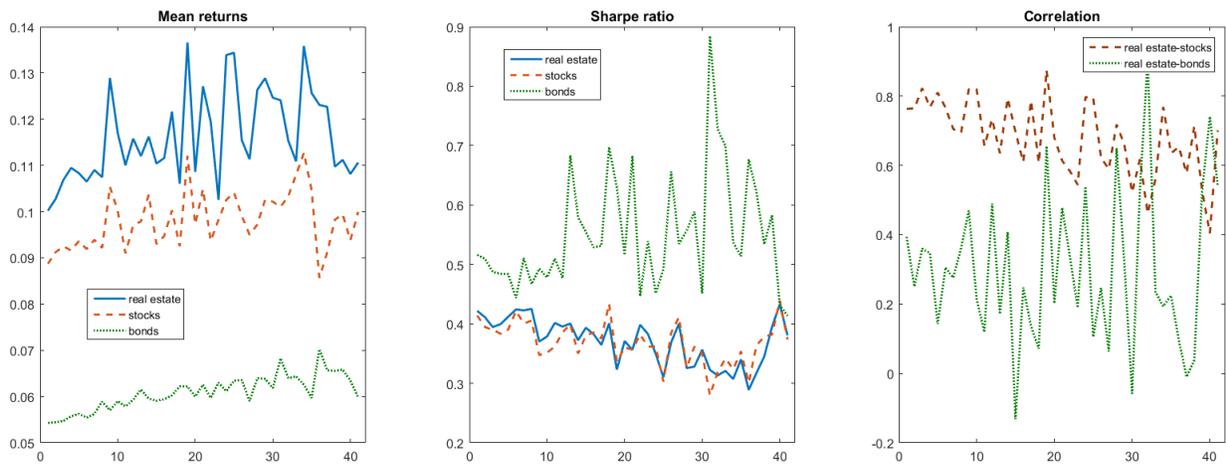
Figure 14: Investment Horizon

The left panel plots the average return on global real estate (CREI), stocks, and bonds as a function of the investment horizon (x-axis, expressed in months). Each average return is annualized. The middle panel divides the annualized average return in excess of a similarly annualized risk-free average return by the annualized standard deviation to arrive at an annualized Sharpe ratio. The right panel shows the correlation between real estate and stocks and between real estate and bonds, measured over 1- through 40-month horizons. The sample period is December 1994 until June 2015.

and correlations further, we compute 60-month rolling window volatilities and correlations. This analysis can shed light on the question whether the correlation between real estate and stock markets is likely to normalize to lower levels as the financial crisis disappears in the rear view mirror.

Volatility   The top left panel of Figure 15 shows that global real estate returns had volatilities that averaged around 14% before the financial crisis. These volatilities increased dramatically during the crisis to levels twice as high. Over the past two years, the 60-month windows no longer include the financial crisis period and real estate volatility has come back down to 15%. The middle panel shows similar volatility dynamics for stocks. Stock volatility is plotted on the same scale as real estate volatility. While stock volatility also doubled from pre-crisis to crisis, it started from a lower level of around 10% and never went above 22%. Bond return volatility, plotted in the right panel against a much smaller vertical axis range (4–8%) also rose during the crisis and fell back down at the end of the sample. These figures illustrate that volatility, even when measured over five-year horizons, can move substantially and in coordinated fashion across asset classes.

Correlation: real estate with stocks   The bottom three panels of Figure 15 show correlations between the three asset classes measured over 60-month windows. The top left panel shows that real estate had a correlation of around 70% with the overall stock market in the late 1990s. The correlation then fell precipitously to 50% because of the technology sector's boom and bust which had little effect on real estate stocks. During and right after the financial crisis, stocks and real estate correlated very strongly with each other, reaching levels above 90%. Over the last two years, the correlation between real estate and stocks has fallen from a high of 93% back down to 80%.

Combining the correlation dynamics with the volatility dynamics, we obtain the covariance dynamics which are an important input to understand the gains from diversification. The covariance between global real estate and stocks increases eight-fold from .07 on the eve of the financial crisis (August 2002–07) to 0.56 in 2012 (August 2007–12). By the end of the sample when the financial crisis is no longer in the 60-month window (May 2009-14), the covariance is back down to 0.16. The covariance is fast approaching the levels seen in the early 2000s (around 0.14).

These dynamics underscore that there can be substantial stretches of time where an asset class fails to offer much gain from diversification. Taking a longer-term perspective, real estate capital markets have become more integrated with overall capital markets over the past two decades. This inevitably means higher co-movements than in the past, somewhat

bounding the scope for diversification.
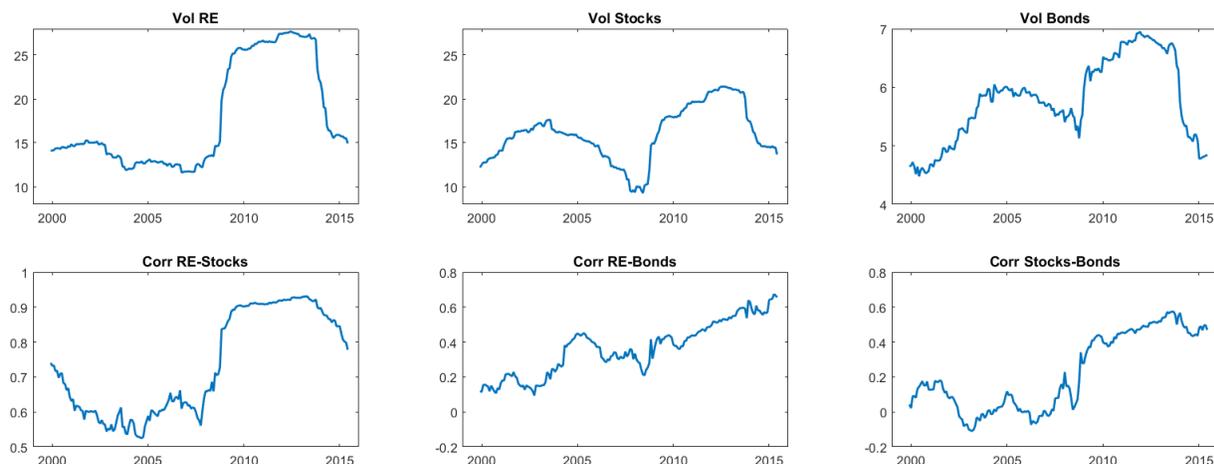
Figure 15: Time-Varying Volatilities and Correlations

The three top panels plot annualized standard deviations of global real estate (CREI), stocks, and bonds. Each observation measures the standard deviation of the past 60 months of returns, multiplied by $\sqrt{12}$. The three bottom panels report the correlations between these return series, also measured over 60-month horizons. The sample period is December 1994 until May 2015. Because 60 months of returns are needed to compute the first point, the graphs start in December 1999.

Correlation: real estate with bonds   The middle panel in the bottom row of Figure 15 shows the correlation between REITs and bonds. Interestingly, this correlation increases steadily from a low of 10% to a high of almost 70% at the end of the sample period. These dynamics mirror the correlation between the overall stock market and bond market in the bottom right panel, plotted on the same scale. The correlation between real estate and bonds is higher than that between all stocks and bonds throughout the sample. The graphs show that real estate has more bond-market risk than overall stocks and that its bond market risk has risen steadily but substantially over time.

The covariance between real estate and bonds quintupled from pre-crisis (.02) to crisis (.10), before falling back to .05 by the end of the sample by virtue of the decline in volatility of bond real estate and bond returns (and despite the rise in correlations). The recent fall in covariance restores some of the diversification benefits between real estate and bonds, but the covariance remains more than twice as high as before the crisis.

Stocks and bonds were great complements in a portfolio over the 1994–2009 period. Their correlation increased to much higher levels since then. The flight-to-safety feature

43

of U.S. Treasuries and the continuing decline in bond yields over the past 7 years (in part due to Quantitative Easing policies) alongside a stock market rally accounts for the rise in correlation. With bond yields at historically low values, and bound to go up, and with stock markets trading at high multiples, it seems conceivable that both stock and bond prices may fall in unison. This would lengthen the regime of positive correlation between stocks and bond.

5.1.5 Factor analysis

Next, we investigate the performance of publicly traded commercial real estate in the U.S. using standard asset pricing factor models. The analysis serves to advance the univariate correlation analysis. In addition to understanding what risks commercial real estate is exposed to, we can investigate whether real estate stocks have displayed abnormal performance (alpha) relative to the factors considered.

The natural starting point is a two factor model which contains a stock market factor and a bond market factor. Stock market and bond market risk are the two key sources of systematic risk, ubiquitous both in the academic literature and in investment practice. They are the only two sources of risk included in the Opportunity Cost Model, discussed at length in Section 6. Therefore, a first important question that the factor analysis answers is: What is the exposure of real estate to stock and bond market risk. And is there outperformance relative to the appropriate combination of stocks and bonds?

We then consider additional asset price factors, notably a size, value, and momentum factor, which have been shown to command positive premia. In recognition of such factor premia, Norges Bank employs factor-harvesting strategies as part of its investment strategy. We ask how real estate returns covary with these factors and whether they have outperformed the appropriate combination of stock, bond, size, value, and momentum factors.

Two-factor model Column 1 of Table 3 studies the two-factor model for global real estate returns over the full sample (December 1994–June 2015). It estimates the following regression:

$$R_t^r - R_t^f = \alpha_r + \beta_r^s(R_t^s - R_t^f) + \beta_r^b(R_t^b - R_t^f) + e_{t,r},$$

where β_r^s measures the exposure of real estate to the stock market factor and β_r^b measures the exposure of real estate to the bond market fact. The regression intercept, α_r, measures the average abnormal excess return on real estate. It is the compensation real estate investors have earned over and above the compensation for bearing stock and bond market risk.

We find that global real estate has a global equity beta β_r^s of 0.82. This beta is measured

very precisely; it has a t-statistic of 11.2.[18] The equity beta of 0.82 implies that for every 1% increase in global stocks, real estate stocks increase on average by 0.82%. Global real estate has a global bond beta β_r^b of 0.68, which is again measured very precisely. Put differently, the opportunity cost of investing $1 in global real estate is $0.82 in global stocks, $0.68 in global long-term bonds, and a short position of $0.50 in U.S. T-bills. This portfolio of stocks, long-term bonds, and T-bills can be thought of as the replicating portfolio for global real estate. Under the opportunity cost model, discussed below, this is the portfolio that would be sold to purchase a dollar of real estate.

Combined, fluctuations in returns on stocks and bonds explain 62.4% of variation in global real estate returns. One-third of variations in global real estate returns are uncorrelated with stocks and bonds. That orthogonal component can be interpreted as a "global real estate factor." The fact that a significant fraction of real estate returns is "unspanned" by stock and bond returns implies that there may be gains from diversification from adding real estate to the portfolio.

The first row shows the abnormal return on real estate, or α_r. This abnormal return is essentially zero. This means that real estate performed in line with its replicating portfolio. It also means that there seems to be no additional compensation for exposure to the global real estate factor.

Finally, the last row of the table reports the expected return on real estate:

$$E[R_t^r] = R_t^f + \hat{\beta}_r^s E[R_t^s - R_t^f] + \hat{\beta}_r^b E[R_t^b - R_t^f],$$

where the hats indicate that we use the estimated value for the betas. This expected return omits the alpha and therefore asks what return an investor in global real estate ought to earn to compensate her for the systematic exposures of real estate to stock and bond risk. The equity risk premium, $E[R_t^s - R_t^f]$, and the bond risk premium, $E[R_t^b - R_t^f]$, are the compensation per unit of stock and bond market risk, respectively. The expected return on global real estate is 9.55% per year. This cost of capital is the sum of a risk-free interest rate of 2.57%, compensation of 5.08% for exposure to stock market risk (the product of a beta of 0.82 and an equity risk premium of 6.17%), and compensation of 1.90% for exposure to bond market risk (the product of a beta of 0.68 and a bond risk premium of 2.82%).

Comparison to the U.S. To help interpret these results, it is instructive to repeat the analysis for the United States, as reported in Column (3) of Table 3. We estimate the same two-factor model regression, but this time the left-hand-side variable is the excess return on

---

[18]Recall that a t-statistic below -1.98 or above $+1.98$ indicates that the coefficient estimate is different from zero with 95% probability, assuming the estimator is normally distributed.

U.S. REITs (data from NAREIT) and the right-hand side variables are the excess return on U.S. stocks (value-weighted return from Ken French) and the excess return on 10-year U.S. Treasury bonds. We see that the stock beta of real estate is a bit lower, at 0.75, and the bond beta is substantially lower (and insignificant) at 0.27. The lower bond beta is in part due to the use of a global bond index that combines government and corporate bonds. It suggests that the pure interest rate risk of real estate — as measured by exposure to government bond returns — is lower than what the global bond beta suggests. The higher global stock and bond betas may also be partially due to expressing all returns in a common currency.

In the U.S., the two factor model explains only half as much of the variation in real estate returns as globally (31.9% R^2 versus 62.4%). This large discrepancy may be due to the aforementioned difference in the bond factor and to the common currency. Another possible explanation is that there are fewer diversification benefits left from adding real estate to a stock-bond portfolio once real estate, stock, and bond portfolios are all globally diversified.

The first row shows that in the U.S., there is a substantial abnormal return to real estate of 26 basis points per month or 3.1% per year. Given the volatility of real estate returns, the alpha estimate is not statistically different from zero. The last row shows a U.S. two-factor expected return for U.S. REITs of 9.66%, similar to the corresponding global number.

Adding risk factors Column (2) of Table 3 adds three well-known asset pricing factors to the pricing model: a global size (smb), a global value (hml), and a global momentum (mom) factor. These factors are constructed by Fama and French (2012) and are based on stock market data for the same 23 countries that are reflected in the global real estate portfolio. There is a long tradition in academic research, as well as in practice, to use size, value, and momentum factors as additional risk factors, and to gauge investment performance relative to these factors. What results is a global five-factor model.

The results show that global real estate has significant exposure to the global size and value factors, but not to the global momentum factor. The exposure to the value factor is particularly strong and well-measured. This confirms, in a global context, earlier findings in the REIT literature that real estate stocks tend to behave like small value stocks. The five-factor model explains a substantially larger fraction of global real estate return variation. The R^2 statistic rises to 71.2%. Thus, the global real estate factor is correlated with the global value factor (42% correlation) and to a lesser extent with the global small stock factor (11% correlation). As a practical matter, real estate strategies look to some extent like value tilts in equity portfolios. Of course this leaves open the question of which of the two factors, the global value factor or the global real estate factor is the more "fundamental" source of

Table 3: **Factor Models for Real Estate**

The dependent variable is the excess return on the Core Real Estate Index (CREI) in columns (1) and (2). The independent variables are a constant, the excess return on the Fama-French global stock market index and the excess return on the Barclays Global Aggregate Bond Index in Column (1). Column (2) adds the global size (smb), value (hml), and momentum (mom) factors as independent variables. The dependent variable is the excess return on the U.S. NAREIT All-Equity REIT index in columns (3) and (4). The independent variables are a constant, the excess return on the Fama-French U.S. stock market index and the excess return on the ten-year U.S. Treasury bond in Column (3). Column (4) adds the U.S. size (smb), value (hml), and momentum (mom) factors as independent variables. The first row reports the intercept α, the other rows report risk factor exposures β. The R^2 and model-implied expected return (without alpha but with risk-free rate) are in the last two rows of each panel. The sample is December 1994–June 2015.

	Global		U.S.	
	(1)	(2)	(3)	(4)
α	0.02	−0.13	0.26	−0.01
t-stat	0.09	−0.63	0.75	−0.03
β^s	0.82	0.87	0.75	0.75
t-stat	11.21	14.73	5.97	8.99
β^b	0.68	0.55	0.27	0.45
t-stat	5.01	4.69	1.23	3.01
β^{smb}	–	0.33	–	0.54
t-stat	–	4.24	–	6.84
β^{hml}	–	0.58	–	0.87
t-stat	–	7.16	–	9.09
β^{mom}	–	−0.07	–	−0.15
t-stat	–	−1.41	–	−2.15
R^2	62.43	71.19	31.85	57.77
Exp. ret.	9.55	11.32	9.66	12.87

risk. Funke, Gebken, Michel, and Johanning (2010) argue that that exposure to real estate risk goes a long way towards accounting for the value premium. More likely in our opinion is that both reflect a deeper, structural source of economic risk. An extensive literature on the value premium in equities has debated the question of what that source of aggregate risk is, arguably without final resolution (see, for example, Koijen, Lustig, and van Nieuwerburgh, 2015, for a recent contribution). We return to this discussion below.

Because value stocks have higher average returns than growth stocks, exposure to this value factor adds to the expected return on real estate. As the last row of the table shows, the expected return increases to 11.3%. This expected return reflects compensation for time value of money (2.57%), stock market risk exposure (5.33%), bond market risk exposure (1.55%), small stock exposure (0.11%), value stock exposure (2.34%), and momentum stock exposure (−0.58%).

Finally, the first row shows that global real estate has underperformed. The five-factor alpha is negative: −0.13% per month or −1.5% per year. However, the estimate is not statistically different from zero. We conclude that the underperformance is neither statistically nor economically significant.

Adding risk factors in U.S.   Column (4) repeats the exercise for the U.S. The results are similar, and if anything more extreme. Adding size, value, and momentum factors adds substantially to the R^2. U.S. REITs load more heavily on all three risk factors than in Column (3). The (negative) momentum exposure becomes statistically significant. After controlling for the additional factors, the outperformance of U.S. REITs disappears entirely. The five-factor alpha is essentially zero. The cost of capital according to the U.S. five-factor model is 12.9% and reflected the added compensation for value risk (2.0%), small stock risk (1.29%), and momentum risk (−0.87%). The conclusions from the five-factor model are consistent for the global and the U.S. analysis.

The appendix extends the U.S. analysis in two directions. First, it discusses a longer sample that starts in 1972. We find results that are broadly consistent with the one presented here. Second, it considers a sixth risk factor, an illiquidity risk factor proposed by (Pástor and Stambaugh, 2003). It is the return on a tradable strategy that goes long illiquid stocks and short liquid stocks. It is only available for the U.S. We find that U.S. REITs have a zero exposure to this factor after including the other five factors. This null result suggests that REITs may not be the best vehicle for tapping into an illiquidity premium for real estate assets and/or that the illiquidity in stocks is different from that in real estate assets.[19]

---

[19]The price of the REIT will deviate from that of the underlying assets until its expected returns correctly reflect the high liquidity of the REIT rather than the low liquidity of the assets it contains (see Cherkes

Implications Global and U.S. listed real estate have neither out- nor underperformed a properly matched funding mix of stocks, bonds, and SMB, HML, and MOM factors. The zero alpha result implies that the optimal portfolio should not overweight real estate *relative to the optimal portfolio of stocks and bonds*. This leaves open the possibility that real estate could help improve the Sharpe ratio of a *sub-optimal* portfolio of stocks and bonds, for example a portfolio with certain position limits on the stock and bond positions. Of course, an unconstrained stock-bond portfolio would also be able to achieve the higher Sharpe ratio.

Time variation in risk exposure As highlighted by the dynamic volatility correlation analysis above (recall Figure 15), risk exposures of real estate to the underlying stock and bond risks are not constant over time. To understand what constitutes fair compensation for bearing real estate risk today, it is important to consider a dynamic factor analysis. We estimate the two- and five-factor global risk models over rolling 60-month windows. The first 60-month window ends in 1999 because the data start in 1994. Figure 16 plots the resulting dynamic risk factor exposures. Several interesting patterns emerge:

1. The abnormal return, α_r, plotted in the first panel fluctuates substantially between $+1\%$ and -1% per month. The graph shows that the two-factor model alphas are usually below the five-factor model alphas. While there is a long stretch of time where the two-factor alpha is positive and economically large, this is much less the case for the five-factor alpha. The last reading (for the five-year period ending in June 2015) on the five-factor alpha is essentially zero. Estimated alphas over horizons of 60 months are to be interpreted with caution.

2. The stock market beta of real estate, plotted in the top right-hand side panel, also fluctuates substantially. The two-factor beta fluctuates between 0.3 and 1.3, while the six-factor beta is more stable and moves between 0.6 and 1.2. There was a strong increase in the equity risk of real estate during the financial crisis. This is a reflection of the rise in correlation between stocks and real estate as well as a rise in the relative volatility of real estate stocks relative to the overall stock market. The stock market beta of real estate has decreased back down by the end of the sample to pre-crisis levels.

3. The bond market beta of real estate fluctuates dramatically between 0 and 1.2, and shows a steady increase over the last five years to reach an all-time high in June 2015. Both 2- and 6-factor models produce similar bond betas. The rising beta reflects the rise in the correlation between real estate and bonds as well as the relative rise in

---

et al., 2009).

volatility of real estate relative to bonds. The high bond beta suggests that real estate is currently particularly vulnerable to an increase in interest rates.

4. The exposure of real estate to the small stock factor fluctuates between -0.1 and 0.5, and has declined as of late. The exposure to value stocks has an even wider range (between -0.1 and 1.1), and has also been on a steady decline over the past five years. The momentum risk of equity REITs is modest and moves in a fairly narrow range around zero. Recently, real estate stocks have gone from being a momentum hedge to being exposed to momentum. Combined, these additional factor exposures have gradually diminished in importance over the past five years insofar that the 2- and 6-factor models explain nearly the same variation in real estate returns at the end of our sample (73.9% versus 71.8%).

The time-varying beta model can be applied to calculate the fair expected return on real estate. Figure 17 plots the risk premium for real estate, i.e., the cost of capital without the risk-free interest rate. The left panel is for the two-factor model, while the right panel is for the five-factor model. To construct the risk premium of real estate, we hold the average risk premium on each factor constant at its full-sample average and only allow the betas to fluctuate over time. We do this given the difficulty in estimating average returns and in order to avoid misinterpreting negative realized returns on factors as periods with low risk premia. The two-factor risk premium goes as low as 3.9% per year and as high as 8.8%. The five-factor risk premium is between 3.3% and 10.5%. On average, the five-factor risk premium is about 1.4% higher, but the difference can be as high as 4.2%. The graph clearly shows the increased importance of bond market risk at the end of the sample, even as the stock market risk and the small value risks have declined. In June 2015, the risk premium on real estate was 7.38% according to the two- and 7.50% according to the 5-factor model. One still needs to add in a risk-free interest rate to turn these numbers into costs of capital.

5.1.6   Inflation and Growth Exposure

The mandate asks for an analysis of whether, and to what extent, real estate provides a hedge against inflation risk as well as how it correlates with economic growth. We perform this analysis in the context of the United States since it is difficult to define what constitutes world inflation.

Inflation hedging   There is a long line of academic research in finance that investigates the extent to which stocks hedge inflation, obtaining mixed results. As we shall see, this is not surprising in light of the large variation over time in the correlation between stock returns and inflation.
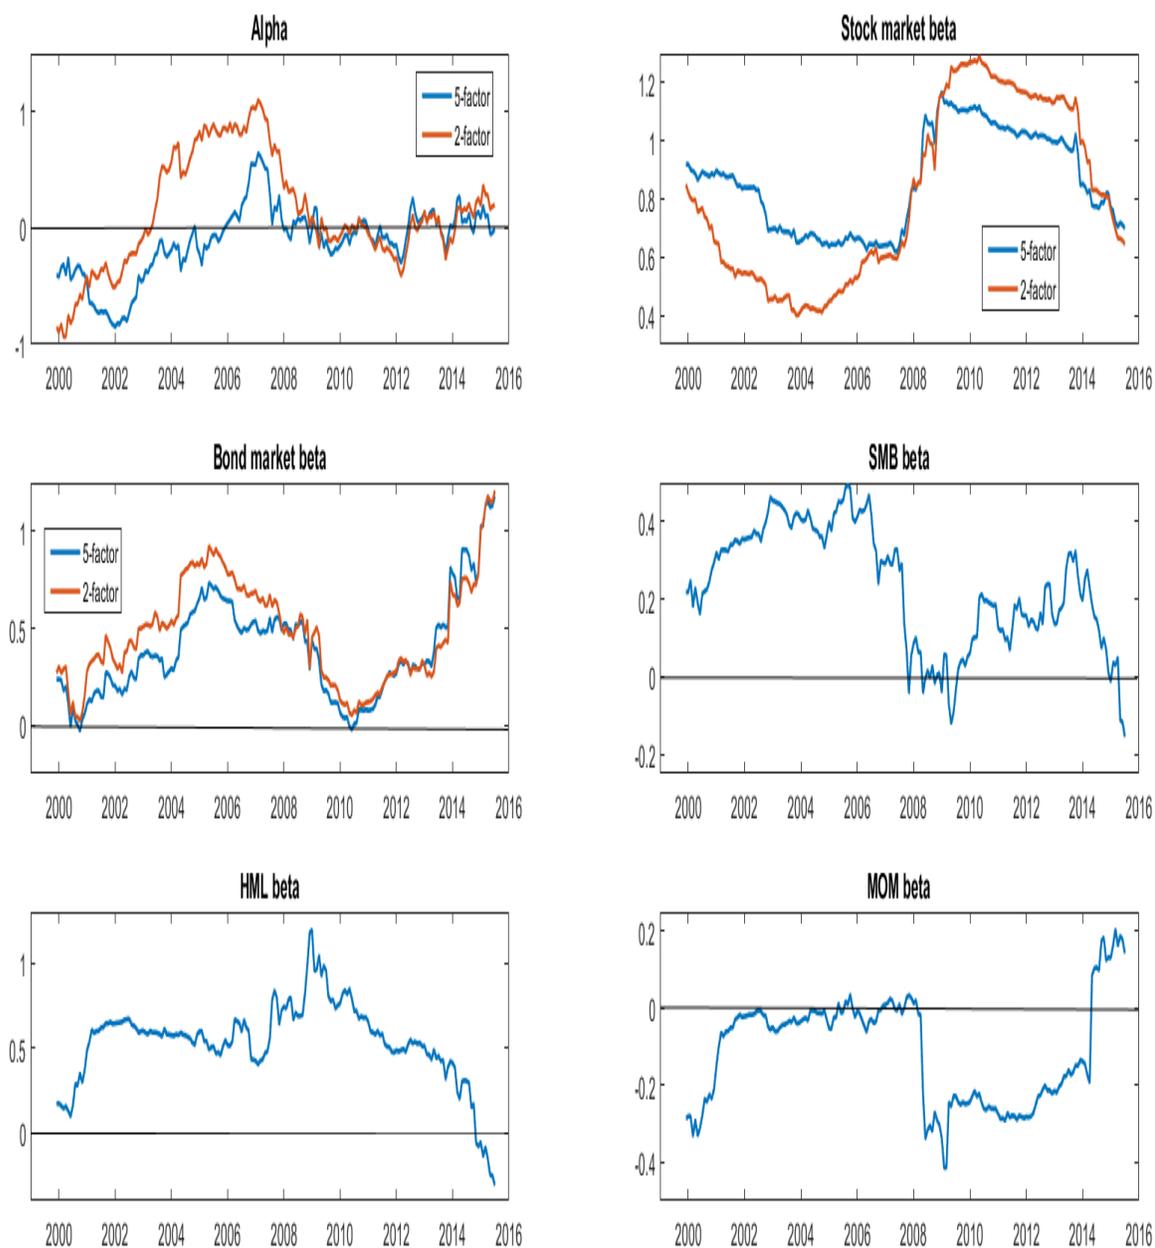
Figure 16: Time-Varying Betas for equity REITs

The figure plots the exposures (betas) of global listed real estate to five global risk factors: the stock market excess return, the bond market excess return return, the size (SMB) factor, the value (HML) factor, and the momentum (MOM) factor. Each risk-factor exposures is estimated via a multivariate regression using the most recent 60 months of data. The sample period is December 1994–June 2015.
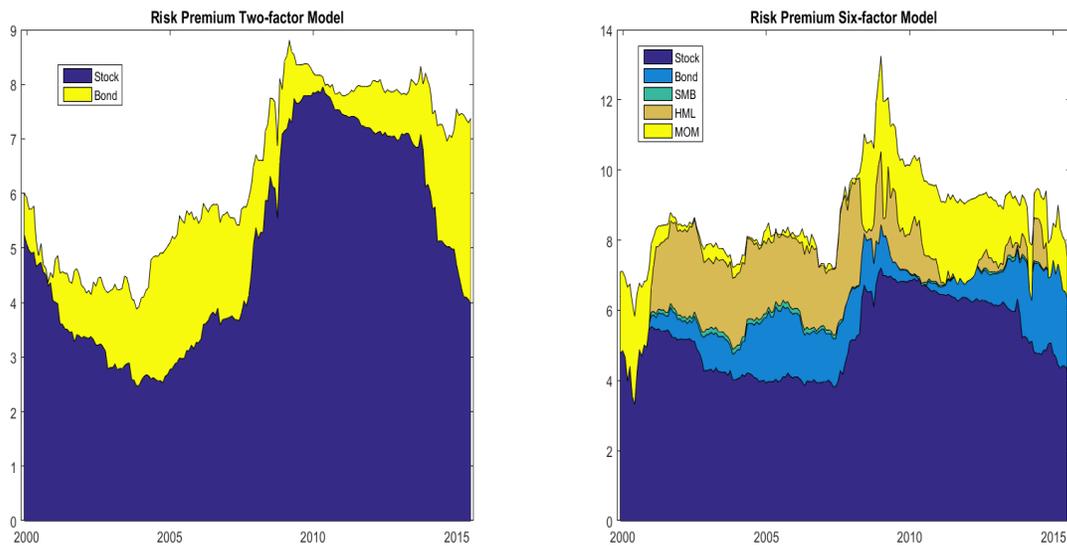
Figure 17: Risk Premium Decomposition for equity REITs

The left panel plots the risk premium, or expected excess return, on global real estate stocks as implied by the two-factor model. The two factor model contains the excess return on the global stock market and the global bond market as its factors. The right panel plots the expected excess return on global real estate as implied by the five-factor model. It adds the global size (SMB), value (HML), and momentum (MOM) factors. All factor exposures (betas) are estimated based on 60-month rolling windows. To calculate the risk premium, we multiply each beta with the average return on each factor, where the averages are computed over the full 1994–2015 sample.

Main results Because inflation has seen dramatic swings over time, it is useful to consider a long sample. Since we focus on the U.S., we are able to go back to 1972. Over the January 1972–June 2015 sample, Consumer Price Inflation (CPI) in the U.S. was 4.06% per year with a standard deviation of 1.33% annualized. We study the correlation between real estate returns and inflation at various horizons to account for the possibility that monthly data may be too noisy to reflect the true underlying correlation of the asset class with inflation. We find that the correlation of real estate returns and inflation is −4.0% at the monthly frequency, −0.2% at quarterly frequency, −5.3% at the annual frequency, 7.0% at the 24-month frequency, and +20.0% at the 36-month frequency. Thus, at short horizons, there is little inflation-hedging ability, but at longer horizons, real estate returns do start to comove positively and more meaningfully with inflation.

For comparison, we also calculate the correlation between stock returns at large an inflation. It is as follows: −8.5% at the monthly frequency, −7.4% at the quarterly frequency, −10.0% at the annual frequency, −3.8% at the 24-month frequency, and +9.5% at the 36-month frequency. Naturally, long-term nominal Treasury bond returns are negatively correlated with inflation: −12.1% at the monthly frequency, −23.4% at the quarterly frequency, −16.9% at the annual frequency, −17.9% at the 24-month frequency, and −23.0% at the 36-month frequency. Real estate has the least negative, or most positive, correlation with inflation across the three asset classes considered. While its correlation with inflation is not that different from that of stocks, real estate does hedge inflation risk a lot better than long-term bonds.

The return on real estate is the sum of a dividend growth component and a component that reflects changes in the price-dividend ratio. As the horizon lengthens, the dividend growth component starts to account for a greater fraction of the return variability. Thus, the inflation hedging properties of real estate returns are increasingly accounted for by the hedging properties of real estate cash flow growth. Since rents on commercial real estate assets are often adjusted to reflect changes in the Consumer Price Index, we expect to see a strong positive correlation between real estate dividend growth and inflation. Indeed, we find a 33.2% correlation at the annual frequency.

Rolling correlations Finally, we investigate how the correlation of inflation with RE-ITs and with stocks has evolved over time, employing 120-month rolling windows. The middle panel of Figure 18 shows that the correlation of REITs with inflation has trended up and is around +10% for the most recent decade (2005–2014). A similar pattern can be observed for stocks as a whole (right panel). The increase in the correlation between stocks and inflation is consistent with the decline in the correlation between stocks and nominal

bond returns observed earlier. Stocks and equity REITs are as good an inflation hedge now as they have ever been. Yet, the correlation with inflation is still not that different from zero. Finally, the left panel shows that inflation volatility has been on the rise recently. The quantity of inflation risk and the need for a hedge have risen.
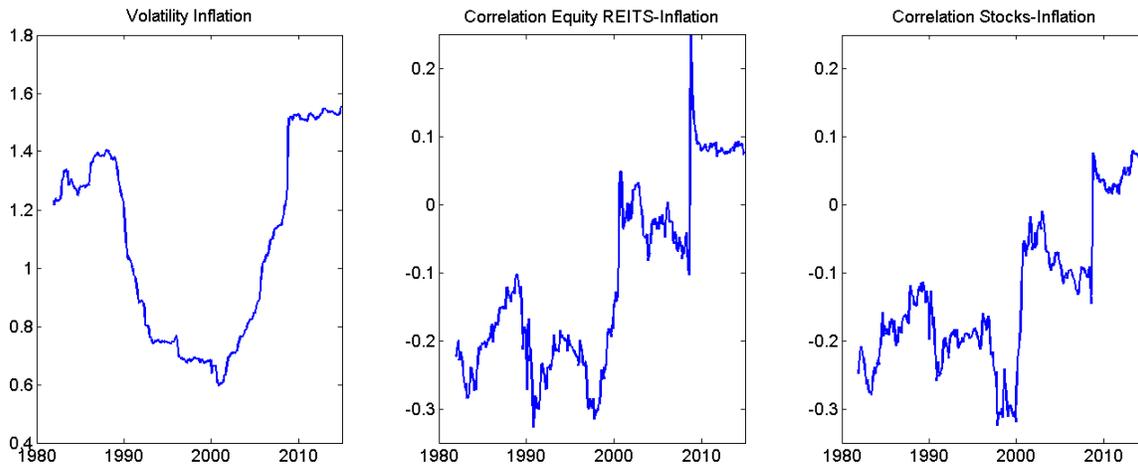
Figure 18: Time-Varying Inflation Volatility and Correlations

The figure plots the volatility of U.S. CPI inflation (left panel) and its correlation with equity REITs (middle panel) and with stocks (right panel). Each point reflects the volatility or correlation measured over the previous 120 months. The sample is January 1972 until December 2014.

Economic Growth Risk   Next, we study the correlation of real estate returns with economic growth. We use the Chicago Fed National Activity Index as our measure of economic activity, *CFNAI* for short. It summarizes the information about the level of economic activity in about 80 time series of macro-economic activity. It has the convenient feature that it is normalized to zero when the economy grows according to trend, it is positive when the economy grows above trend and negative when it grows below trend. The series has standard deviation of 1. The series is available monthly from 1967 onwards. We use the main series, which is the three-month moving average series.

Main results   Over the January 1972–June 2015 sample, the correlation between real estate returns and economic activity is 1.1% at the monthly, 9.3% at the quarterly, 42.4% at the annual, 10.8% at the two-year, and 77.4% at the three-year frequency. The same correlations between stocks at large and economic activity are -2.1% at the monthly, 1.1% at the quarterly, 50.2% at the annual, 21.0% at the two-year, and 71.0% at the three-year

frequency. Bonds are negatively correlated with economic activity for frequencies up to one year. Treasury bonds do well in recessions, which is a well-known flight-to-safety effect. At lower frequencies, their returns start to comove positively with economic activity. In these calculations, the level of economic activity is always calculated at the end of the period. In sum, the evidence for positive correlation between economic activity and real estate returns is weak at high frequencies, but becomes much stronger over the types of frequencies typically associated with business cycles (one- through three-year fluctuations).

One reason for this stronger comovement is that returns over longer horizons are starting to get more heavily influenced by their dividend growth component and less by their discount rate component. Real estate cash flows (e.g., rents) are cyclical. The right panel of Figure 19 shows the dividend growth rate on U.S. REITs. The graph shows that dividend growth falls substantially during bad macro-economic times, such as the Great Recession of 2008-09 but also the recessions of 1974, 1982, 1991, and 2001. The large drop in cash flows, especially during prolonged recessions, is is a trait real estate stocks have in common with value stocks (see Koijen et al., 2015). Looking at dividend growth data, we indeed find that dividend growth on U.S. REITs has a positive correlation of 21.4% with economic activity (CFNAI) at the annual level.

Because REITs must pay out such a large fraction of their net operating income for tax reasons, and because NOI is highly cyclical, this correlation of real estate cash flows with macroeconomic activity makes sense. It implies that a fundamental source of risk for real estate is exposure to prolonged macro-economic declines. Value stocks are exposed to this same risk factor. Because deep recessions are bad states of the world for the average investor, it is natural that investors would demand a risk premium for this exposure.[20] In the U.S., real estate investors have earned 3.6% per year over and above what they should have earned to reflect the regular stock and bond risk (recall column (3) of Table 3). One interpretation of this 3.6% is as a premium for macro-economic recession risk. Value stocks also earn this premium. Consistent with this interpretation, once a value factor is added to the stock and bond factors, the outperformance disappears. The five-factor alpha in the U.S. is 0.17% per year and not different from zero (Column (4) of Table 3). In the global data, we also find essentially a zero five-factor alpha for real estate stocks, and we find a strong value-factor exposure for real estate stocks (Column (2)). One difference with the U.S. analysis is that the two-factor model generates no outperformance either (Column (1)). One possible reconciliation is that the global stock and bond market factors better reflect

---

[20]Note that while the stock market as a whole also suffers in bad macro-economic times, a regular stock market factor may not (fully) pick up this exposure. Koijen et al. (2015) provide evidence along these lines for value stocks.

the exposure to macro-economic risk than the U.S. stock and bond market factors.
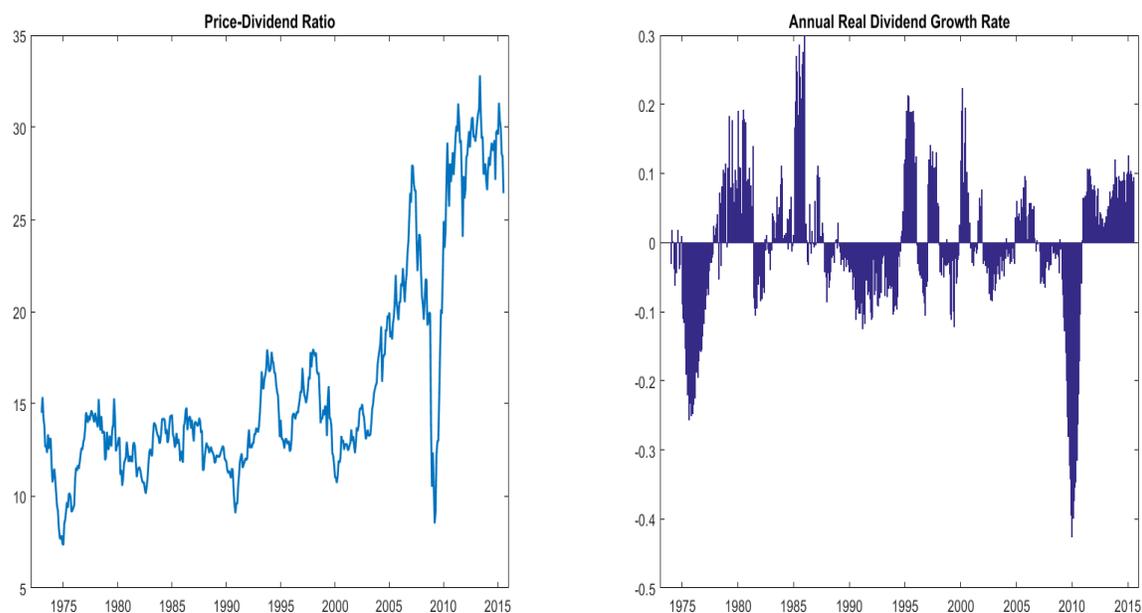
Figure 19: Price-Dividend ratio and Dividend Growth on U.S. equity REITs
The left panel plots the price-dividend ratio on the U.S. All Equity REIT index from NAREIT. Data are monthly from December 1972–June 2015. The right panel plots dividend growth on the All equity NAREIT index. Dividends are invested within the year at the T-bill rate.

5.1.7 Portfolio Analysis

What is the role of real estate in the global optimal portfolio? Because of its simplicity, we limit our analysis to a simple mean-variance optimization approach.[21] It is useful to recall that if we started from the mean-variance efficient portfolio of stocks and bonds, and the alpha of real estate with respect to a stock-and-bond factor model were zero, then the optimal portfolio would have exactly zero weight on real estate. Because the global stock market portfolio already includes real estate stocks, the allocation to real estate is to be interpreted as an over-allocation.[22] A non-zero (over-)allocation to real estate thus emerges only if either the alpha of real estate is positive or if the current allocation to stocks and bonds is not mean-variance efficient.

---

[21]We recognize the limitations of this type of analysis and the need to complement it with other approaches. Below we consider imposing various portfolio constraints.

[22]Norges Bank's real estate portfolio consists of the direct real estate investments and the over-allocation to real estate stocks.

We use the returns on global real estate, global stocks, and global bonds over the period 1994–2015 to form both expected returns and the return covariance matrix, which form the inputs in this analysis. We find that the unconstrained mean-variance efficient portfolio of the three risky assets consists of 2.5% real estate, 18.0% stocks, and 79.5% bonds. It has an annualized average return of 6.1% with volatility of 6.0% and a Sharpe ratio of 0.586. The small allocation to real estate assets in the tangency portfolio reflects global real estate's small abnormal return relative to the appropriate combination of stocks and bonds (recall Column (1) of Table 3); the average return of global real estate is 9.79% per year, which is only 0.25% above the expected return implied by the stock and bond combination of 9.55%. Indeed, if the expected return on real estate were set exactly equal to 9.55% (implying a zero two-factor alpha), we would find a zero (over-) allocation to real estate.

Repeating this analysis on U.S. data, we find a tangency portfolio that allocates 24.8% to stocks, 65.3% to bonds, and 9.9% to real estate. This tangency portfolio has a mean return of 7.5% with a volatility of 6.4%, and a Sharpe ratio of 0.773. The much larger two-factor alphas of U.S. real estate relative to U.S. stocks and bonds implies a substantially larger (over-)allocation to real estate than in the global sample.

One potentially undesirable feature of the portfolios reported above (in light of potential increases in interest rates) is that they feature a very large allocation to bonds. This may reflect a historically "lucky" sample where interest rates trended down and average realized bond returns were high. Such a scenario may be unlikely to repeat itself given the currently historically low bond yields. Therefore, we also ask what the maximum Sharpe ratio (tangency) portfolio is imposing that the weight on bonds not exceed 35%, which is the current target weight in bonds of the GPFG. With global data, the tangency portfolio invests 27.9% in real estate, 37.1% in global stocks, and 35% in bonds. Because of the constraint that is imposed on the bond position, the Sharpe ratio must be lower than that of the unconstrained tangency portfolio. The loss in Sharpe ratio in our sample of global data is 0.092 (0.585-0.494). However, this portfolio has a higher mean return than the (65,35,0) portfolio without real estate (7.86% versus 7.56%) and a substantially higher Sharpe ratio (0.494 versus 0.474); it adds 20 basis points to volatility (10.71% versus 10.53%). In sum, once the maximum constraint on the bond position is imposed, real estate takes a prominent place in the portfolio.

Can one add real estate without increasing portfolio volatility and without sacrificing return? In the same vein, we can ask whether moving the portfolio away from 65% global stocks and 35% global bonds by adding real estate can help improve the mean return of a global portfolio without adding to portfolio volatility. The answer is yes. We find

that, for any real estate portfolio weight between 0 and 52%, there exist portfolio weights in stocks and bonds such that the new portfolio has both a higher mean return and a lower portfolio volatility than the 65-35% stock-bond portfolio. Figure 20 plots all portfolios that satisfy this criterion (limiting ourselves to portfolio weights that have 1% increments). At the far left side of the graph, we find the portfolio the (65,35,0) portfolio that invests nothing in real estate. This portfolio has a Sharpe ratio of 0.474, indicated by the thin solid line. Moving to the right, we increase the allocation to real estate. For example, with a 5% real estate allocation, we find a (60,35,5) and a (59,36,5) portfolio which both have higher mean return and lower volatility than the (65,35,0) portfolio. The Sharpe ratios on these portfolios are 0.481 and 0.484, respectively. With a 10% real estate allocation, portfolios ranging from (52,38,10) to (55,35,10) have higher mean and lower volatility, with Sharpe ratios ranging from 0.494 to 0.487. As a final example, with a 28% share in real estate, as we found in the constrained optimization, portfolios ranging from (29,43,28) to (35,37,28) increase return and lower volatility. Sharpe ratios range from 0.510 to 0.497. Note that these portfolios all violate the 35% maximum in bonds we imposed in the constrained optimization. Not surprisingly, they end up with higher Sharpe ratios than the 0.494 we found for the (37,35,28) portfolio. The Sharpe ratio of 0.51 is still substantially below the maximum unconstrained Sharpe ratio of 0.586 that we found in the unconstrained portfolio analysis. The unconstrained tangency portfolio does not satisfy the criterion that the mean return be at least as high as the 10.53% on the (65,35,0) portfolio. Indeed, because the unconstrained tangency portfolio invested so much in bonds, its mean return was only 6.1%. Returning to the graph, as the portfolio weight of real estate increases, the weight on stocks is reduced and the portfolio weight on bonds is increased. Thus, to preserve return and not increase risk, a large real estate allocation must be funded out of stocks more than one-for-one. None of the portfolios in the graph invest less than 35% in bonds.

These results convincingly show that adding real estate to a portfolio of stocks and bonds need not increase total portfolio volatility, while it can increase mean returns and Sharpe ratios. This ultimately reflects the diversification benefits of real estate, a benefit that is present even though its two-factor alpha is zero. Needless to say, these results are based on historical data and may reflect the favorable environment bonds have enjoyed over the past several decades.

5.1.8 Valuation Ratios

The GPFG currently has a fixed portfolio target of 5% in real estate and is in the middle of acquiring assets to achieve that target. With roughly 2.7% of its current portfolio allocated to real estate, it is aiming to add about 1% per year over the mid-2015 to end-2017 period. An
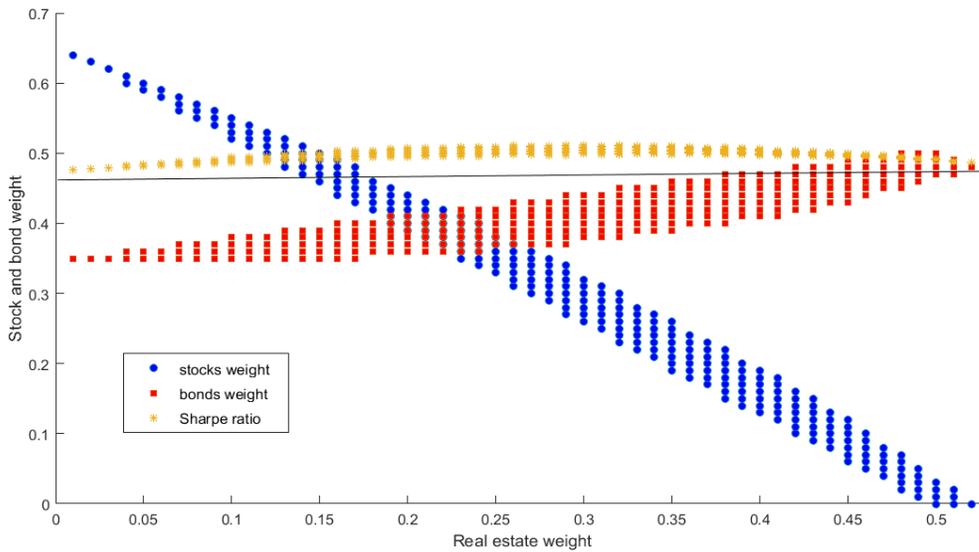
Figure 20: Portfolios with Real Estate

The horizontal axis indicates the weight of real estate in the portfolio, in one percent increments. The vertical axis plots the portfolio share in stocks (blue circles) and in bonds (red squares) of those stock-bond-real estate combinations that have both a higher mean return and a lower portfolio volatility than the 65-35% stock-bond portfolio. The graph also plots the Sharpe ratios (yellow stars) of the corresponding portfolios.

important question arises as to the wisdom of a fixed and relatively rapid glide path towards that target allocation. The analysis below suggests that real estate assets may currently be expensive by historical standards. Recent market commentators share this concern.[23] The high valuations may be especially pronounced for the "trophy assets" the GPFG is interested in.

As they have for hundreds of years, real estate markets go through cycles, and the timing of investments is crucial for return performance over medium-term horizons. As a case in point, private equity funds invested aggressively in real estate in 2006–08. These vintages performed very poorly and dragged down the overall return performance of real estate private equity over the entire history of that asset class, in large part because so many more dollars were invested in these vintages compared to earlier and later ones (Fisher and Hartzell, 2013). While the academic literature on market timing suggests that it is difficult to time markets (including the real estate market), there are times where the opportunity cost of waiting for valuation levels to normalize seems particularly low. The analysis below suggests that now may be such a time. One of the structural advantages of the GPFG is precisely that it does *not* have to adhere to a strict time line of filling asset-allocation buckets.

We start with a discussion of various common valuation metrics applied to U.S. commercial real estate. The high quality of data make this a good place to start. We then proceed to a more sophisticated present-value model analysis. Finally, we expand this present-value analysis to the global data.

Repeat-sales price indices First, we look at price levels for the universe (listed and unlisted) of all U.S. commercial real estate buildings with transaction prices above $2 million. We use the Real Capital Analytics & Moody's repeat-sales commercial property price index, which only considers actual transactions of buildings that sell at least twice in the sample. The data go back to the end of 2000.[24] The left panel of Figure 21 shows that U.S. commercial property prices saw a large boom of 80% between 2000 and 2007 (solid line). Nationwide, prices increased by 80%. They fell from a level of 180 in November 2007 to 107 in January 2010, a decline of 40%. Since then, they have risen back to 200.5, an increase of 87%, with price levels in May 2015 surpassing the 2007 peak in the Fall of 2014.

---

[23]The Federal Reserve Bank recently investigated the risk of a bubble in commercial real estate markets, while the CEO of Vornado noted that this was a time to harvest rather than to sow (see "US commercial property hotspot draws Fed attention," Financial Times, July 16, 2015, and "Norway Funds Bulks Up on Real Estate," Wall Street Journal, July 21, 2015). 63% of surveyed investors in private equity real estate funds believe rising valuations is the key issue for the private real estate market in 2015 (see Preqin, 2015a). The authors have had recent private conversations with several real estate market participants who are of the same opinion.

[24]There are drawbacks to repeat-sales indices, particularly the small samples sizes they often imply. For an alternative in the residential area, see Stanton and Wallace (2015).

The dashed line in the figure plots the price for office buildings in Central Business Districts. Since the Fund invests a substantial fraction of its holdings in office buildings in major markets, this is a relevant series. U.S. CBD Office prices rose 99% between 2000 and February 2008, or almost 20 percentage points more than the nationwide all-property type index. CBD Office fell all the way back to 100.9 by September 2009, a 49% drop over 20 months. It saw a very strong recovery to 270 by May 2015, a 167% increase, exactly twice the overall commercial property price increase. A further breakdown (not shown in the graph) reveals that CBD office *in major markets* rose 160% between September 2009 and March 2015 to reach 278.5. In contrast, non-major market CBD office increased by substantially less over the same period, 112%, to reach 172. Both series were 100 in December 2000.

In the year to May 2015, CBD office prices rose 23.6% (compared to 16.4% overall). This followed an increase of 20.9% in the year to May 2014 (17.3% overall). Combining the recent price appreciation with the record transaction volumes, there is no doubt that the U.S. commercial property market is on fire. There is also no doubt that the increasing interest from large (foreign) institutional investors in real estate has fueled price increases over the past several years, especially for class A properties in prime locations of the major cities of the world. Indeed, property prices in the major U.S. markets have increased a lot more than nationwide prices, as the dash-dotted line in the left panel of Figure 21 indicates.

We observe that CBD Office has displayed larger average price appreciation (8.5% per year compared to 5.6% overall) but also much higher price volatility. The standard deviation of (overlapping) 12-month price growth rates was 17.2% over this period compared to 12.5% for the nationwide series. While cash flows on major CBD office buildings may be resilient due to the robust demand for quality office product in such locations, prices are certainly *not* less volatile.

As an aside, a strikingly similar picture arises in the U.K. Nationwide commercial property prices, measured using the same repeat-sales methodology and by the same data provider, increased by 62.9% between 2002.Q4 and 2007.Q2 (the comparable increase in the U.S. was 64.0%). Central London office prices rose by a similar 62.4%. During the crisis, nationwide commercial property prices fell by 41.1%, below 2002.Q4 levels. Central London office prices fell by somewhat less during the crisis, −34.7%. Since 2009.Q2, nationwide prices rose by 75.2% while Central London office prices rose by 98.6%. In the year to 2015.Q2, Central London office prices rose by 23.2%, very similar to CBD Office in the U.S.

We conclude that prices have risen to high levels in both the U.S. and the U.K., and more so in the top-tier locations and core properties such as Class-A CBD Office buildings.
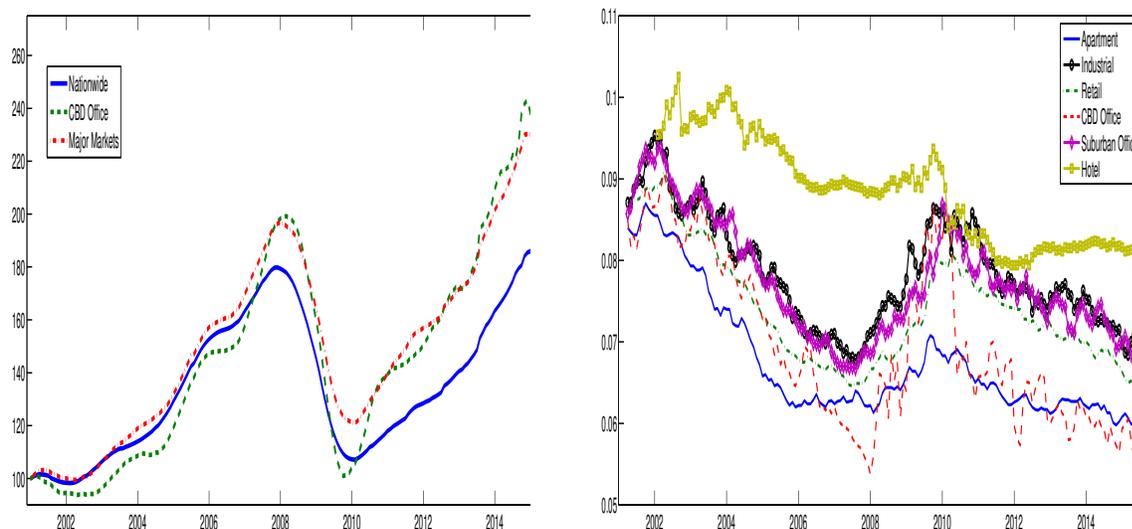
Figure 21: Commercial property prices and Cap Rates U.S.

The left panel plots the RCA/Moody's Repeat-Sales Commercial property price Index. Data are from Real Capital Analytics. The CPPI is available monthly from December 2000 until May 2015. Cap rates are available monthly from March 2001 until May 2015.

Capitalization rates While the repeat-sales price index uses sales of the same properties, and controls (at least to some extent) for the quality of the building over time, it does not keep fundamentals constant. Increasing prices may thus reflect improving fundamentals such as net operating income (rental income minus operating expenditures). One valuation metric the industry often turns to is the capitalization or "cap" rate, defined as current net operating income divided by current price. As such, it is like a dividend yield on a stock or a yield on a bond. Low cap rates indicate low yields or high prices per dollar of net income. The right panel of Figure 21 plots cap rate data from Real Capital Analytics for several U.S. commercial property types. The office sector is broken down into CBD and suburban office.

The pattern is similar across all property types: cap rates fell from a series high at the start of the graph in 2001 until late 2007. Cap rates then increased as property prices fell (by more than NOI). Since late 2009, cap rates have been trending back down. Cap rates on U.S. apartment buildings stood at 5.9%, industrial property at 6.9%, retail at 6.5%, CBD office at 5.1%, suburban office at 6.8%, and hotels at 8.2% as of May 2015. For CBD office and apartments, the cap rate series reach an all-time low, implying prices at an all-time high relative to current fundamentals. Industrial, Retail, and Suburban Office are virtually back to their valuation peaks of July 2007. Only Hotels is bucking the trend with nearly constant cap rates of 8% over the past 4 years. In sum, prices have been rising not only in absolute

62

terms but also relative to fundamentals.

Naive risk premium   One can decompose the cap rate into a long-term safe bond yield and a spread. The spread between the cap rate and the government bond yield gives an investor a metric of the extra current income yield over and above that on similar-maturity Treasury investments. The spread is often referred to as "the risk premium" on real estate. This terminology is misleading since it ignores the effect of future price changes, which in part reflect future income changes. Over the same March 2001–May 2015 period, the 10-year U.S. Treasury yield fell from about 5% to about 2%. This decline largely accounts for the decline in the cap rates. As of May 2015, investors earned an extra 3.4% yield over the Treasury rate to compensate for the risk in CBD Office. For apartments, the spread was about 4%, for retail properties 4.5%, suburban office 4.7%, and industrial 4.9%. Hotel is the outlier with a yield spread of 6.1%. These yield spreads are 0.7% above their May 2001–May 2015 averages for all property types except CBD office, which currently pays its historical average spread. The spread between cap rates and bond yields is substantially above its 2007 value. While real estate prices are back to peak or have surpassed their peak, real estate valuations are not exceptionally expensive *relative* to long-term bonds. Of course, long-term real bonds might be exceptionally expensive right now. Nevertheless, given the exceptionally low bond yields, investors "searching for yield" in riskier asset classes receive risk compensation for real estate exposure above the historical average.[25]

Present-value model   Since REITs pay out essentially all of their net income, and to the extent that they hold a representative portfolio of institutional-grade buildings, the dividend yield on REITs should be a good measure of the underlying cap rate on commercial real estate.[26] Studying REITs has the advantage of a long sample that starts in 1972.

The left panel of Figure 19, introduced above, plots the price-dividend ratio on U.S. equity REITs. Each data point refers to the dividends paid out over the course of the year divided by the price at the end of the year.[27,28] The graph shows that over the past 5 years

---

[25]An analysis that uses the dividend yield on equity REITs as the measure of the cap rate confirms that the difference between that cap rate and the Treasury yield is currently 1.2% points above its 1972–2015 historical average.

[26]More precisely, since REITs are levered (about 40% for U.S. REITs), the dividend yield reflects the cap rates of the underlying real estate portfolio minus the yield on the corporate debt minus management expenses. Strictly speaking then, this analysis investigates the variation in dividend yields on REITs rather than commercial real estate cap rates. This seems like a trade-off worth making given that REIT prices are higher-quality and timelier measures of real estate prices than the appraisal- and transaction-based price series typically used when valuing commercial real estate buildings.

[27]We assume that all dividends are reinvested at the 1-month T-bill rate. Koijen and van Nieuwerburgh (2011) argue why this is the most appropriate assumption to make.

[28]We believe there to be three data errors in the NAREIT ex-dividend equity REIT return series. In May

real estate stocks have reached valuations not seen since at least 1972. Valuation ratios have roughly doubled since the end of 2000. The average price-dividend ratio was 13.1 from December 1972–December 2000. It was 28.8 from June 2010–June 2015. As of June 2015, the price-dividend ratio is 26.4, down from a peak of 31.3 in January 2015. Further corrections took place over the summer of 2015.

What can we learn from studying dividend-price ratios (or equivalently cap rates)? A low cap rate or high price-to-NOI ratio must reflect the market's expectation of (i) lower future returns on real estate (i.e., future price declines), (ii) higher future net income (NOI) growth, or (iii) a combination of the two. Using the present-value relationship pioneered by Campbell and Shiller (1988), the cap rate dp is the difference between the expected discounted sum of future returns r and the expected discounted sum of future dividend (NOI) growth Δd:

$$dp_t = \overline{dp} + E_t \left[\sum_{j=1}^{+\infty} \rho^{j-1} \left(r_{t+j} - \bar{r} \right) \right] - E_t \left[\sum_{j=1}^{+\infty} \rho^{j-1} \left(\Delta d_{t+j} - \bar{d} \right) \right], \qquad (1)$$

where the dividend yield (cap rate), the return, and the dividend growth rate are all measured in logs. The first term is the long-term mean dividend yield; $\bar{r}$ and $\bar{d}$ are the long-term mean return and dividend growth rates, respectively.[29] Using this framework, we can ask how much of the observed variation in cap rates reflects fluctuations in discount rates (the first term) versus fluctuations in future cash flow growth rates (the second term). A large finance literature has found that stock price movements largely reflect movements in future prices (first term), rather than in future cash flows (second term). When the dp ratio is low, stock prices tend to fall to restore the dp ratio back to its long-run mean. Put differently, a low dp ratio predicts price declines rather than high dividend growth rates going forward (see Koijen and van Nieuwerburgh, 2011). We now revisit this evidence for U.S. equity REITs.

We find that 25% the overall variance in dp of U.S. REITs is accounted for by the variance of the discount rates (the first term in (1)), 25% by the variance in growth rates (second term), and 50% by their covariance. Alternatively, we can decompose the variance of dp into the covariance with future returns minus the covariance with future growth rates.

1980, the income return is 12.26%, in September 1984 it is 3.42%, and in June 1990 it is 5.45%. These data points are extreme outliers since the average monthly income return is 0.61% and the 99th percentile 1.53%. Further, using these data results in a dividend yield series that does not correspond to the one provided by NAREIT. We assume that these three observations are typos and should be an order of magnitude smaller. We adjust the ex-dividend returns accordingly, while leaving the cum-dividend returns as is. The resulting dividend yield series has a correlation with that reported by NAREIT of 97% instead of 76% pre-correction.

[29]This follows from log-linearizing the definition of a return $R_{t+1} = \frac{P_{t+1}+D_{t+1}}{P_t}$, to obtain $r_{t+1} = k + \Delta d_{t+1} + \rho \, pd_{t+1} - pd_t$, where $dp_t = d_t - p_t = -pd_t$ and all lowercase letters denote natural logarithms. The constants k and ρ are related to the long-term average log dividend-price ratio: $\rho = (1 + \exp(\overline{dp}))^{-1}$. By iterating forward on the return equation, adding an expectation operator on each side, and imposing a transversality condition (i.e., ruling out rational bubbles), we obtain the equation in the main text.

We find that each term accounts for 50% of the variance. This is an interesting finding because it is different from the consensus view for stocks. There, researchers have found that about 90% of the variance of returns comes from variation in expected returns and only 10% from variation in expected growth rates. The visual evidence on dividend growth and the regression evidence on dividend growth predictability by the dp ratio are both consistent with the result that there is a large predictable component in REIT dividend growth.[30]

What do we conclude from this exercise? Above-average cash-flow growth and below-average returns will pull the dp ratio back towards its long-run average from its current low value. The present value model suggests that lower future returns and higher future dividend growth rates will each account for half of the adjustment. In other words, as long as there are no structural changes in the economy that would lead to a permanently lower mean cap rate, we should expect cap rates to move back up.[31] The good news (at least relative to the way stock valuations adjust) is that a substantial portfolio of the adjustment is likely to come in the form of stronger cash flow growth on commercial real estate properties.

Interpreting recent real estate valuation ratios The five-factor model for real estate returns, discussed above, implied an average expected return of 11.5% per year (in logs). At the end of the sample, in June 2015, the expected return was 10.3% as a reflection of the higher risk premium but also the very low risk-free interest rate at that point. Despite the high expected return, we observed a high price-dividend ratio of 26.4 on U.S. real estate in June 2015. We ask what assumptions on future dividend growth rates are necessary to reconcile the observed price-dividend ratio with the five-factor estimate of the expected return. In our calculation, we assume that the expected return gradually reverts back to the full sample mean, over a 10-year period. We also assume that dividend growth is constant for the first ten years and returns to its long-term mean of 3.18% per year (in logs) after 10 years. We compute the annual dividend growth rate (over the first 10 years) that investors must be expecting to justify the current valuation ratio. Figure 22 reports this number for each month starting in December 1976, always using the corresponding price-dividend ratio and expected return for each month (thick line). The thin straight line shows the observed average annual dividend growth rate, as a point of comparison. When the thick line is above

---

[30]In fact, the evidence we find for dividend growth predictability for REITs is consistent with the pre-WW II evidence for stocks as a whole (Chen, 2009). One argument for the disappearance of dividend growth predictability post-WW II among stocks is earnings smoothing. Since REITs cannot engage in dividend smoothing because of the REIT rules, the result of substantial dividend growth predictability makes considerable sense.

[31]One potential argument for permanently lower cap rates is permanently lower growth rates in the economy — a Great Stagnation. See Eggertsson and Mehrotra (2014) and Kozlowski, Veldkamp, and Venkateswaran (2015) for a discussion. See Lettau and Van Nieuwerburgh (2008) to deal with return predictability when the mean of dp exhibits regime changes.

the thin line, the market expects dividend growth rates to be above average. To the extent that positive deviations appear excessive, one could interpret these findings as support for a "bubble."
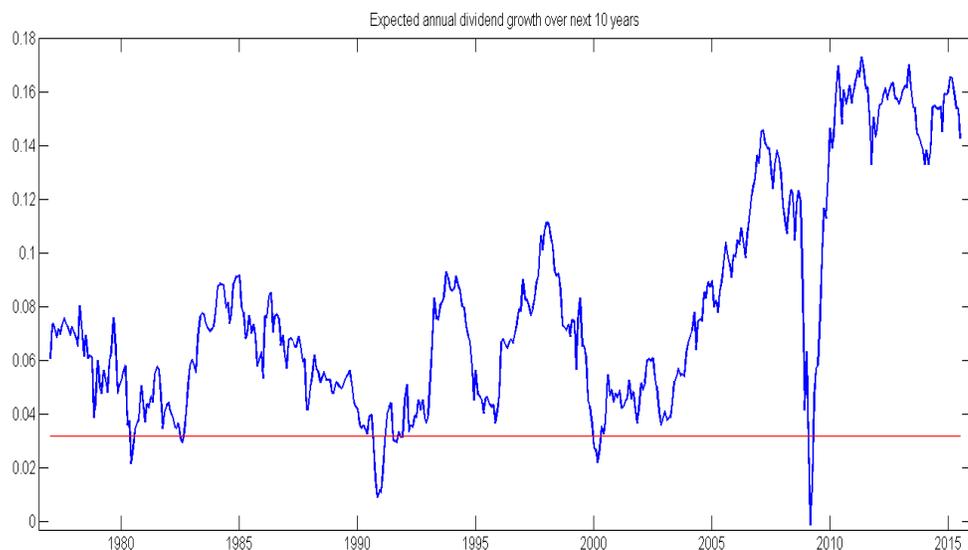
Figure 22: Expected Dividend Growth Implied in U.S. REIT Valuation
In each month, the graph plots the expected dividend growth over the next 10 years, expressed as a per year quantity, that is implied by the price-dividend ratio in that month and the expected return on U.S. REITs in that same month according to the 5-factor model (estimated over a 60-month window ending that month).

Looking at the last observation, for June 2015, we see that investors expect real estate cash flows to grow at 14.3% per year for the next 10 years. Put differently, investors must believe that (discounted) NOI growth will be 78% over the next ten years rather than the 24% implied by the historical average growth rate. Despite the strong cash flow growth observed in 2014 and 2015, we do not deem these implied growth rates particularly plausible. If dividend growth turns out to be only average over the next 10-years, then the price-dividend ratio ought to be 12.1 as opposed to the observed 26.4. That is, REITs would be overvalued by a factor of 2.19 (119%). The graph also shows that real estate was more expensive in January 2015 than it has ever been (16.5% implied growth), according to this metric. For comparison, the implied expected dividend growth number on the eve of the financial crisis in February 2007 was 14.5%. In February 2009, the measure bottomed out at a slightly negative value, suggesting that real estate was cheap.

66

International real estate valuation ratios   We repeat the exercise for the global Core Real Estate index.[32]   Figure 23 shows that the CREI price-dividend ratio is remarkably stable, especially compared to that on the MSCI ACWI stock index.
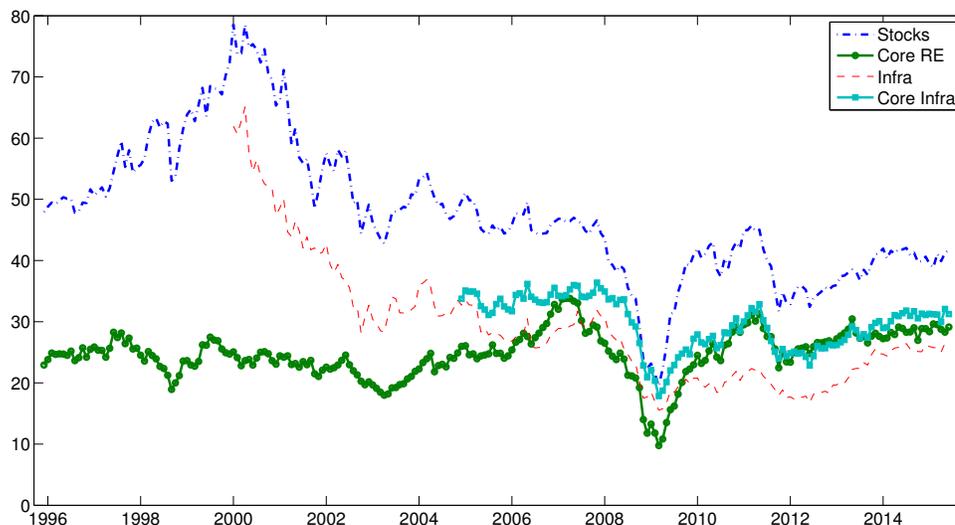
Figure 23: Price-Dividend Ratios Global Real Estate, Infrastructure, and Stocks
The figure plots the price-dividend ratio on the MSCI World Core Real Estate Index (circles), MSCI World Core Infrastructure Index (squares), the MSCI World Infrastructure index (dashed line), and the MSCI ACWI stock index (dash-dotted line). Data are monthly from November 1995–May 2015.

The top panels of Figure 24 plot the *pd* ratio on the CREI and the corresponding dividend growth rate. There is a remarkable alignment suggesting that a lot of the movements in the pd ratios reflect movements in dividend growth rates. The average nominal dividend growth rate on the CREI was 3.5% per year for the 1995–2015 sample. The CREI reaches a price-dividend ratio of 29.1 in May of 2015. This valuation ratio is very similar to that of U.S. REITs (28.4) in the same month. Global real estate valuations are still below their early-2007 peak of 33.8 (CREI), but approaching the previous peak rapidly.

How can current valuation ratios be justified? The CREI pd ratio was 28.3 in June 2015. As we did for U.S. REITs, we ask what cash flow growth rate beliefs justify these prices. Again we use the expected return given by the 5-factor model. Given an expected return of 7.5% in June 2015, investors must believe that dividend growth will be 12.4% each year

---

[32]We use the "price" and "gross" series of the MSCI Core Real Estate indices to construct ex- and cum-dividend returns. From those, we build dividend-price ratios, and the rest of the analysis proceeds as above. We sum dividends within the year by reinvesting them at the U.S. T-bill rate. The resulting Core Real Estate price-dividend ratio is available monthly from November 1995 until May 2015 (235 months).
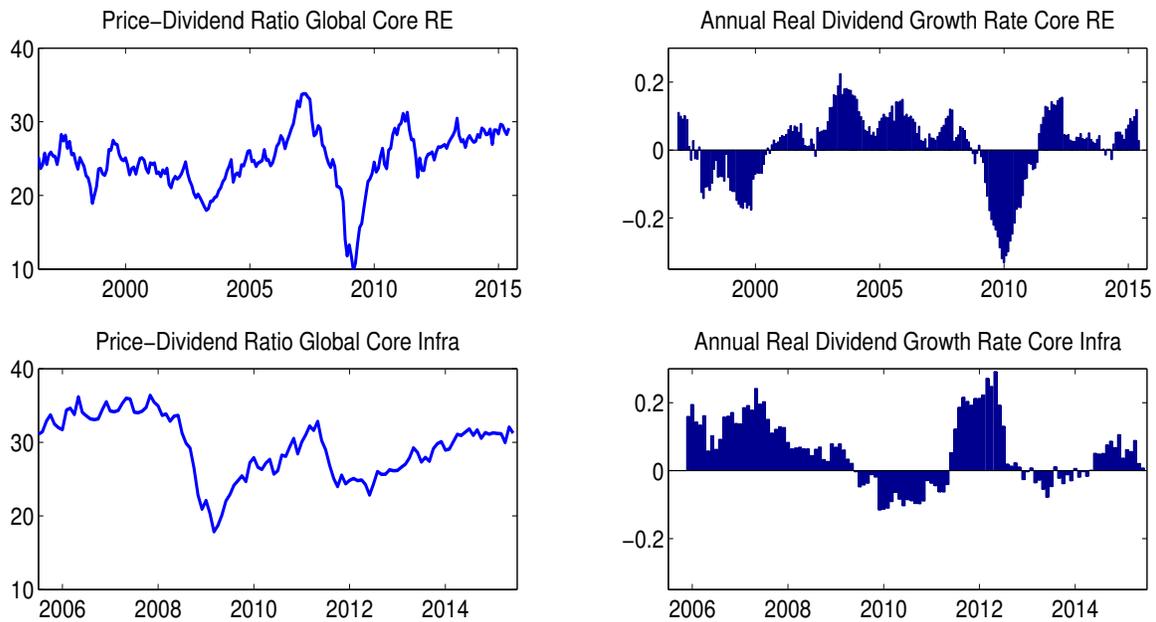
Figure 24: Price-Dividend Ratios and Dividend Growth Rates on Core Real Estate and Core Infrastructure Indices

The left panels plot the price-dividend ratio on the MSCI World Core Real Estate Index (top) and MSCI World Core Infrastructure Index (bottom). The right panels plot the corresponding annual dividend growth rates. Monthly dividends are reinvested within the year and 12-month growth rates are computed from these annualized dividends. Data are monthly from November 1995–June 2015 in the top panels and November 2004–June 2015 in the bottom panels.

for the next ten years before reverting back to the long-term average of 3.0%. Relative to the U.S. numbers, this implied expected growth rate is more than 2% points lower. Also, it is lower than the pre-crisis peak of 16% in February 2007 and lower than the local peak of 14.5% in April 2011; see Figure 25. Nevertheless, it remains implausible that real estate cash flows would grow at 12.4% per year for the next 10 years. If dividend growth were at its historical average instead, real estate would be overvalued by a factor of 1.97 (97%).
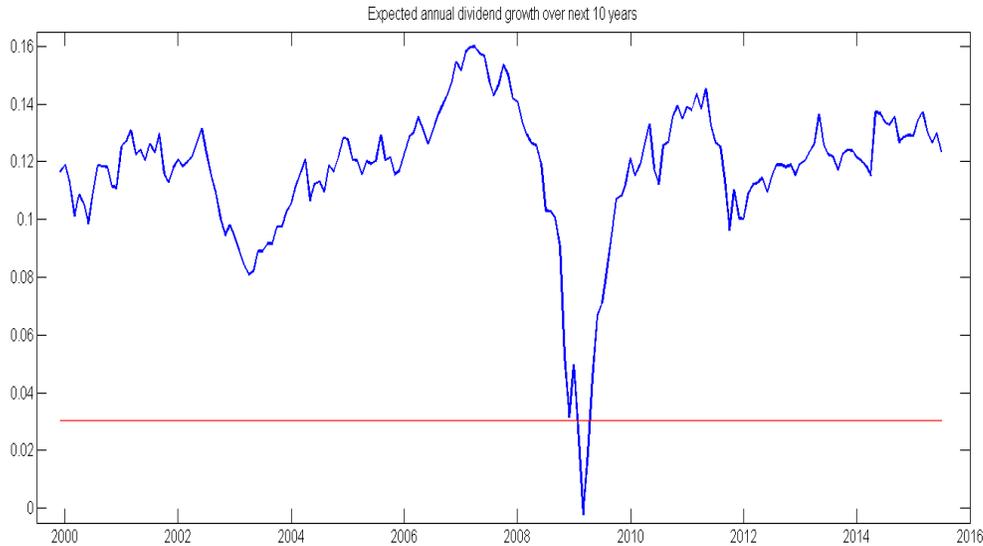
Figure 25: Expected Dividend Growth Implied in CREI Valuation
In each month, the graph plots the expected dividend growth over the next 10 years, expressed as a per year quantity, that is implied by the price-dividend ratio in that month and the expected return on global real estate in that same month according to the 5-factor model (estimated over a 60-month window ending that month).

5.1.9 Conclusions

Global and particularly U.S. real estate stocks have enjoyed high returns over the past 20-plus years. Globally, real estate returns were commensurate a properly-matched portfolio of stocks and bonds, while in the U.S., REIT outperformed such a portfolio. Real estate is more stock-like than bond-like, contrary to folk wisdom. Real estate risk evolves dynamically. The interest rate risk of real estate has risen significantly over the last five years.

In both the U.S. and globally, stocks and bonds imperfectly explain real estate returns, leaving a large "real estate factor" unaccounted for. Because of this real estate factor, there are gains from diversification that can be obtained when adding real estate to a portfolio

of stocks and bonds. A portfolio that is constrained to have no more than 35% in bonds, prominently features real estate at the expense of stocks and has a higher Sharpe ratio than the portfolio without real estate.

We conjecture that the real estate risk factor is related to macro-economic activity. Real estate cash flows fell sharply during deep macro-economic downturns and investors require compensation for this risk. This is a feature real estate stocks have in common with (small) value stocks. Consistent with this observation, real estate returns are correlated with (small) value stock returns, both globally and in the U.S., and controlling for small and value stock factors reduces the size of the unexplained return component in real estate significantly. Therefore, it makes sense for the management of the Fund to conceptually treat real estate investing in similar fashion to its investments in small value stocks.

Real estate does not appear to be a good inflation hedge in the U.S. at short horizons, but its inflation hedging capability improves substantially at longer horizons.

Finally, our valuation analysis suggests that commercial real estate looks expensive in absolute terms, relative to net operating income, and especially when taking into account current levels of risk. Current prices can only be justified under aggressive, arguably implausible, assumptions on long-term NOI growth.

5.2 Global infrastructure

Next, we study risk and return of global infrastructure. As we did for real estate, we focus the risk-return analysis on the listed equity component of infrastructure. As shown in the market size section, a much larger fraction of the investable infrastructure universe is publicly listed, compared to real estate. In the interest of space and due to the somewhat shorter sample period, we limit attention to a full-sample, static analysis.

5.2.1 Data

We study three global infrastructure indices. Since the composition of infrastructure indices is quite heterogeneous, it is important to precisely define what is covered. Our main series, the MSCI World Core Infrastructure Index (abbreviated as CII) combines publicly listed companies active in two industries: Utilities and Transportation Infrastructure. Utilities consists of five sub-industries: Electric, Gas, Water, Multi-Utilities, and Oil & Gas Storage & Transportation. Transportation Infrastructure consist of the following sub-industries: Railroads, Airport services, Highways and railtracks, and Marine ports and services. The CII also contains a small component referred to as Specialized REITs, which are two U.S. firms that operate cell phone towers. We consider this to be part of Telecommunications

Infrastructure. The weight on Utilities is capped at 60% and the Infrastructure sector weight at 40%. Each subsector weight is capped at 15%. All securities are weighted by their free-float market capitalization, capped at 5% for an individual security. The stocks included in the Core Infrastructure index make up about 3.2% of the global stock portfolio.[33] Panel A of Table 4 provides the composition of the CII as of May 6, 2015. The CII returns are only available from December 2003 until May 2015, a total of 138 months, so a comprehensive analysis of the same statistical quality as the preceding analysis for real estate is not possible and results are subject to this caveat.

To mitigate the statistical issues and to highlight differences in the definition of what constitutes infrastructure, the second index we consider is the MSCI World Infrastructure Index, which we abbreviate as II to distinguish it from its "Core" CII counterpart. Its returns are available for a longer period from January 1999 on (197 months). The II has a different sectoral composition than the CII, as listed in Panel B of Table 4. It owns almost entirely the same Utilities companies, but they represent only 51% of the II compared to 60% for the CII. The Transportation Infrastructure component represents the largest difference. It accounts for only 2.7% (12 positions) of the total versus 30.2% for the CII (32 stocks). Telecommunication Infrastructure (Alternative Carriers, Integrated Telecommunication Services, and Wireless Communication Services) accounts for 44.4% of the II, but only 9.6% of the CII. The II also contains Social Infrastructure in the form of Education Services and Health Care Facilities, accounting for the remaining 1.9%.

The third and last infrastructure index we consider is the MSCI Emerging Markets Infrastructure Index (EMII). We include this index since we were specifically asked to comment on emerging markets infrastructure in our mandate. The index currently holds 110 positions, chosen from the emerging markets equity universe (mid- and large-cap securities across 23 Emerging Markets), and data coverage also starts in January 1999. In terms of sector composition, Telecommunications Infrastructure makes up 63.2% of the index, more than in the II (where Telecom represents 44.3%) and far more than in the CII (9.6%). This is due to the large Wireless Telecommunication position of the EMII (47%). China Mobile, the largest position in the fund, represents 18% of the index, and the second largest position, America Movil from Mexico, represents 7.2%. Utilities represent a far smaller 21.9% of the EMII, followed by Transport Infrastructure (8.5%), Social Infrastructure (4.1%), and Other (2.3%).

---

[33]Their weight fluctuates between 2.49% and 3.98% of the FTSE Global All Cap Index over the period December 2005 to May 2015.

Table 4: **Composition Infrastructure Indices**

The Table shows the sector composition, in terms of index weights, of the MSCI World Core Infrastructure Index (CII), the MSCI World Infrastructure Index (II), and the MSCI Emerging Markets Infrastructure Index (EMII). It lists the Global Industry Classification Standard (GICS) sub-sector codes and names, as well as the weight of each subsector on May 6, 2015.

GICS code	GICS sub-industry	CII	II	EMII
55101010	Electric Utilities	14.4	21.3	12.4
55102010	Gas Utilities	11.1	2.9	4.7
55103010	Multi-Utilities	13.9	15.4	
55104010	Water Utilities	4.1	1.1	2.3
10102040	Oil & Gas Storage & Transportation	16.7	10.5	2.5
	Utilities	**60.2**	**51.1**	**21.9**
20304010	Railroads	14.1		
20305010	Airport Services	3.6	0.8	2.8
20305020	Highways & Railtracks	10.4	1.6	2.5
20305030	Marine Ports & Services	2.0	0.3	3.2
	Transport Infra	**30.2**	**2.7**	**8.5**
40402070	Specialized REITs	9.6		
50101010	Alternative Carriers		1.1	
50101020	Integrated Telecommunication Services		32.4	15.8
50102010	Wireless Telecommunication Services		10.8	47.5
	Telecom Infra	**9.6**	**44.3**	**63.2**
25302010	Education Services		0.1	
35102020	Health Care Facilities		1.8	4.1
	Social Infra	**0.0**	**1.8**	**4.1**
	Other	**0.0**	**0.0**	**2.3**

5.2.2 Return Analysis

Panel B of Table 2 shows the longest sample of infrastructure return data: for 1999-2015, we have both the II and the EMII available (last two columns). Over this period, the II for (the developed market) performed poorly with an average annual return of only 4.1% while the EMII performed very well, with average annual returns of 10.8%. For comparison, global stock returns were in between with an average return of 6.4%. The volatility of II returns was 14.4%, below that of stocks (16.0%), and EMII (20.7%). The Sharpe ratio on EMII was 0.41 far above the 0.11 on the II. Driving the weak II and global stock returns over this sample are the technology bust of 1999-2001 and the ensuing recession of 2001-02. As we discuss below, the II performs much better once the 1999-2003 period is excluded. Maybe surprisingly, there seems to be less negative skewness for EMII than for the other asset classes (except bonds). The II has a correlation of 82% with global stocks, 33% with bonds, and 58% with real estate (CREI). Similarly, the EMII index has a correlation of 83% with global stocks, 26% with bonds, and 63% with real estate (CREI). These correlations leave open benefits from diversification when adding infrastructure to a portfolio of stocks, bonds, and real estate. Over this period, the II and EMII had a correlation of 75%, so that there are diversification benefits from combining developed and developing markets infrastructure as well.

Panel C of Table 2 looks at the most recent period (December 2003–May 2015), when we also have the Core Infrastructure index (CII) available. The CII performed exceptionally well over its short history, with an average annualized return of 12.1% and a low volatility of 13.4%. The Sharpe ratio on the CII of 0.80 exceeds the high Sharpe ratios on the other two infrastructure indices (0.61 for the II and 0.62 for the EMII), that on real estate (0.47), stocks (0.50), and bonds (0.47). As we noted above, the CII contains a much higher weight on Transportation Infrastructure and a lower weight on Telecommunication Infrastructure. This difference resulted in 330 basis points higher average returns per year for the CII but only 0.8% higher volatility. The EMII performs very well with 12.7% returns but also higher volatility than the other infrastructure indices. Despite the differences in industry composition, the CII and II returns have a correlation of 93%. The CII and the EMII have a correlation of 81%, while the EMII and the II have a correlation of 78%. The CII also has high correlation with stocks (90%), with real estate (84%), and with bonds (57%) over this period. Judged by the correlations, the CII seems to offer fewer gains from diversification than the II or EMII when added to a portfolio of global RE, stocks, and bonds, but it has had better return performance (per unit of volatility).

5.2.3 Factor analysis

Table 5 studies the two- and five-factor models for global infrastructure. Panel A studies the 1999-2015 period. In Column (2), we find that the II has a stock beta of 0.72 and a bond beta of 0.35. Both are measured precisely. Interestingly, the II has a significantly negative exposure to both the size and value factors. The latter two factors help raise the explanatory power of the model from 68.7% in Column (1) to 75.4% in Column (2). Unlike real estate, infrastructure behave more like large growth companies than small value companies. Above, we commented on the poor return performance of the II. This poor performance shows up as a negative two-factor alpha of −18bps per month in Column (1). However, once the other factors are gone, we notice that the underperformance of the II disappears. Indeed the last row shows that the expected return in the 5-factor model is 4.3%, which is very close to the observed average return but 1.8% lower than in the two-factor model. So what appears to be underperformance according to the two-factor model is not in the five-factor model.

The EMII has more stock market and less bond market risk than the II. EM infrastructure stocks also behave like growth companies rather than value companies, but they have a positive exposure to the size factor. These exposures result in an expected return of 6.0% according to the five-factor model. Since average returns were 10.4%, there is a very large alpha of 4.4% per year for the EMII. Because of the volatility of EMII returns, the alpha is hard to distinguish from zero.

Panel B of Table 5 studies the 2003–2015 period, when we have all three infrastructure indices available. First, without the 5 years of data between 1999 and 2003, the outperformance of all infrastructure indices is positive. It is even statistically different from zero for the CII. The 3.2% five-factor alpha for CII is also economically large. Second, the interest rate risk of infrastructure is substantial in this period (bond beta of 0.5; a very similar value to the one for real estate over this period), echoing our earlier finding that bond market risk of real estate has been rising over the recent past. Third, infrastructure loses its growth tilt over this period. The II and CII indices still have a significant large stock tilt. Both the global value premium (1.3% per year) and the global small stock premium (0.9%) are small over this episode so that size and value factor exposures do not affect expected returns much. Fourth, infrastructure stocks load positively (and in two of the three cases significantly) on the momentum factor. Given the large global momentum factor premium over this period (5.2%), the small exposure increases expected return modestly. The five-factor model explains 85% of return variation in CII returns in this period and suggests a cost of capital of 8.8%. The relatively small unspanned factor in infrastructure assets suggests that the gains from diversification are modest. But it makes the strong return performance less likely to be due to an omitted risk factor.

Table 5: **Factor Models for Infrastructure**

The dependent variable is the excess return on the Infrastructure Index (II) and Emerging Markets Infrastructure Index (EMII) in Panel A and the excess return on the II, EMII, and Core Infrastructure Index (CII) in Panel B. The independent variables are a constant, the excess return on the Fama-French global stock market index, and the excess return on the Barclays Global Aggregate Bond Index, and the returns on the global size factor (smb), value factor (hml), and momentum factor (mom). The first row reports the intercept α, the other rows report risk factor exposures β. The R^2 and model-implied expected return (without alpha but with risk-free rate) are in the last two rows of each panel. Panel A is for January 1999–June 2015, while Panel B is for December 2003–June 2015.

	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: January 1999– June 2015 sample					
Index	II	II	EMII	EMII		
α	−0.18	−0.03	0.24	0.37		
t-stat	−1.06	−0.18	0.91	1.45		
β^s	0.73	0.72	1.08	1.03		
t-stat	13.99	18.24	20.68	19.83		
β^b	0.24	0.35	0.07	0.14		
t-stat	2.04	3.40	0.46	1.08		
β^{smb}	–	−0.46	–	0.25		
t-stat	–	−4.41	–	2.50		
β^{hml}	–	−0.28	–	−0.39		
t-stat	–	−4.49	–	−4.44		
β^{mom}	–	0.07	–	−0.03		
t-stat	–	1.45	–	−0.45		
R^2	68.65	75.44	68.82	72.80		
Exp. ret.	6.13	4.30	7.50	5.98		
	Panel B: December 2003–June 2015 sample					
Index	II	II	EMII	EMII	CII	CII
α	0.13	0.04	0.26	0.23	0.33	0.27
t-stat	0.85	0.31	1.03	0.93	2.36	1.94
β^s	0.61	0.66	0.91	0.94	0.69	0.70
t-stat	11.94	13.46	15.94	14.73	19.17	18.19
β^b	0.47	0.51	0.45	0.45	0.52	0.54
t-stat	3.53	4.05	2.80	2.94	4.94	5.38
β^{smb}	–	−0.43	–	0.03	–	−0.18
t-stat	–	−5.81	–	0.24	–	−2.51
β^{hml}	–	0.04	–	−0.22	–	0.15
t-stat	–	0.38	–	−1.31	–	1.54
β^{mom}	–	0.18	–	0.05	–	0.10
t-stat	–	4.35	–	0.72	–	2.35
R^2	75.57	80.55	70.50	71.15	84.26	85.47
Exp. ret.	7.37	8.47	9.61	9.92	8.09	8.81

5.2.4 Portfolio analysis

We now turn to the global mean-variance portfolio analysis with stocks, bonds, real estate, and infrastructure. We start with the Jan 1999–May 2015 sample. We use the CREI and II. The unconstrained tangency portfolio contains 27% stocks, 83% bonds, 33% real estate, and −43% infrastructure. It has an annual return of 7.3%, a volatility of 9.3%, and a Sharpe Ratio of 0.576. The negative weight on infrastructure is due to its weak performance over this period relative to real estate and stocks. If we constrain the position in bonds to no more than 35%, then the tangency portfolio contains 39% stocks, 35% bonds, 106% real estate and −80% infrastructure. Again, the result is a short position in infrastructure this time financing a very long position in real estate rather than bonds. This portfolio has a much higher mean of 12.5% but a lower Sharpe ratio of 0.536. Under no shorting constraints on II, we would be back to the portfolio with only stocks, bonds, and real estate, discussed above.

We find very different results for the EMII over the same sample period. Given its strong performance, the portfolio wants to go long EMII and short stocks. In the unconstrained case, we obtain −55% stocks, 70% bonds, 46% real estate, and 39% infrastructure. Constraining the bond allocation to 35%, the portfolio weights become −105% stocks, 35% bonds, 97% real estate, and 73% EM infrastructure. These are obviously too extreme to be practically useful, but they illustrate the desire to have large positions in both real estate and infrastructure, at the expense of stocks. The extreme positions arise also because of the high correlations among stocks, real estate, and infrastructure. They make it difficult to precisely pin down the portfolio shares.

Using *Core* Infrastructure (December 2003–May 2015), the unconstrained portfolio prescribes −107% in stocks, −26% in bonds, −33% in the CREI, and 266% in the CII. This portfolio has a Sharpe ratio of 0.95. The constrained portfolio of −56% in stocks, 35% in bonds, −23% in the CREI, and 144% in the CII has a Sharpe ratio of 0.93. Again an extreme portfolio driven by the desire to exploit the positive alpha of CII. With little unexplained variation in the CII index, there is a near-arbitrage opportunity that manifests itself in extreme portfolios. But the main message is the same: the CII has been a very attractive asset and therefore it shows up so prominently in the portfolio.

Given these extreme outcomes, it makes sense to constrain the portfolio further. We ask what the minimum variance portfolio is of global real estate, stocks, bonds, and the CII that gets the same average return as the 65-35 stock-bond portfolio (7.33%) and that caps the portfolio weights in real estate at 15%, stocks at 95%, bonds at 45%, and infrastructure at 15%. We find that the lowest minimum variance portfolio has volatility of 9.69%. It holds 4% CREI, 36% stocks, 45% bonds, and 15% infrastructure. Both bond and infrastructure

positions hit their constraint. This portfolio has a Sharpe ratio of 0.62, which is much higher than the 0.54 on the 65-35% stock-bond portfolio. We also note that if the portfolio weight on bonds is tightened to 35%, then there is no portfolio that has lower variance than the 65-35 portfolio that puts positive weight in either infrastructure or real estate. The bottom line is that the historical data provide a compelling rationale for shifting the portfolio towards real estate and infrastructure. Doing so in a way that maintains return and increases Sharpe ratio requires reducing the stock position and increasing the bond position.

5.2.5 Valuation levels infrastructure

We now repeat the valuation analysis for global infrastructure. We use the "price" and "gross" series of the II, EMII, and CII indices to construct ex- and cum-dividend returns. From those, we build dividend-price ratios, and the rest of the analysis proceeds as above. We sum dividends within the year by reinvesting them at the U.S. T-bill rate. The II and EMII pd ratio starts in December 1999 (187 months), while the Core Infrastructure pd ratio starts in November 2004 (128 months). Since infrastructure stocks face no mandatory dividend distributions the way many real estate stocks do, we expect the dividend yield to be lower (price-dividend ratio to be higher) for infrastructure stocks.

Figure 23 shows the price-dividend ratios of II and CII, alongside those of global stocks and real estate. The II price-dividend ratio follows the downward price adjustment of the stock index closely, in sharp contrast to real estate. From 2005 onwards, the CII is also available. Infrastructure valuation ratios are substantially above those of real estate in 2005–2006, but real estate catches up by 2007. All pd ratios declined precipitously in 2007–2008 as the global financial crisis hit, before staging a strong recovery starting in early 2009. The graph also shows the outperformance of the CII relative to the II in the post-crisis period. The CII stood at 29.9, the EMII at 27.4, and the II at 26.4 at the end of our sample in June 2015. These valuation ratios are in line with those of global real estate at 28.3 and below the broad equity market's valuation ratio of 41.7. Infrastructure valuations are still below their 2007 peak of 36 (CII) and 43 (EMII). In terms of simple valuation ratios, infrastructure and real estate valuation ratios track each other quite closely over the past decade.

One difference with real estate however, is the much higher average dividend growth rates on infrastructure. For the CII, nominal dividend growth averaged 7.5% over the period 2005–2015, compared to only 2.3% for real estate (CREI). Over the same period, dividend growth was 9.9% for the EMII and 4.7% for the II. Figure 24 shows large fluctuations in dividend growth on the CII (bottom right panel). Like for real estate, dividend growth on infrastructure is highly cyclical. A substantial fraction of fluctuations in the valuation ratio can be traced back to these cash flow growth fluctuations (bottom left panel).

Justifying current valuation ratio on global infrastructure We now ask what the valuation ratio on infrastructure can tell us about the market's expectation about future cash-flow growth, given a reasonable expected return. As we did with real estate we use the dynamic five-factor model as our risk model. We also assume that the long-term dividend growth rate is the observed average growth rate over the longest available sample. We assume that dividend growth is constant over the next 10 years before returning to the long-term mean. It is the annual dividend growth rate during these first 10 years we report below. Under these assumptions, we find that all three infrastructure indices are currently cheap. The implied dividend growth rate as of June 2015 (for the next 10 years, per year) is 3.9% for the II, -4% for the CII, and -33.6% for the EMII. The first number is close to the historical average dividend growth rate, while the last two numbers are far below it.

This finding may be driven by unrealistic expectations for long-term mean dividend growth. While the mean dividend growth rate on the II of 4.5% per year seems reasonable, it seems less likely that the observed mean growth rates of 7.5% for the CII and 9.9% for the EMII are appropriate long-term growth rates (that can be sustained forever). To be conservative (and in light of the short sample we have over which to estimate means), we lower the long-term mean growth rate to 4.5% for CII (to equal that of the II) and to 6.9% for the EMII, and repeat the exercise. Figure 26 shows the results. The main message is that the current valuation ratios remain very reasonable, with implied dividend growth rates which are rising for II and especially the CII but remain below their pre-crisis peak values. Implied dividend growth rates are much more modest for infrastructure than for real estate: around 6% for CII compared to double that for CREI. They are also much closer to their historical mean growth rate. Emerging markets infrastructure has become substantially cheaper because expected returns have fallen faster than prices have risen.

We are aware of reports that recent infrastructure transactions have occurred at aggressive valuations. These reports come mostly from private transactions done by private equity firms. One may wonder how this is consistent with our findings that infrastructure appears fairly to cheaply valued. It is possible that these private transactions are unusual, maybe because they reflect high prices paid for trophy infrastructure assets that make the news. Or it could be because cash-flow growth rates for the specific infrastructure projects is below that for the average infrastructure project represented in the indices.

5.2.6 Conclusion

Infrastructure performed poorly in the late 1990s and early 2000s, but has seen strong performance since then. We found some evidence that infrastructure outperformed suitably matched stock-bond portfolios in at least some of the periods we studied. The attractive
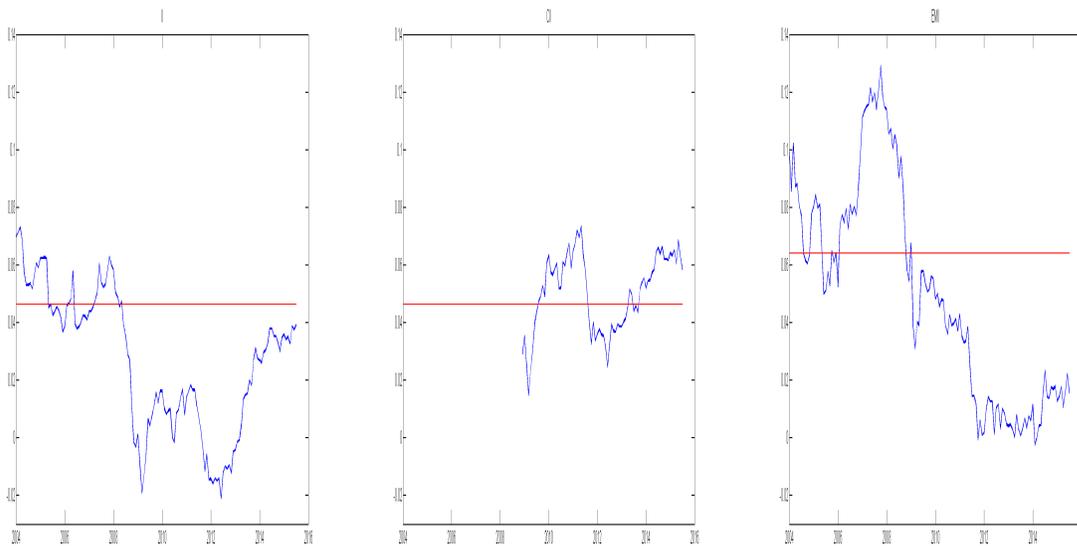
Figure 26: Expected Dividend Growth Implied in Infrastructure Valuation
In each month, the graph plots the expected dividend growth over the next 10 years, expressed as a per year quantity, that is implied by the price-dividend ratio in that month and the expected return on infrastructure in that same month according to the 5-factor model (estimated over a 60-month window ending that month). The left panel uses the Infrastructure Index (II), the middle panel the Core Infrastructure Index (CII), and the right panel uses the Emerging Markets Infrastructure Index.

returns lead to a prominent place of infrastructure in the portfolio alongside stocks, bonds, and real estate. While infrastructure valuation ratios are on the rise and are about at the same level as for real estate, infrastructure appears to more attractively valued than real estate because of its lower expected returns (reflecting its lower systematic risk) and its higher dividend growth rate. Emerging markets infrastructure has performed strongly and appears particularly attractively valued given the growth potential for infrastructure development in that part of the world. These findings need to be seen in light of the fact that infrastructure is a relatively new asset class with only about 15 years of return data. It is hard to draw robust conclusions from such a short sample.

6 Regulating real estate and infrastructure in the management mandate to Norges Bank

A second major question in our mandate is how the Ministry should regulate real estate and possible infrastructure investments in the management mandate to Norges Bank, and how Norges Bank's performance should be reviewed. Specifically, we comment on (1) the usefulness of the opportunity cost model as the model of delegation compared with the current model, (2) the use of a public versus private benchmark for real estate and infrastructure assets, (3) risk regulation, (4) reporting requirements, and (5) performance evaluation.

The key challenge, in our opinion, is how to give Norges Bank sufficient latitude to invest successfully in real estate and infrastructure assets, while at the same time regulating risk. This question gets substantially more complicated once direct asset investments rather than listed securities are in play.

We begin our analysis by assuming that a decision on a (long-term) portfolio allocation to stocks, bonds, real estate, and infrastructure has been reached, and that the Fund is in steady state after a ramp-up phase towards that allocation (we discussed earlier the separate issue of the optimal timing to reach that long-term target).

6.1 Current mandate: tracking error

Norges Bank's current mandate is to "seek to achieve the highest possible return, net of cost," subject to constraints on the allowable investment universe, restrictions on the use of leverage, and a requirement that the volatility (tracking error) of its deviations from a benchmark portfolio imposed by the Ministry of Finance not exceed 100 basis points.[34]

Advantages and disadvantages Tracking error penalizes managers for deviating from the benchmark portfolio, so if the benchmark portfolio is optimal, tracking error is a measure of how far from the optimal portfolio you are. In addition, if the benchmark portfolio is well diversified, any non-diversified portfolio will tend to have a high tracking error. However, as noted in a January 2014 letter from Norges Bank to the Ministry of Finance,[35] in practice there are often significant problems in practice with the use of tracking error. We elaborate on these difficulties here.

---

[34]The limit was originally 150 basis points, but was lowered to 100 basis points in 2009, in the wake of the financial crisis.

[35]See http://www.nbim.no/en/transparency/submissions-to-ministry/2014/framework-for-the-management-of-the-government-pension-fund-global/.

1. Tracking error is useful only when the benchmark portfolio is *optimal*, i.e., has both correct assets and correct weights. Otherwise it may lead to erroneous investment decisions. For example, the Fund's current benchmark portfolio does not include any real estate or infrastructure. Excluding asset classes from the benchmark can lead to very poor investment decisions, because the better a new asset class diversifies an existing portfolio, the *higher* the resulting tracking error. To see this, consider a new asset class (real estate, say) that is being added to the investment portfolio, but is not a part of the benchmark portfolio. Consider adding a fraction w of some new asset class, x (with expected return and volatility μ_x and σ_x respectively) to a benchmark portfolio, b, with expected return and volatility μ_b and σ_b respectively. Assume the correlation between x and b is ρ. Then the expectation, volatility, and tracking error of the return on the combined portfolio, $r_p = wr_x + (1-w)r_b$, are

$$\mu_p = w\mu_x + (1-w)\mu_b,$$
$$\sigma_p = \sqrt{w^2\sigma_x^2 + (1-w)^2\sigma_b^2 + 2w(1-w)\rho\sigma_x\sigma_b},$$
$$\approx \sigma_b(1-w) + w\rho\sigma_x \quad \text{for small } w.$$
$$\text{TE} = \sqrt{\text{var}(r_p - r_b)} = \sqrt{\sigma_x^2 + \sigma_b^2 - 2\rho\sigma_x\sigma_b}.$$

   As ρ decreases, keeping all other parameters constant, the volatility of the portfolio goes down due to the greater diversification. At the same time, the portfolio's tracking error *increases.* In other words, when we are considering adding a new asset class to achieve diversification benefits, the use of tracking error drives us in exactly the wrong direction — the higher the diversification benefits of adding the new asset class, the worse the tracking error of the overall portfolio. This can clearly be see in Figure 27, which shows the tracking error and portfolio standard deviation that result from adding 5% of a new asset class, with standard deviation 20%, to an existing benchmark portfolio, also with standard deviation 20%, for different assumptions about the correlation between the two returns.

2. It is hard to measure the return on unlisted real estate and infrastructure frequently and precisely enough to make tracking error an operational risk management tool.

3. Even if such return measures were available, the long-term nature of these unlisted investments may make the volatility of short-run deviations from a benchmark not all that meaningful. This argument may apply to listed real estate and infrastructure investments as well.
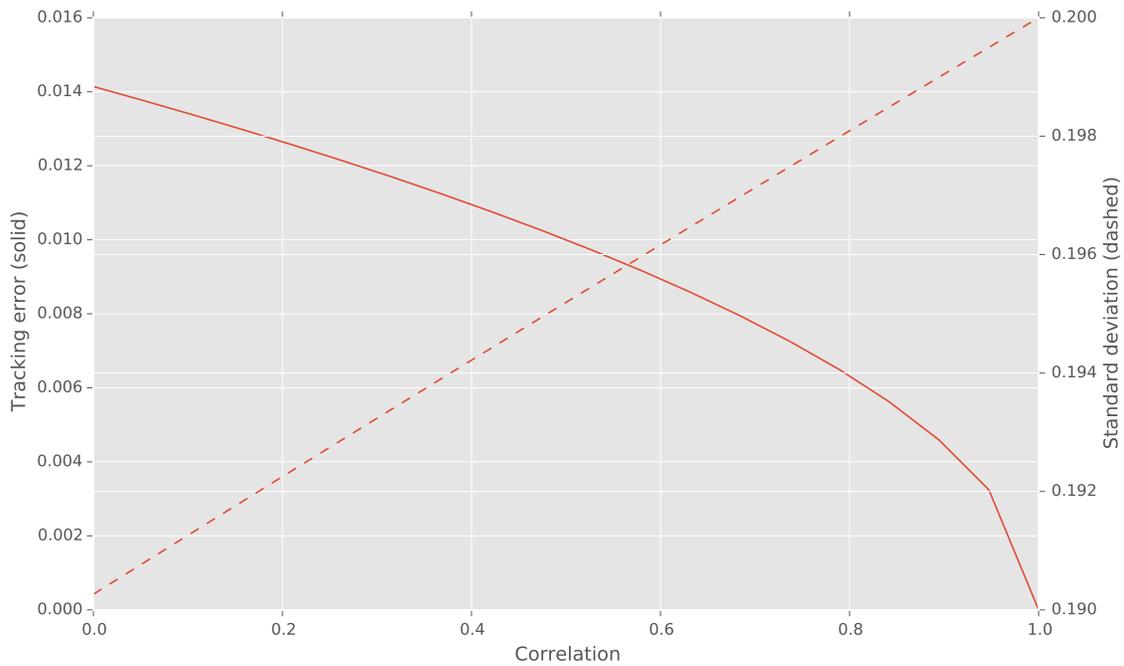
Figure 27: **Standard deviation and tracking error versus correlation.** This figure shows the tracking error and portfolio standard deviation that result from adding 5% of a new asset class, with standard deviation 20%, to an existing benchmark portfolio, also with standard deviation 20%, for different assumptions about the correlation between the two returns.

For all these reasons, we do not believe that tracking error is an appropriate risk management tool, especially when it comes to unlisted infrastructure and real estate investments.

6.2 Opportunity Cost model

In their recent review of GPFG, Ang et al. (2014a) recommend evaluating fund performance using the "opportunity cost" (OC) model, currently in use by the Canada Pension Plan Investment Board (CPPIB) and GIC Private Limited (formerly known as the Government Investment Corporation of Singapore).[36] This approach is also referred to as the Reference Portfolio approach.

The reference portfolio approach is top-down, and initially assumes all assets are invested in a liquid stock-bond index portfolio. That becomes the starting point for measuring the success of efforts to earn a higher return. Additional asset categories like real estate and infrastructure, and individual assets in these categories, are assumed to replace, dollar for dollar, some combination of stocks and bonds in the reference portfolio with the same degree of systematic risk. The goal is to produce a better return on a new asset (class) than the stock-bond combination it replaces, without adding much to total risk.

Using the Opportunity Cost Model   To see how the OC model works (see also Ang et al., 2014a, Appendix C), suppose we start with a 65%/35% reference portfolio of stocks and bonds, to which we wish to add an allocation of 1% in the Core Real Estate Index (CREI). From the estimates in Table 3, we can write the return on this index as

$$R_t^{\mathrm{RE}} = 0.02 + 0.82R_t^s + 0.68R_t^b + \epsilon_t. \tag{2}$$

In other words, $1 of the real estate index has the same amount of systematic risk as $0.82 of stocks and $0.68 of bonds. According to the OC model, to add 1% in real estate to your portfolio, you would have to sell 0.82% of stocks and 0.68% of bonds. To keep the total value of the portfolio unchanged, you would also need to invest $0.82 + 0.68 - 1.00 = 0.50\%$ in short-term riskless securities (e.g., Treasury bills). The new portfolio would thus have weights of $65\% - 0.82\% = 64.18\%$ in stocks, $35\% - 0.68\% = 34.32\%$ in bonds, 1% in real estate, and 0.50% in short-term riskless securities, and it would have exactly the same amount of systematic (i.e., stock and bond) risk as the original portfolio.

The betas from this regression also tell us how to calculate the *appropriate required return*, or discount rate, to apply when evaluating a particular investment. This is given by the expected systematic return, since there is no compensation required for idiosyncratic

---

[36]For a description, see Ang et al. (2014a) and the 2012 CPPIB Annual Report.

risk:

$$R^{\text{benchmark}} = R_t^f + \beta^s E[R_t^s - R_t^f] + \beta^b E[R_t^b - R_t^f]. \tag{3}$$

It is the sum of the risk-free rate (time-value-of-money compensation) and a risk premium for each systematic risk component, each of which is the product of the beta and the factor risk premium. An investment has outperformed over some period if its return exceeds this benchmark return.

As an aside, the model in equation (3) extends the OC model discussed in Ang et al. (2014a) because it adds a long-term bond factor. Ang et al. (2014a) only include a stock market factor and a short-term risk-free bond. They estimate a one-factor model of excess returns on the asset in question on excess stock market returns. The funding mix for \$1 of real estate is \$$\beta^s$ in stocks and \$$(1-\beta^s)$ in short-term bonds. We estimate a two-factor model for excess returns. The funding mix for \$1 of real estate is \$$\beta^s$ in stocks, \$$\beta^b$ in long-term bonds, and \$$(1-\beta^s-\beta^b)$ in short-term bonds. We think this is an important extension given the importance of long-term bond market risk for real estate and infrastructure investments. We note that the exact choice of a short-term risk-free rate is probably not crucial. A natural choice is the U.S. T-bill rate, which is widely considered to be the world's risk-free rate (at least when expressed in dollars).

Advantages and disadvantages   Ang et al. (2014a) discuss the advantages of the Opportunity Cost Model in some detail. In summary, it allows for investments to deviate from a simple fixed-weight benchmark (e.g., at times when an entire asset class is deemed relatively cheap or expensive), without being immediately penalized as in a tracking-error setting, but all new investments are correctly evaluated relative to a equivalent (in terms of systematic risk) benchmark investment in listed assets (stocks and bonds). The approach also has the advantage of piercing through asset-class labels such as real estate or infrastructure, because it deconstructs each individual investment in terms of the amount of risk exposure it has to well-understood risk factors such as equity markets and bond markets. The OCM approach prevents equity risk or interest rate risk from creeping into the portfolio via unlisted real estate or infrastructure investments. The approach induces senior management to only undertake real estate and infrastructure investments if an extra return can be achieved (in expectation) holding constant the overall portfolio's systematic risk.

However, while it enforces keeping the systematic risk of the portfolio constant, the OCM says nothing about idiosyncratic risk (or sources of systematic risk besides stocks and bonds), or about the optimal scale of any new investment. For example, suppose one replaced the entire 65% in stocks in the reference portfolio with a 65% investment in a single stock with a stock beta of 1, a bond beta of 0, and an R^2 from regressing its returns on the stock and

bond indices of 50%. Then the new portfolio's systematic risk would be the same as it was before, but one would now be exposed to large amounts of idiosyncratic risk that was not there before. Indeed, the total portfolio volatility would have gone up by over 40%.[37] The OCM thus needs to be used in conjunction with some restrictions on allowable investments to prevent exposure to excessive idiosyncratic risk.

Estimation Uncertainty   We have described how to apply the Opportunity Cost model once armed with the parameter estimates in Equation 2. However, it is important to note that these values are only estimates of the true regression coefficients, each with significant uncertainty. For the regression reported in Table 3, for example, while the three regression coefficient point-estimates are 0.02, 0.824 and 0.675 for α and the two betas, respectively, as used in Equation 2, the corresponding 95% confidence intervals for these parameters are, respectively, $[-0.412, 0.452]$, $[0.680, 0.968]$ and $[0.411, 0.939]$, so in reality we are far from certain about these values. The factor risk premia are also subject to estimation uncertainty, and are in fact harder to estimate reliably than are the betas (see Merton, 1980). One mitigation is to measure factor risk premia over (much) longer periods than the period over which the beta is estimated, since the latter period is constrained by availability of return data.

Despite this uncertainty, if everyone were running the same regression shown in Table 3, it is reasonable to assume that everyone would use the same parameter estimates in their implementation of the OCM. However, in practice there are a lot of implementation decisions that go into calculating these parameter estimates. For example, what estimation period should one use? What frequency of data (e.g., weekly versus monthly)? What indices should one use for the factor returns? The exact results obtained will be different for every possible answer to these questions, and may vary substantially. For example, the three parameter estimates from the regression above become $(0.0790, 0.8109, 0.7278)$ if we change the estimation period to 12/97–6/15, and $(0.1448, 0.6651, 1.1530)$ if we use the 5-year period 6/10–6/15. A pragmatic approach is to use as much data as possible, but to compute the required rate of return under a range of assumptions. A wide range of funding-cost estimates indicates uncertainty about the investment's potential for outperforming the appropriate stock-bond benchmark.

Below we give a detailed example of how the OCM could be implemented to assess a hypothetical Manhattan office purchase.

---

[37]An R^2 of 0.5 means that the variance of the systematic and idiosyncratic components of the stock's return are the same, so its total variance is twice that of the stock index. Equivalently, its volatility is $\sqrt{2} \approx 1.41$ times that of the stock index.

Governance and Implementation Issues   Since these parameter estimates determine the benchmark return to apply to a given investment, they are very important in determining whether an investment should be made in the first place, and also for evaluating its performance after the fact. This leads to potentially serious conflicts of interest if the same people are estimating these parameters and making investment decisions. As a result it is important that these functions be performed by different, completely independent groups within Norges Bank. The importance of independence is also stressed by CPPIB (see CPP Investment Board, 2015, p. 29), who note that "The Total Portfolio Management department identifies benchmarks that best represent each active program and are operationally feasible. They do this independently of any investment department. They recommend these benchmarks for approval by the Investment Planning Committee and then the Human Resources and Compensation Committee."

We find ourselves in agreement with Ang et al. (2014a, pp. 75–77) who advocate for a portfolio-construction group that would determine the funding basket for each investment opportunity independently of the investment teams. "The group needs to be familiar with each investment program and have sufficient resources to operate within the time constraints inherent in transaction-oriented situations. There should be a culture of collaboration between the groups so that information flows freely between them to enable the best long and short decisions for the total portfolio. This structure creates an inherent tension between the portfolio group and the investment teams—a tension that is healthy as long as there is a balance of capabilities and respect across the groups, and there is active support for the portfolio construction function from the senior management team." In our view, the portfolio-construction or management group could be part of the risk-management department of Norges Bank. or at least work closely with that department.

How is real estate funded in the OCM?   As just explained above, the Opportunity Cost model prescribes how to adjust the portfolio weights in stocks, (long-term) bonds, and T-bills for a given weight in real estate, in order to give the portfolio's overall systematic risk constant. We now apply this idea to data on the global Core Real Estate Index for the period December 1994–June 2015.

Figure 28 varies the portfolio weight on real estate from 0 to 25% (as indicated on the horizontal axis), each time adjusting the other portfolio weights according to the OC model. The left panel plots the resulting portfolios' mean return and volatility in the left panel and portfolio Sharpe Ratio in the right panel. Both portfolio mean return and return volatility rise as one increases the real estate portfolio share. The increases are small: returns move from 7.56% to 7.62% as the RE weight increases from 0% to 25%. Volatility increases slightly

Figure 28: **Portfolio characteristics versus real estate holdings**

faster, from 10.53% to 10.86%. The modest return changes are a reflection of the small alpha of RE relative to the appropriate combination of stocks and bonds. The modest increase in portfolio volatility arises from two off-setting forces: extra diversification benefits lower volatility, while unspanned real-estate risk increases volatility. The latter effect dominates but the effect is quantitatively small.

The Sharpe Ratio in the right panel is non-monotonic and peaks at a RE weight of 5% at a value of 0.4750, compared to 0.4744 for the 65%-35% Reference Portfolio. This maximal SR portfolio invests 60.88% in stocks, 31.62% in bonds and 2.5% in T-bills. These weights imply that the real estate allocation is mostly funded by a reduction in the equity allocation. The 5% weight in RE that maximizes the SR happens to be the weight that the GPFG currently has allocated to real estate. The (in-sample) Sharpe ratio is fairly flat in the real estate weight, however, ranging only between 0.465 and 0.475 for a large range of real estate weights.

How is Infrastructure Funded? Using the exact same procedure, we can add Infrastructure to the portfolio. We use the Core Infrastructure Index as our infrastructure investment. We use the results from Table 5, Panel B, Column (1) which show that to fund a $1 investment in the CII, the OC model suggests selling $0.61 stocks, $0.47 long-term bonds, and

buying $0.08 in T-bills. Let w^r be the chosen weight in real estate and w^i the chosen weight in infrastructure, then the OC model implies the following weights for stocks (w^s), bonds (w^b), and T-bills (w^f), starting from the 65-35 Reference Portfolio without risk-free assets:

$$w^s = .65 - .8240w^r - 0.6146w^i,$$
$$w^b = .35 - .6751w^r - 0.4661w^i,$$
$$w^f = \quad 0 + .4991w^r + 0.0807w^i.$$

It is easy to very that the new portfolio weights sum to one for every choice of (w^r, w^i) and that all such portfolios have the exact same amount of systematic risk (defined as stock and bond risk).

Figure 29 shows the results, constraining both w^r and w^i not to exceed 15% (so the maximum total allocation to "alternative assets" is 30%). The top left panel plots the mean portfolio return; it increases both in the RE direction and in the Infra direction, but more so for Infra. The portfolio return ranges between 7.2% and 8.2% per year.
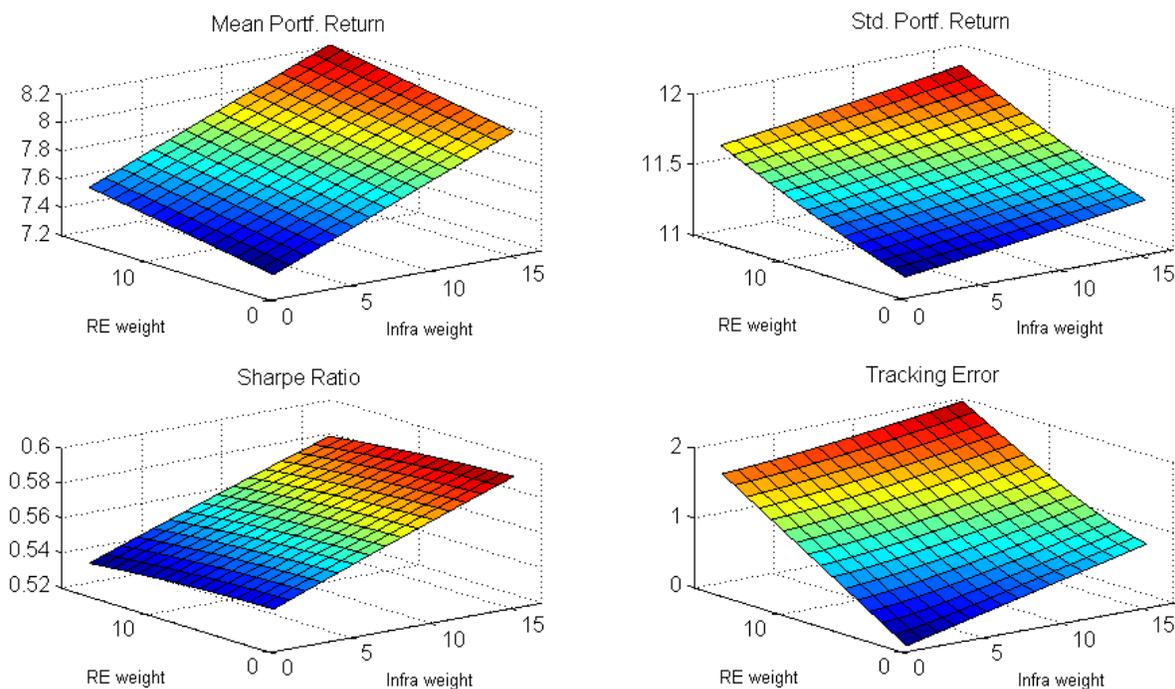
Figure 29: **Effect of varying real estate and infrastructure proportions.**

The top right panel is portfolio volatility. It again increases in both directions, but more so in the RE direction. Total portfolio volatility is lowest at 11.1% for the 65-35 Reference

Portfolio and highest at 11.9% for the portfolio that invests 15% in RE and 15% in Infra. It turns out that when one uses the OC model to determine portfolio weights on stocks, bonds, and T-Bills, one ends up increasing portfolio volatility by adding either real estate or infrastructure. However, the increase in volatility is modest; even for a 30% alternatives portfolio weight, the volatility is only 0.78% per year higher. By keeping the systematic risk of the portfolios unchanged, one is constraining total portfolio volatility substantially. Of course, buying individual RE or Infra assets may add more volatility than buying a well-diversified global listed RE and Infra index.

The bottom left panel shows the Sharpe Ratio. It is flat in the RE dimension and sharply increasing in the Infra direction. This is consistent with our earlier finding that the CII has strong positive alpha in sample. The SR is maximized by going to the corner of 15% infra and 0% RE. At that allocation, it is 0.592 compared to 0.541 for the 65-35% stock-bond portfolio. The SR gains are substantially larger from adding Infra than from adding RE, at least in the 2003–2015 sample period.

Finally, in the bottom right panel, we calculate the tracking error (TE) of the portfolios relative to the 65-35% Reference Portfolio. We see that TE rises more steeply in the RE than in the Infra dimension. This is because of the higher volatility of real estate and its larger unspanned component. A maximum tolerance for TE can help pin down the maximum permissible RE and Infra portfolio weights. The TE ranges from 0 to 2.04% per year (204bp). Because we are implicitly assuming that all stock, bond, RE, and Infra investments are in the form of listed indices, this TE only measures the *incremental* TE from a changing asset mix. If one wanted to allow, for example, no more than 100bp TE from such asset class tilt, then the graph implies maximum weights on RE and Infra. For example, an investment of 10% in RE and 0% in Infra would remain just below the 100bp TE limit. A tilt of 15% Infra and 0% RE would remain well below the 100bp TE threshold. Likewise, an investment of 12% infra and 5% RE would remain just under 100bp TE.

6.3   Our recommendation

6.3.1   Stocks and bonds

The use of tracking error, with or without the opportunity cost model, is useful in enforcing discipline in liquid, listed markets where it is easy to measure asset prices, and where it is a relatively simple matter to choose a good, well-diversified benchmark portfolio. As such, we suggest that the fund continue using its current benchmark/tracking error targets for stock and bond components of its portfolio. We also recommend considering the use of the Opportunity Cost Model (OCM) suggested by Ang et al. (2014a) and endorse their call for

slightly higher tracking-error limits.

Funding real estate and infrastructure investments with stocks and bonds, as prescribed by the OCM, could result in higher tracking error for the remaining stock and bond portfolio. For example, if a Swiss real estate investment is funded by selling a mix of Swiss stocks and bonds, the remaining stock and bond portfolio may no longer have the same weight on Swiss stocks and bonds as the Reference Portfolio. This would result in higher tracking error for the stock portfolio and for the bond portfolio. A tight tracking error limit could ultimately affect the usefulness of the OCM as an investment approach. This observation argues for higher tracking error limits on the stock and bond components of the portfolio. Alternatively, every real estate and infrastructure investment could be funded by selling the appropriate combination of the global stock and bond market *indices* that make up the Reference Portfolio, rather than by investment-specific stock and bond positions. Such an approach would have the advantage of simplicity and it would have no implications for tracking error of the stock and bond portfolios. For these reasons, this alternative is our recommended approach.[38]

6.3.2 Real estate and infrastructure

As discussed above, there are significant practical problems associated with the use of tracking error for a portfolio of (especially unlisted) real estate or infrastructure. For these asset classes, we propose using the Opportunity Cost model, as discussed above. The OC model turns the focus on the total amount of systematic risk in the real estate and infrastructure portfolios, rather than, say, tracking error.

We now describe how the OC model can be implemented, starting at the level of individual real estate and infrastructure assets, and then aggregating up to the level of the real estate or infrastructure portfolio.

Listed investment For *listed* real estate or infrastructure investments, the above discussion for the CREI applies directly. The OC model provides both the required return for the listed RE investment and the abnormal return (alpha). Only listed investments with a sufficiently high (positive) alpha and information ratio (alpha divided by the residual volatility) should be undertaken. The variance of the appropriate required return in equation (3)

---

[38]Even though the tracking error of the stock component and the tracking error of the bond component of the portfolio would both remain unchanged, the tracking error of the combined stock and bond portfolio may increase. If the real estate portfolio is less stock- and more bond-like than the Reference Portfolio, or vice versa, the stock-bond portfolio that remains after funding the real estate investments would have additional total tracking error relative to the Reference Portfolio. This is an additional argument for increasing the total tracking-error limit on the combined stock-bond portfolio.

measures the systematic risk of this investment.

Direct investments The same logic applies to unlisted investments. However, there are some practical considerations that arise from the unlisted nature of the investment that require us to adjust the performance and risk measurement from that we discussed for listed assets.

One practical difficulty with unlisted real estate and infrastructure assets is to determine the appropriate risk, i.e., the asset's betas. There usually exists no monthly or quarterly return time series for an individual unlisted real estate or infrastructure asset, asset sales occur infrequently, and historical cash-flow information may be hard to come by. There may even be no return information at all if the project is a new development, greenfield infrastructure project, or even a major re-purposing of the asset. Even in a best-case scenario, the available return series would suffer from the same appraisal-induced biases we highlighted before.

The conceptual solution is a straightforward extension of what we teach MBA students about using the CAPM: Find a good comparable, listed investment with the same characteristics as the investment in question and with a long enough return history. Since these are equity returns, they need to be unlevered to obtain the appropriate return. Then estimate the OC model and calculate the expected return, which is the expected return to be applied to the project in question.

We apply this approach below to an imaginary Manhattan office purchase. The difficulty, of course, lies in finding an appropriate comparable. This problem is not specific to real estate or infrastructure, but would also arise in the context of private equity investments, for example. While it may not be possible to find the perfect comparable, the only requirement is that the project under evaluation has the same *systematic* risk attributes. Matching geography, property type, level of development (mature projects versus development), and currency between the direct asset and the comparable listed investment would all seem to be important to arrive at the right level of systematic risk.

These criteria should also be relatively straightforward to satisfy in practice. For example, NAREIT publishes REIT sector indices (office, retail, industrial, etc.) for the U.S. with long return histories (1994–2015). Likewise, public infrastructure indices can be broken down by geography and industry (utilities, transport infrastructure, telecom infrastructure, social infrastructure, etc.) or even sub-industry (gas utilities, wireless telecom, etc.). MSCI Barra is in the process of designing appropriate benchmark return series for direct real estate and infrastructure by asset type and geography that would also be quite useful for this purpose. It is important to note that since the comparable will typically be levered, we will need to

de-lever the equity-return series to obtain a time series of asset returns.[39]

Numerical illustration   To illustrate this method, we provide a concrete example. We imagine the Fund is evaluating an individual investment in a large Manhattan class A office building. Since this is the kind of building that SL Green, a publicly listed REIT that specializes in class A Manhattan office buildings, owns, SL Green constitutes a good listed comparable.

We note that SL Green owns a large portfolio of office buildings, so its portfolio is much better diversified than the individual Manhattan office building. This difference does not matter, however, since we are only interested in using SL Green's returns to obtain the *systematic* component of the individual office asset's return.

We obtain cum-dividend stock returns for SL Green from January 1998 (the first quarter after SL Green became publicly listed is the first quarter of 1998) until June 2015 (210 months). Average annualized stock returns are 19.2% with an annualized volatility of 37.4%. We also obtain its book value of debt in each quarter. We assume debt stays constant within the quarter. We calculate leverage as the ratio of the book value of debt each month to the sum of book value of debt plus market value of equity at the end of the month. Average leverage over the full sample is 43.8%. We then construct the return on assets each month by multiplying the stock return by one minus the leverage ratio that month. The annualized average return on assets is 11.8% with a volatility of 14.6%.[40] The sharp reduction in volatility of assets relative to that of equity comes about because leverage was very high (80%) in the last four months of 2008 and the first six months of 2009, exactly when equity returns were very low or very high (April and May 2009). In such periods of 80% leverage, only 20% of the extreme equity returns are transmitted to asset returns.

Next, we estimate a regression of the excess asset returns (unlevered equity returns net of the U.S. T-bill rate) of SL Green on a constant and the excess returns on stocks and bonds. Panel A of Table 6 reports results for the excess return *on equity* for SL Green. This would be the relevant panel for an investor who is deciding on an overweight or underweight position in the stock of SL Green itself. Over the full sample (Column (1)), we see that the 2-factor model explains about 32% of the return variation in SL Green, with a large alpha of 67 basis points per month, or 8.0% per year. We recall that this alpha reflects both skill and compensation for exposure to a real estate risk factor and/or other risk factors, orthogonal to stock and bond returns. The model implies an opportunity cost of equity capital of 11.2% for SL Green. Column (2) shows results for the more recent 2003–2015 period (139 months).

---

[39]Leverage information for REITs, for example, is available from public data sets such as Compustat in the U.S. Weights of the various REITs in the sector index are available from NAREIT.

[40]We are treating the debt as riskless for simplicity.

The R^2 over the shorter period is substantially higher, as is the estimated cost of capital.

Table 6: **Applying the Opportunity Cost Model to Direct Real Estate Assets**

The dependent variable is the excess return on equity for SL Green in Panel A and the excess return on assets for SL Green in Panel B. The dependent variables are the Fama-French global excess stock return and the excess return on the global Barclays All Bond Index, collectively referred to as the two-factor model (2F). The first row reports the intercept, α, the other rows report risk factor exposures β for stocks and bonds. The last but one row reports the R^2 of the factor model. The last row reports the expected return according to the factor model. It adds in the mean risk-free rate on a U.S. T-bill over the sample period. The data are monthly from January 1998–June 2015 (210 months) in column 1 and December 2003–June 2015 (139 months) in column 2.

	(1)	(2)
	2F 98-15	2F 03-15
	Panel A: ROE, SL Green	
α	0.67	0.42
t-stat	1.15	0.61
β^s	1.22	1.88
t-stat	5.40	6.42
β^b	0.83	0.34
t-stat	1.77	0.46
R^2	32.19	47.79
Exp. ret.	11.20	16.94
	Panel B: ROA, SL Green	
α	0.54	0.61
t-stat	2.32	2.24
β^s	0.44	0.64
t-stat	6.98	9.64
β^b	0.26	0.08
t-stat	1.69	0.38
R^2	26.75	42.62
Exp. ret.	5.28	6.53

Panel B is the relevant panel to evaluate the investment in an individual Manhattan office building. It uses the unlevered return on SL Green, i.e., the return on *assets*, as a proxy for the return on the unlisted asset in question. Comparing Column (1) between panels A and B makes clear that the stock beta is almost 3 times as high for the ROE as for the ROA, a direct consequence of leverage. Thus the assets are substantially less risky than the equity. This is reflected in a much lower expected return on assets of 5.3% compared with the expected return on equity of 11.2%. This difference is even larger in the shorter sample because leverage for SL Green increased during the financial crisis, which weighs

more heavily on the short sample.[41]

Based on the full sample asset return regression results, the OC model prescribes that a $1 investment in the Manhattan Office building should be funded by selling $0.44 in stocks, $0.26 in (long-term) bonds, and $0.30 in (short-term) T-bills. Put differently, the required return to be used in the DCF analysis for the Manhattan office building is 5.3%. This required return can be constructed reliably and at high frequency. Its volatility provides in our view the best available measure of the *systematic* risk of the individual asset. In the example, the monthly volatility of the benchmark return for the office building is 2.2% over the full sample (Column 1), which annualizes to 7.6%. As we suggested above, calculating volatility at the monthly frequency may be in appropriate for long-lived real estate and infrastructure assets. Annual benchmark returns have a volatility of 8.9% over the 17 full years in our sample. This is actually a bit higher than the 10.4% annualized monthly volatility. Over the shorter sample, the annualized monthly volatility is 10.1% while the annual volatility is 9.5%. We consider these to be reasonable numbers for the systematic risk of real estate investments.

The second column of Panel B repeats the exercise for the 2003–15 sample. The funding cost is now 6.5%, higher than the full-sample estimate because of a higher stock and a lower bond beta. As we mentioned above, it is important in practice to check the robustness of such cost of capital calculations to different estimation samples. In this example, the full-sample and short-sample estimates are reasonably close.

Lessons from this example   This example illustrates how to practically implement the OC model for an unlisted investment. It makes several choices which we find appealing, and which we believe apply more generally:

1. It uses the global stock and bond indices as the only factors to determine the systematic risk exposure. These are the exact same factors that feature in the Reference Portfolio. This approach eliminates arbitrariness with respect to which stocks and bonds ought to be sold to fund a specific investment. This approach also does not affect the tracking error of the stock and bond portfolios since the composition of these portfolios is unaffected.
2. It matches the specific unlisted investment to a good listed comparable, based on geography (Manhattan), asset class (class A office), currency, and development stage (mature).

---

[41] As an aside, Panel B shows further evidence in favor of an investment in SL Green since the firm seems to be able to generate an abnormal return on assets that is economically and statistically large: 6.5% extra ROA per year (with t-statistic of 2.32) in Column 1 and 7.3% in Column 3 (with t-statistic of 2.24). This alpha is of course not the alpha of the individual office building under consideration, since the estimated abnormal return could simply be due to the management prowess of SL Green.

3. It used the maximum amount of data and performed robustness with respect to other samples.

An independent portfolio construction group, discussed above, would perform such an analysis for every candidate real estate or infrastructure investment.

Aggregating up   Having studied individual listed and unlisted real estate investments, it is straightforward to compute risk and return at the level of the overall real estate (or infrastructure) portfolio. The return on the real estate (infrastructure) portfolio is simply the weighted average of the various (listed and unlisted) investments in the portfolio. Just like the individual investment returns, the portfolio return can be decomposed into a systematic and idiosyncratic component. Likewise, the variance of the portfolio return can be decomposed into the variance of the systematic component of the portfolio return and the variance of the idiosyncratic component. Standard diversification logic suggests that overall risk of the portfolio will be largely systematic, as long as there are a large enough number of investments in the real estate/infrastructure portfolio (at least 20). That is, even if the average investment has considerable idiosyncratic risk, the portfolio's idiosyncratic risk may be very small.

Other sources of risk   The Opportunity Cost Model specifies two key sources of aggregate risk when it defines the Reference Portfolio: global stock and bond market risk. Because every investment that the fund manager chooses is funded with the appropriate mix of stocks and bonds, the overall portfolio's stock and bond market risk is always the same as that of the Reference Portfolio. This basic fact is a major advantage for risk management. However, deviations for the Reference Portfolio may expose the portfolio to other sources of factor risk. As we discussed before, like value stocks, real estate investments may be exposed to deep macroeconomic recession risk, for example. There may be several other sources of systematic risk besides aggregate stock and bond market risk that real estate and infrastructure are exposed to. For risk management purposes, Norges Bank should measure those additional sources of systematic risk, and these should also be part of the Bank's reporting (see below).

Limiting idiosyncratic risk   While the Opportunity Cost Model keeps the systematic risk of the investment portfolio constant, it imposes no restrictions on idiosyncratic risk. Moreover, we cannot even measure the idiosyncratic risk of our unlisted investments. Nevertheless, in most cases our inability to measure idiosyncratic risk is not all that problematic since much of it will be diversified away within the real estate/infrastructure portfolios, and additional idiosyncratic risk will be diversified away when these portfolios are combined with

stocks and bonds. Appendix B provides a simulation exercise that sheds light on this question, showing that the overwhelming majority (95%) of the total variance of the real estate portfolio comes from the systematic return component, the very component that can be measured precisely and at high frequency.[42]

Of course, without some constraints, it is possible to take large, undiversified positions whose idiosyncratic risk does not diversify away. We therefore propose that the Ministry of Finance add constraints on the possible range of real estate investments that may be taken by Norges Bank. To limit the purely idiosyncratic risk associated with individual assets (say, the purchase of a very large wind farm or a very large building), there should be a restriction on the concentration of the real estate and infrastructure portfolios. Specifically, a maximum Herfindahl index — a standard measure of concentration — for the real estate and for the infrastructure portfolios should be specified.

Restricting portfolio weights   The Opportunity Cost Model also imposes no restrictions on the allocation to any given asset class. To limit very large shifts in the asset class mix, we recommend that the Ministry of Finance put maximum weights on the infrastructure and real estate allocations. These maximum weights should be set generously enough to give the senior management of Norges Bank enough leeway to deviate from the Reference Portfolio when they have a strong case for increasing the allocation to real estate and/or infrastructure. Figure 29 provided some numerical detail on the amount of tracking error that various real estate and infrastructure allocations imply. Based on those calculations, we recommend that the *maximum* weight on real estate and infrastructure should be 10% each. Alternatively, the Ministry of Finance could set *minimum* weights on both stocks and bonds.

We stress that we are *not* recommending 10% as a target allocation for either real estate or infra, but merely as an upper bound even when managers think there is enormous scope for achieving high risk-adjusted returns in these sectors. As we discussed in Section 5, we found little evidence prompting us to recommend an overweight allocation to real estate. For infrastructure, the evidence was more mixed and depended on the infrastructure index and sample period under study.

---

[42]The simulation exercise also shows that, for the typical amount of idiosyncratic risk in a project, a portfolio with a modest number of individual investment projects (say 50) will still have a non-trivial tracking error relative to the systematic return of the portfolio.

6.3.3 Comparison of current situation to OCM

In summary, we advocate a switch from the current approach, in which the management mandate to Norges Bank stipulates fixed portfolio weights for stocks, bonds, and real estate as well as a benchmark for each of these asset classes, to the Opportunity Cost Model. We end this section by recapitulating the advantages of the proposed system relative to the current one:

- The best expertise for judging portfolio allocations may lie with the investment manager rather than with the asset owner. In particular, if the investment environment changes dynamically, the party closest to the market is best suited to change allocations. By putting the responsibility for making the investment decisions with the asset manager, the OCM encourages such shifts in asset allocation in pursuit of the highest return on risk. The current model with fixed portfolio weights lays the responsibility of portfolio adjustments with the asset owner (or the Ministry as its delegate), which increases the distance from the market place and makes adjustments slower and more cumbersome. Since unlisted (and listed) real estate and infrastructure investment decisions require additional market expertise, they strengthen the case for the OCM as the better delegation framework.

- The OCM obviates the need for a real estate benchmark like the IPD. Rather, the appropriate combinations of stocks and bonds deliver effectively a new benchmark for the real estate and the infrastructure portfolio. The benchmark is easy to understand and measure because it focuses on well-understood sources of risk in stock and bond markets. The current approach relies on an undesirable benchmark, Investment Property Databank (IPD) Global Property Benchmark (excluding Norway). We find ourselves in agreement with Norges Bank's qualms with the current benchmark for the real estate portfolio.[43] As we discussed at length in Section 4, property price indices like the IPD rely on self-reporting and have a series of biases induced by the appraisal process, which make them unfit for risk regulation purposes.

- The OCM approach eliminates the need to fill asset class buckets, a problem with the current approach. Rather, it forces the manager to decompose real estate and

---

[43]From a January 2014 letter of Norges Bank to the Ministry of Finance: "Like other indices for private investments, the IPD index has a number of shortcomings. The composition of the index does not necessarily reflect investment opportunities, but will depend on which owners choose to report return data to the index supplier. The index is not replicable. It will not be possible for the individual investor to buy a small share of all of the properties included in the index. Our experience is also that the IPD index is ill-suited as an instrument in our public communication of the results of our management of the Fund's real estate investments." See http://www.nbim.no/en/transparency/submissions-to-ministry/2014/framework-for-the-management-of-the-government-pension-fund-global/.

infrastructure investments into their component risk exposures and only undertake the investments if an additional return on risk can be earned. The discipline imposed by the approach should lead to better investment performance of the overall portfolio. The current approach may induce the asset manager to invest in real estate in order to meet the target portfolio weight, even if the expected return is lower than would be appropriate under the OCM.

- Combined with maximum portfolio weights on real estate and infrastructure, the OCM facilitates risk management. The systematic risk of the overall portfolio (stock and bond risk) remains the same, irrespective of the allocation to real estate and infrastructure. Tracking error for the stock portfolio remains unchanged, as does tracking error for the bond portfolio (though not necessarily tracking error on the combined stock and bond portfolio). Maximum portfolio weights limit the overall portfolio risk compared with that of the Reference Portfolio of stocks and bonds. The current system exempts the real estate portfolio from the 100bp tracking error limit, providing little formal guidance for risk management of the real estate allocation.

- While conceptually straightforward, the OCM is more challenging to implement than the current system. It must be adapted to deal with unlisted real estate and infrastructure investments. We offer several detailed recommendations for how to do this in practice.

6.4 Reporting

Accurate and informative reporting is crucial to the transparent operation of the Fund, and has long been a hallmark of its operation. In order for the Norwegian public to be able to understand and evaluate the Fund's performance along a number of important dimensions, we recommend that Norges Bank report the following attributes of the Fund's real estate and infrastructure investments to the public every year:

- The overall portfolio weights allocated to both real estate and infrastructure investments. Since some of these investments may be levered (e.g., REITs) while others are not (e.g., direct real estate holdings), we recommend reporting these weights both in terms of equity holdings and in terms of the value of the underlying assets controlled by those equity investments.

- Returns on the real estate and infrastructure portfolios.

- Systematic returns on the real estate and infrastructure portfolios. The latter are the returns on the the systematically equivalent portfolio of stocks, bonds and T-bills, as described in Equation 2 above.

- The exposures (betas) of the real estate and infrastructure portfolios to stocks, bonds and other relevant factors (e.g., value and liquidity). These exposures could be measured dynamically, for example using the most recent 60 months of data.
- The volatility of the real estate and infrastructure portfolios. This should be calculated as the volatility of the quarterly systematic returns (multiplied by 2 to express them as an annual number). To help smooth out short-run volatility, we recommend calculating these volatilities based on the most recent 20 quarters of data. In addition, ex-ante measures of the volatility of systematic returns should be reported.
- The Herfindahl index of the real estate and infrastructure portfolio and/or the proportion of each portfolio made up by the top-5 and top-10 individual holdings.
- The total costs incurred for managing the real estate and infrastructure portfolios, as well as a breakdown of these costs into (a) listed versus direct investments, (b) external versus internal management costs, and (c) by external counter-party.
- Potential conflicts of interest with external counterparties related to the governance of unlisted assets.

While outside of the mandate of this report, we also recommend reporting a break-down of the costs associated with active and passive investment for the stock and bond portfolio components.

We believe that the proposed reporting at the portfolio, rather than asset-specific, level strikes a good balance between maintaining a high degree of transparency and not compromising the investment strategy of the fund.

6.5   Determining return expectations

What return expectations should the Ministry of Finance and the Norwegian people have for the real estate and infrastructure portfolios? And how should Norges Bank's investment performance be evaluated for these portfolios?

The above analysis provides a framework for answering this question. If the Opportunity Cost Model is adopted as we envision, the return on real estate and infrastructure portfolios (after all expenses) should be at least as large as the relevant combination of the global stock and bond index investments that were are sold to fund that investment. Norges Bank would report the stock and bond exposures (betas) of the real estate and the infrastructure portfolios periodically (see our reporting recommendations in Section 6.4). Combining the relevant exposures with the realized stock and bond excess returns results in the *systematic return* for the real estate and the infrastructure portfolios, which is to be reported as well. The difference between the actual return earned and the systematic return is a measure of

the out- or under-performance of the real estate and infrastructure portfolios.[44]

As mentioned elsewhere in our report, real estate and infrastructure investments take substantial time to come to fruition. This is an important consideration, especially in the build-up phase of the portfolio, when asset development, repositioning, or simply early-stage growth may all contribute to low returns. Therefore, we recommend that the quarter-to-quarter variations in the difference between actual and systematic return be interpreted with caution and with a view towards the asset maturity of the portfolio. A long-horizon measure, such as a 20-quarter moving average of the actual-minus-systematic return differential, may be a better measure of Norges Bank's performance in its real estate and infrastructure portfolios.

Our recommendation defines out- or underperformance relative to the two-factor OCM. As we have shown in Section 5, global real estate and infrastructure also have some exposure to other factors such as the global size, value, and momentum factors. This additional risk factor exposure results in a different expected return for the five- than for the two-factor model. For example, the CREI has a two-factor expected return of 9.55% and a five-factor expected return of 11.32% (Table 3). The CII has a two-factor expected return of 8.09% and a five-factor one of 8.81% (Table 5). As discussed in Section 5, we view the compensation for such additional factor exposure as similar to the extra compensation earned when investing in value rather than growth stocks. Since the value factor is not part of the Reference Portfolio, performance due to exposure to the value factor is treated as outperformance relative to the Reference Portfolio. Real estate has similar risk characteristics as value stocks. Since there is no real estate index in the Reference Portfolio, compensation for a "real estate risk factor" should be similarly treated as outperformance relative to the Reference Portfolio. The uncertain and time-varying nature and magnitude of the exposures of real estate and infrastructure to size, value, and momentum factors (recall Figure 16) bolster the case for this approach. However, we recommend that Norges Bank report the exposure of its real estate and infrastructure portfolios to such additional factors (e.g., value, size, momentum, liquidity). This will allow for a break-down of out- or underperformance into a component due to factor exposure and due to asset selection.

How do real estate and infrastructure allocations affect the overall portfolio return? The numerical analysis above gives some indication of the types of returns that can be earned when the OCM is used to fund investments in listed real estate and infrastructure. The top left panel of Figure 29 shows that shifting the portfolio mix towards real estate and infras-

---

[44]The approach is sufficiently flexible to deal with any combination of listed/levered and unlisted/unlevered investments. Unlisted asset purchases without leverage will have lower stock and long-term-bond exposures and higher short-term-bond exposures. This will automatically give them lower expected returns than listed (or levered unlisted) investments.

tructure over the period Dec. 2003–June 2015 would have added up to 90bp to the overall portfolio return without changing its systematic risk (as measured by its stock- and bond market risk). Constraining the choice set to a maximum 10% real estate and a maximum 10% infrastructure weight, reduces the maximum incremental return to 60bp. Constraining the choice set to a maximum 5% real estate and a maximum 5% infrastructure weight, reduces the maximum incremental return to 30bp. These numbers are incremental return from a pure asset class tilt towards listed real estate and infrastructure assets. Incremental return from superior asset selection, for example via unlisted asset purchases, could raise these numbers further. Anecdotal evidence from peer funds, discussed in section 8, suggests that additional returns of about 1% may be possible.

Finally, we note that unlisted real estate and infrastructure investments increase the complexity of the overall portfolio as well as its illiquidity. Portfolio rebalancing becomes more difficult and expense if large real estate or infrastructure assets are involved. While our discussion in Section 4 found limited support for the existence of an illiquidity premium in real estate, the Fund's asset owners ought to earn additional compensation for the lower liquidity of (unlisted) real estate and infrastructure investments.

6.6   Managerial compensation

While the issue of managerial compensation takes us somewhat outside the mandate for this report, and falls within the discretion of Norges Bank, we nevertheless feel it is important to make note of some considerations that should be taken into account. We distinguish between managers deciding on the asset allocation and managers responsible for real estate and infrastructure projects.

Management deciding on asset allocation   The two-factor Opportunity Cost model described above represents a good metric for evaluating the sector-allocation decisions. It holds the management team responsible for the asset-allocation decision to a higher standard because it measures them on the value of earning incremental return relative to the Reference Portfolio by tilting the portfolio towards real estate and infrastructure. The performance should be measured over a long enough horizon for the investment hypotheses to come to fruition, especially in light of the long lags in developing or converting new real estate and infrastructure assets. Estimating the two-factor alpha over a rolling 60-month/20-quarter period seems like a reasonable approach to reduce short-term noise in performance and provide sufficiently long-term incentives.

Real estate and infrastructure managers At the portfolio-manager level, it does not make sense to evaluate performance relative only to stocks and bonds, since this runs the risk of penalizing (rewarding) managers for simply being in real estate at the wrong (right) time. Rather, it seems appropriate to evaluate portfolio manager performance relative to what would have been achieved by a passive allocation to a liquid real-estate index. Thus, we advocate evaluating real-estate performance by calculating a manager's alpha based on a three-factor model that includes stocks and bonds, as in the two-factor OCM discussed above, but also *returns on a listed real estate index*, for example the CREI we used before,

$$R_t - R_t^f = \alpha + \beta^s(R_t^s - R_t^f) + \beta^b(R_t^b - R_t^f) + \beta^r(R_t^r - R_t^f) + e_t,$$

where R_t^r is the return on the CREI in period t, and the regression is run using the most recent 60 months (5 years) of monthly returns. Managerial compensation would increase in this three-factor alpha. Each year, the 60-month window over which the alpha is computed would shift by 12 months. The logic is that if the real estate manager cannot add a three-factor alpha over a reasonably long period, the listed real estate portfolio should be passive, a strategy which is cheaper and easier to implement.

For infrastructure investments, we recommend the same approach, this time calculating the manager's alpha relative to stocks, bonds and a liquid infrastructure index such as the MSCI World Core Infrastructure Index, CII:

$$R_t - R_t^f = \alpha + \beta^s(R_t^s - R_t^f) + \beta^b(R_t^b - R_t^f) + \beta^i(R_t^i - R_t^f) + e_t,$$

where R_t^i is the return on the CII in period t. The infrastructure manager's compensation would be increasing in the α in the above equation.

Competition for talent One important challenge associated with running an unlisted real estate and infrastructure portfolio is to recruit and retain talent. This challenge particularly pronounced in some of the most competitive labor markets where Norges Bank has an office, such as London, New York, and Tokyo. Many of the best real estate and infrastructure investment managers work in these markets for private equity (style) firms and their annual compensation can be in the seven-to-eight digits (in US dollars). Even if such salaries are justified by value added, they would probably not be politically feasible in the institutional context of Norges Bank. As discussed further in Section 6, some Canadian pension plans have circumvented these objections by setting up captives who are independent of the asset manager but only work on the real estate and infrastructure investments of that asset manager. Others have hired external managers to overcome this barrier, even if it

meant substantially higher management costs. Before expanding further into real estate and infrastructure, Norges Bank must critically assess how its compensation structure affects the talent it can hire and retain, and ultimately, how this will affect the performance of its direct asset portfolios.

7 Renewable energy and emerging markets infrastructure

As the last part of our mandate, we study the question of whether the GPFG ought to devote special attention to two specific areas in infrastructure investing: renewable energy and emerging market infrastructure. On both subjects, providing a rigorous answer is a difficult task. While equity indices (and ETFs) that specialize in Clean Energy and in Emerging Markets Infrastructure do exist, they generally have a short history. It is also a volatile history, typical of the uncertainties associated with the early stages of a new sector of the economy or a new asset class. The same is true for fixed income investment products like green bonds or emerging market project bonds for infrastructure, which we only briefly touch upon. Based on the available history, hard conclusions on risk and return are unwarranted and, in particular, ignore the possibility that these asset classes may be maturing and entering into a different regime. A historical analysis also ignores the fact that a new political climate and a new investor base may be developing in both areas, which brings new opportunities. Therefore, this section necessitates a more forward-looking and qualitative approach that relies more on projected demands and less on historical risk-return analysis. It would be a mistake in our view not to carefully consider these two emerging investment universes; both offer great promise. However, given the myriad of risks involved, a gradual approach seems appropriate.

7.1 Alternative energy, green building, and green bonds

The world is increasingly turning towards renewable power generation. Ambitious targets have been set in Europe, and more recently in the United States and China, to generate a larger and larger fraction of energy from clean sources in an effort to reduce greenhouse gas emissions. While it took the world decades to make progress on the science and conjure up the political will to do something about it, there does now seem to be accelerated progress on both counts.

The transition to a new energy future promises to be major and will bring both risks and opportunities for investors. The complexity arises, first, from the two-way feedback between the traditional and alternative energy sectors. For example, a low price of oil and gas lowers the incentives to invest in renewable energy and the value of existing renewable energy assets. Conversely, increasing alternative energy availability will cut into the demand for traditional energy generation and reduce the value of those traditional energy assets. Second, returns from clean energy investments are often crucially dependent on government

incentives. Uncertainty about future subsidy policies represents an important and hard to quantify source of risk.

On the real estate side, the world is increasingly turning to green building. This is an important development because buildings account for 41% of energy consumption, more than any other sector; 73% of electricity consumption; and 38% of all CO_2 emissions. Buildings also represent an *increasing* share of GHG emissions as the world shifts away from capital intensive growth towards service-based economies and the population moves from rural to urban areas. Improving the energy efficiency of commercial and residential real estate assets can therefore be a powerful engine of environmental progress. Better energy efficiency of buildings can significantly lower operational costs and increase resale value for property investors. A commercial real estate strategy that fully incorporates the philosophy of "real estate as an energy play" is not only good for the environment, it can also be an important driver of the performance of the real estate portfolio, increasing returns and lowering risk. We discuss the relevant literature on this topic below.

A common theme in alternative energy and green building is that most investment projects have both high up-front costs and high potential long-term economic (and social) returns, even compared to real estate and infrastructure (which are always relatively long-term in nature). Given its long-term investment horizon and ample liquidity, the Fund would seem to have a comparative advantage in such projects. This needs to get balanced against the risk inherent in clean energy investments. Good risk management can mitigate both the short-term and long-term risks, for example by pursuing a diversified portfolio strategy that invests in many clean energy technologies simultaneously, invests through both equity and fixed-income products, and invests at the different stages of the product cycle, from R&D to mature firms with a proven commercial track record.

7.1.1   Renewable energy

Stylized facts on global clean energy investment   Renewable energy accounted for 9.1% of world electricity generation in 2014. According to Bloomberg's New Energy Finance, global new investment in renewable energy was $270bn in 2014, up from a $45bn investment in 2004, which implies a 20% compound annual growth rate over the past decade. 2014 marked the fifth consecutive year with more than $200 billion in new investments. A record 103GW of renewable power capacity came online, with wind power accounting for 49GW ($100bn invested) and solar accounting for 43GW ($150bn invested). All other clean energy technologies such as biomass, biofuels, geothermal, and small hydroelectric power only attracted a few billion each, and investment in several of them shrank compared with 2013. Large hydroelectric power is excluded from these numbers. Of the $270bn in new investment,

$15bn was new equity raised by publicly listed renewable energy companies. The strong equity raising experience in 2014 may be partly driven by the 100%-plus rally in clean energy stocks between summer 2012 and March 2014. The bulk of the investment came from asset finance of utility-scale projects ($170bn), followed by $74bn from small distributed capacity, largely rooftop solar.

Renewable energy investment started in Europe but the geographic landscape is shifting rapidly. Europe's renewable energy investment was $57bn in 2014, and did not grow at all compared with 2013. The U.S. grew modestly (+7%) to $38bn. The biggest growth came in Asia where renewable investment in China was $83bn, up 33%, and the rest of Asia (excluding India) stood at $49bn, up 9%. Together, China and Japan invested $75bn in solar in 2014, half of the world's total. In Europe off-shore wind was popular, including a $3.8bn wind farm project, the largest renewable energy-generation plant in the world. Developing countries invested nearly as much as developed countries, with Brazil ($7.6bn), India ($7.4bn), and South Africa ($5.5bn) the largest among them. Several more countries invested more than a billion dollars. Clearly, renewable energy technologies have successfully spread from early-adopter locations such as Europe and the U.S. to developing economies.

The market dynamics are shifting rapidly. The price of solar energy generation has fallen by half over the last 5 years. Government support policies are waning in many countries, or at least becoming less predictable. The dramatic decline in the price of oil makes renewable energy substantially less attractive than it was just a year and a half ago. If the situation persists, gas may become more favored at the expense of coal, but also at the expense of nuclear and renewables. There are structural challenges facing an electricity distribution system which is not set up for majority-renewable energy generation. In addition to the usual risk factors in energy infrastructure investments, clean energy infrastructure faces a substantial amount of technological obsolescence risk and political risk.

Balancing against these challenges is the realization that renewables can play an important role in limiting the increase in global emissions. This problem is becoming more urgent according to the latest science and has moved to the top of the global political agenda. Finally, renewables are attracting ever more interest from large institutional investors.

According to Bloomberg New Energy Finance's 2015 New Energy Outlook, the world's power-generating capacity mix will have transformed dramatically by 2040: from today's system composed of two-thirds fossil fuels to one with 60% from zero-emission energy sources. Renewables will command just under 60% of the 9,786GW of new generating capacity and two-thirds of the $12.2 trillion of investment.

Performance of clean energy investments   Next, we investigate the historical performance of clean energy equity investments. It bears repeating that the short return history must be interpreted with caution.

There exist several equity indices that track the performance of firms engaged in clean energy production and clean energy equipment and technology provision. One example is the S&P Global Clean Energy Index, which provides liquid and tradable exposure to 30 global companies from around the world that are involved in clean-energy-related businesses. It was launched in February 2007, so that there are 102 months of monthly performance data until July 2015. Since it is one of the oldest indices of its kind and one of the broadest, we will use it as our representative clean energy index. Different indices that weigh clean energy sub-sectors differently may display different performance.

Table 7 summarizes return properties of the Global Clean Energy Index and compares them to those of the global core real estate (CREI), core infrastructure (CII), infrastructure (II), stock, and bond indices over the same period. Clean energy stands out as by far the worst performer over this period, with an annual return of -7%. A cursory look at the year-by-year performance sheds light on this weak performance. While 2013 was a blockbuster year, with a total return of 48.4%, 2008 was a disaster, with a -65% return, closely followed by 2011 with a -44% return. Returns in the years 2010 (-28%), 2012 (-16.2%) and 2014 (-5%) were also negative. The contrast with the performance of both infrastructure indices over the same period is striking.

Not only was clean energy's return very low, its volatility was also by far the highest of any asset class and almost 2.5 times that of the infrastructure indices. Finally, the sharp price contractions give clean energy the most negative skewness, indicating substantial downside risk.

In terms of the correlation properties, clean energy has the highest correlation with global stocks (83%), followed by its correlation with the CII (78%) — which has a large transportation infrastructure component — and the II index (74%) — which has a large telecommunications component. It has a substantial positive correlation with real estate (69%) and a modest correlation with bonds (39%). Clean Energy has the lowest bond correlation of all asset classes. These correlation properties suggest that Clean Energy is not a particularly effective hedge against infrastructure return fluctuations. Nor does it hedge real estate returns or general stock fluctuations.

Factor model for clean energy   Over the February 2007–July 2015 period, the global 2-factor model explains 70% of the returns on clean energy (last column in Table 8). The equity beta of clean energy is 1.59, while its bond beta is only 0.09 and not statistically

Means, standard deviations, and Sharpe ratios are annualized. Skewness is the skewness of monthly returns. Data are for February 2007–July 2015.

	CREI	CII	II	Stocks	Bonds	Clean Energy
Mean	4.50	7.45	5.33	5.92	3.76	−6.97
Standard Deviation	22.44	14.68	13.84	17.63	5.91	33.93
Sharpe Ratio	0.17	0.46	0.33	0.30	0.52	−0.23
Skewness	−0.86	−0.85	−0.71	−0.80	−0.04	−1.12
			Correlations			
CREI	1.00	0.87	0.78	0.87	0.51	0.69
CII	0.87	1.00	0.94	0.91	0.56	0.78
II	0.78	0.94	1.00	0.86	0.53	0.74
Stocks	0.87	0.91	0.86	1.00	0.46	0.83
Bonds	0.51	0.56	0.53	0.46	1.00	0.39
Clean Energy	0.69	0.78	0.74	0.83	0.39	1.00

different from zero. As expected, clean energy has a very negative 2-factor alpha of −1.35% per month or −16.2% per year. This alpha is significant despite the short sample. Clean energy stocks looks like technology stocks in the late 1990s and early 2000s with high stock market beta, low bond market exposure, and a large in-sample crash. The alpha on real estate is also negative in this period, but not statistically significantly so. the alpha for the two infrastructure indices is not statistically different from zero. According to the two-factor model, the required rate of return on Clean Energy is 9.25%, much higher than that of Core Real Estate (7.7%), Core Infrastructure (5.7%), and Infrastructure (5.2%). The higher systematic risk of clean energy translates into a higher expected return.

Clean energy in the global portfolio When we combine clean energy with the CREI, CII, stocks, and bonds in a global portfolio optimization, the portfolio puts a −38% weight on clean energy. It also shorts real estate with a weight of −40%. Infrastructure (71%), stocks (57%), and bonds (51%) make up the rest of the portfolio. This portfolio has a mean return of 11.4% per year with a volatility of 8.7% per year, and a Sharpe ratio of 1.23. The presence of clean energy matters in that a portfolio without it would only obtain a Sharpe ratio of 0.43. Needless to say, this portfolio is quite extreme and not desirable either from an economic or a political point of view. The economic issue is that it relies heavily on average returns, which are very difficult to estimate with only 102 months of returns, especially given that this was the early stage of development of this industry and of investors' understanding thereof.

Table 8: **Analyzing Clean Energy Performance**

The dependent variable is the excess return on the MSCI World Core Real Estate index in Column (1), the excess return on the MSCI World Core Infrastructure Index in Column (2), the excess return on the MSCI World Infrastructure Index in Column (3), and the excess return on the S&P Global Clean Energy Index in Column (4). The independent variables are a constant, the excess return on the MSCI ACWI Equity Index, and the excess return on the Barclays Global Aggregate Bond Index. The first row reports the intercept α, the other rows report risk factor exposures β. The last but one row reports the R^2 of the two-factor regression model. The last row reports the expected return according to the two factor model. It adds in the mean risk-free rate on a U.S. T-bill over the sample period. The data are monthly from February 2007–July 2015.

	(1)	(2)	(3)	(4)
	CREI	CII	II	Clean Energy
α	-0.27	0.15	0.01	-1.35
t-stat	-0.88	0.89	0.07	-2.25
β^s	1.04	0.69	0.61	1.59
t-stat	10.95	19.50	10.88	12.96
β^b	0.52	0.45	0.41	0.09
t-stat	2.37	4.05	2.59	0.26
R^2	77.94	85.98	76.38	69.56
Exp. ret.	7.69	5.68	5.16	9.25

Valuation A comparison of cum- and ex-dividend returns teaches us that not only were total returns very low, but so was average dividend growth on the clean energy ETF. Adding up dividends over the past 12 months and computing a 12-month growth rate each month, we obtain mean annual dividend growth of -7.7%, in line with the average return of -7%. Thus, the entire average return performance can be explained by the declining dividend growth performance. The left panel of Figure 30 plots the price-dividend ratio on the clean energy index. The dividend yield is only 1.5% on average, resulting in a very high average pd ratio of 75. There was a huge crash in clean energy between May 2008 and February 2009 when the pd ratio fell by almost a factor of 4. A sharp rebound in the rest of 2009 followed. There was another sharp fall by a factor of 4 between early 2011 and early 2012. As of July 2015, the pd ratio on the clean energy sector stood at 70.6, up sharply from a value of 50 at the start of 2015 due to a combination of a drop in dividends and an increase in ex-dividend prices. The right panel of Figure 30 shows that a lot of the price fluctuations are mirrored in dividend-growth fluctuations.
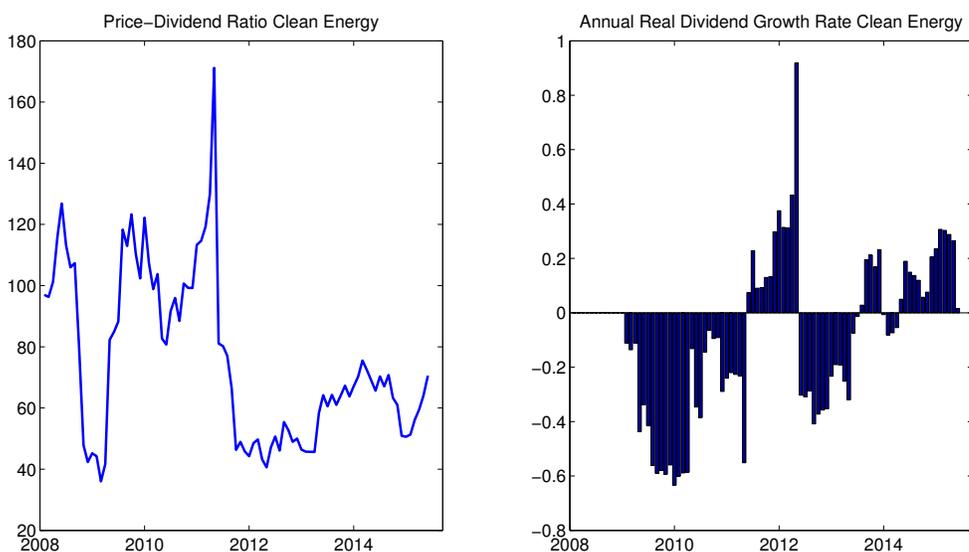
Figure 30: Price-Dividend Ratio and Dividend Growth Rate on Global Clean Energy Index
The left panel plots the price-dividend ratio on the S&P Global Clean Energy Index. The right panels plot the corresponding annual dividend growth rates. Monthly dividends are reinvested within the year and 12-month growth rates are computed from these annualized dividends. Data are monthly from February 2007–July 2015.

Using the two-factor model's expected return for clean energy of 9.25% per year, justifying the pd ratio of 50.6 at the end of 2014 requires a constant expected dividend growth of 7.2%

per year for the indefinite future. Justifying the July 2015 pd ratio of 70.6 requires a permanent growth rate of 7.85%. These are very large permanent growth rates, especially in light of the very low observed growth rate of -7.7% thus far. If growth slows down after 20 years to a long-term rate of 4% per year, 15.5% dividend growth rates for the period 2015–2035 are required to justify today's pd ratio. Put differently, given its high systematic risk, the high implied growth rate for clean energy index indicates the possibility of an overvaluation.

In conclusion, a portfolio allocation to clean energy equity must be based on the belief that its future performance will be radically better than the performance over the past 8 years. Despite the poor past performance, today's prices for clean energy stocks remain elevated and price in an expectation of very strong growth. While there are good arguments for the strong growth potential of clean energy, an open question is whether the sector will deliver the 7%-plus long-term dividend growth rates embodied in today's valuation ratio.

7.1.2 Green building

Stylized facts on green building The three leading certification programs are the LEED, the Energy Star, and the Green Star programs. The U.S. Green Building Council initiated the LEED program in 1998. It provides third-party verification that a building is designed and constructed using strategies aimed at improving performance across the following dimensions: energy consumption, water use, CO_2 emissions, indoor environmental quality, and stewardship of resources. The certification covers six different components of sustainability, including energy performance and material selection. The U.S. Environmental Protection Agency and Department of Energy initiated the Energy Star label in 1992. It focuses solely on energy management and consumption. Finally, the third measure of sustainability is the Global Real Estate Sustainability Benchmark (GRESB). GRESB was launched in 2009 by several large pension funds in an effort to achieve a more comprehensive understanding of their total exposure to environmental, social, governance and energy risks. Most large pension funds participate in the survey in return for their relative ranking. Those in the best quartile on the GRESB score distribution receive a Green Star rating. Norges Bank already supports the work of GRESB to develop a global reporting standard and benchmark for sustainability in the real estate market.

According to the U.S. Green Building Council, 45% of new non-residential construction in 2015 will be green, representing a $130 billion market. This 45% is up from 2% a decade ago. Some 3.6 billion square feet of building space is now LEED certified, and currently growing at a rate of 13%. More than 15% of single-family housing construction is green as well, a number that is projected to rise to 80% by 2018. The number of buildings that has received

an Energy Star has tripled since 2007 and represents about 5% of the U.S. commercial stock by floorspace. The trend towards more green construction is apparent all over the world, and is gaining in popularity as the awareness in the private and public sectors grows.

Green buildings consume less energy. LEED Gold buildings consume 25% less energy and 11% less water; they produce 34% lower greenhouse gas emissions. They have 19% lower maintenance costs. Some 60% of all green construction projects are retrofit projects. Firms that complete green building retrofit projects report a 7 year expected payback time of the investment and expect an increase in the buildings value of 4%. The Energy Star program claims that buildings with the Energy Star label generally consume 35 percent less energy and emit 35 percent less carbon dioxide than average uncertified buildings.

Relationship between green building and financial outcomes There is a small but growing body of work that studies the financial implications of energy efficiency and sustainability on financial performance, both at the asset level and at the portolio/entity level.

At the asset level, research by Eichholtz, Kok, and Quigley (2010) assesses the rents and transaction prices of LEED- and Energy Star-certified office buildings, relative to non-certified, comparable buildings in the United States. Controlling for differences in building quality, rents are documented to be about five percent higher for LEED-certified office properties, and some three percent higher for Energy Star-rated buildings. The reported increments for transaction prices are 11 percent and 19 percent, respectively. These findings are corroborated by Fuerst and McAllister (2011). On the other hand, using more detailed information on building characteristics, net operating income and energy prices, Jaffee, Stanton, and Wallace (2013) find that Energy Star labels do not explain significant variation in property prices once the other variables are controlled for.

At the portfolio level, Eicholtz, Kok, and Yonder (2012) investigate the impact of energy efficiency and sustainability of commercial properties on the financial performance of U.S. REITs between 2000 and 2011. The sample contains 128 REITs, of which 60% own green buildings at some point in the sample. At the end of their sample period in August 2011, they have 71 REITs who own an aggregate 919 Energy Star-certified properties and 44 REITs with a combined 708 LEED-registered properties. Of the combined square footage of property owned by all REITs in the sample, green buildings make up 6% according to Energy Star and 1% for LEED by the end of the sample. For each REIT and year, they calculate the greenness as the weight of green properties in the overall REIT portfolio for both LEED and Energy Star certification.[45] They find that operational performance

---

[45]To deal with reverse causality, whereby stronger firms have more resources to invest in retrofitting or obtaining energy efficiency certification, they employ an instrumental variables technique. They instrument the greenness with measures of local greenness and locational green policy. These measures affect the

is significantly positively affected by the REIT's greenness, after controlling for firm-level characteristics such as price-to-book ratio, property characteristics, and year fixed effects. A one percent increase in the portfolio weight of LEED-certified properties increases the REIT's return on assets by around 3.5 percent, the return on equity by 7.5 percent, and the funds from operations to revenue ratio by about 20 percent. Eicholtz et al. (2012) then ask whether the greenness of a REIT affects its stock performance. In particular, they study the alpha of a REIT, measured using a four-factor model with daily returns.[46] The results do not show any statistically significant effect of the greenness variables as measured by LEED and Energy Star certification on the abnormal stock performance of REITs. The statistically insignificant results for portfolio greenness may imply that the market already incorporates the greenness of REITs. Certification of properties is public information, so the greenness factor is already reflected in the stock price. In contrast, betas on the REIT index are negatively and significantly related to portfolio greenness. A one percent increase in the portfolio weight of LEED-certified properties lowers the market beta by 0.14, which is economically large. The authors argue that properties designated as more efficient consume less energy and other resources, so these properties are less exposed to energy price fluctuations. Second, more efficient properties have higher and more stable occupancy rates. Since the occupancy of commercial buildings is closely related to the business cycle, this may explain the fact that REITs owning more green buildings face a lower market risk, as proxied by beta.

In a very recent (and still somewhat incomplete) paper, Fuerst (2015) extends the previous analysis to a global scale and investigates a much larger sample of REITs. Also, he uses the GRESB sustainability measure, which is a much broader measure of sustainabiity measures firms make, and which may only be imperfectly correlated with the eco-labels. On the other hand, his sample is short and covers only 2011–2014; 422 REITs have GRESB scores. The analysis follows Eicholtz et al. (2012) closely. The paper finds that ROA increases by 1.3% and ROE by 3.3% for each point of improvement on the GRESB score. The GRESB score runs from 0 to 100, with an average score of 45 and a standard deviation of 19. Both results are statistically significant, even after accounting for selection effects whereby the largest REITs are the ones who fill out the GRESB survey. They are also economically large. A one standard deviation increase in GRESB score from the mean (a 34% improvement) increase the ROA by 45%, or from a mean of 2.67% to an ROA 3.88%. The ROE goes from a mean of 5.3% to 11.5%. In contrast with Eicholtz et al. (2012), Fuerst (2015) finds a

greenness of the portfolio but should not be correlated with financial performance, other that through their effect on increased portfolio greenness. They provide evidence in support of the validity of the instrument.

[46]The factors are the NAREIT equity REIT index, and the regular SMB, HML, and MOM factors we defined earlier.

positive effect on alpha but no effect on beta. One possibility is that the GRESB rating is not in the information set of stock investors (except for pension funds), unlike the eco-labels. To the extent that it reflects different value-relevant information from the eco-labels, stock prices may not reflect that information. Alternatively, the results from the global sample for 2011–2014 time period are just different from those in the U.S. for 2000–2011. Finally, the GRESB score combines both operational and managerial aspects. The paper finds it is the "Implementation & Measurement" rather than the "Management & Policy" component of the index that is associated with the better performance.

While this research agenda is very much a work in progress, the results are supportive of the idea that REITs who do good also do well. A strategy that overweights REITs with strong greenness, measured in several ways, may earn positive risk-adjusted returns. Finally, investing in green buildings makes sense from a risk management perspective. There is not only less risk from cash flow shortfalls due to vacancy and less risk in resale value, but also less exposure to volatile energy prices.

We recognize and strongly support that the Ministry of Finance already mandates (in section 2-2 on responsible investment management principles) that Norges Bank shall take into account the green aspects of its real estate investment strategy: "In its management of the real estate portfolio, the bank shall, within the environmental field, consider, among other matters, energy efficiency, water consumption and waste management."

7.1.3   Green bonds

Stylized facts on green bonds   Green bonds date back to 2007–2008, when the European Investment Bank and the World Bank's environmental department started issuing bonds to finance green projects that would mitigate climate change and advance environmental sustainability. Initially the market was small. In 2013, the World Bank's private sector arm, the International Finance Corporation, raised a $1 billion green bond, attracting the attention of institutional investors. Since then, several large companies such as EDF, Toyota, Unilever, and GDF Suez have issued large green bonds ($1–2.5 billion each) to finance renewable energy, energy efficiency, sustainable transportation and other low-carbon projects. The Dutch bank ABN Amro raised €500 million for green buildings. 2015 saw the first issuance of green bonds from India, Brazil, and China. Green bond issuance was $11 billion in 2013, tripled to $37 billion in 2014, and was forecast to hit $100 billion in 2015. However, with the first half of 2015 sales at the same level of 2014, that 2015 forecast is in jeopardy. Since 2008, the World bank has issued $8.5 billion in green bonds in 18 currencies, including a 10-year $600 million benchmark green bond and green growth bonds linked to an equity index and designed for retail investors. Separately, the IFC has issued $3.7 billion to date.

An important driver of the green bond market is the growing number of asset managers with mandates to increase investment in instruments that support low-carbon growth. For example, BlackRock was selected by Zurich Insurance in November 2013, to manage a green bond portfolio of $1 billion. Many pension funds and other institutional investors are now mandated to include responsible and sustainable investments, and green bonds are a natural way to fill some of these quota. As a result of the popularity of green bonds, they have yields about 20-50 basis points lower than those on comparable non-green bonds. From a purely financial perspective green bonds appear to be expensive.

One of the issues the green bond market is currently struggling with is standards: what constitutes a green bond? In March 2015, the International Capital Market Association, a trade group for banks, issuers and asset managers, updated a set of voluntary green-bond guidelines. The guidelines recommend that an outside party review whether a green bond is appropriate and suggest issuers regularly report on their project's environmental impact. Improving transparency and disclosure should make it easier for investors to evaluate the environmental impact of a green bond linked to a specific project.

Governments could stimulate the development of the green-bond market through tax incentives. This would make sense given that timing of investment matters for climate change and that new financial markets need some support in the early stages of development. The World Bank's initiative can be seen in this light. China recently released policy proposals that include green tax incentives and preferential capital requirements for banks with a high share of green assets. Credit enhancements offered by the U.S. Overseas Private Investment Corporation now support solar projects in Chile. Likewise, large institutional investors are making (and can continue to make) a difference in helping to spur on this market.

Green asset-backed securities   Another recent development has been the securitization of green bonds. Green asset-backed securities are green bonds backed by a pool of loans or other revenue-generating financial assets, such operating solar farms. More green asset-backed securitisation would allow banks to get loans to green projects off their balance sheet by packaging them and selling these securities to investors. Similarly, utilities could get renewable energy assets off their balance sheets or use securitization as an alternative source of financing. The easier it is for banks and utilities to get green assets off their balance sheet, the more likely they are to make additional green investments. The diversification and risk-tranching achieved by pooling the bonds is attractive to investors for the same reason it is attractive in the Commercial Mortgage-backed Securities space.

Green securitisation is also crucial because the aggregation process enables smaller scale low-carbon projects to gain access to institutional investors. One example is the Property-

Assessed Clean Energy (PACE) bond. PACE bonds allow governments to finance the up-front cost of energy efficiency improvements on commercial and residential properties, which are paid back over time by the property owners. This allows a property owner to implement improvements such as home insulation or solar panels without the large up-front cash outlay. Improvements are repaid by the home owner over the following 10 to 20 years through property assessments, which are secured by the property itself and paid as an addition to the owners' property tax bills. A PACE assessment is a debt tied to the property as opposed to the property owner(s), so the repayment obligation may transfer with property ownership. PACE bonds eliminate a key disincentive to invest in energy improvements, since many owners are hesitant to make property improvements if they think they may not stay in the property long enough for the resulting savings to cover the upfront costs. Deutsche Bank sold a $104 million PACE bond in 2014. It was rated AA, and extraordinarily well received by U.S. insurance companies and asset managers.

Another example is solar securitization. SolarCity has done several public securitizations starting in late 2013. The pools bundled both commercial and residential assets. The assets are a mix of power purchase agreements and leases.[47]

Here too there may be a role for the government by supporting the development of standard contracts and agreements for low-carbon assets because this facilitates the robust pooling of loans into large-scale securities. The EU Commission and the European Investment Bank (EIB) could enable green securitisation either by setting-up or supporting an aggregation/warehousing facility along the lines of the EIB's proposed Renewable Energy Platform for Institutional Investors. The EIB could apply credit enhancement tools to the securities. Finally, the ECB could allocate a given share of their asset-backed securities purchases to green securities. Again, the presence of a large institutional investors can make a big difference in helping to establish the green asset-backed securities market.

7.2   Emerging market infrastructure

7.2.1   Context

Like renewable energy, emerging market (EM) infrastructure brings both large opportunity and large risk. Much of the world's shortfall in infrastructure is in the developing world.

---

[47]A Solar Power Purchase Agreement (SPPA) is a financial arrangement in which a third-party developer owns, operates, and maintains the photovoltaic (PV) system, and a host customer agrees to site the system on its roof or elsewhere on its property and purchases the system's electric output from the solar services provider for a predetermined period. This financial arrangement allows the host customer to receive stable, and sometimes lower cost electricity, while the solar services provider or another party acquires valuable financial benefits such as tax credits and income generated from the sale of electricity to the host customer. With a lease, customers pay a fixed monthly fee with a guarantee of electricity production.

One major driver of global infrastructure demand is *urbanization* in the developing world. In 2010, the world's urban population was 3.5 billion people, of which 2.6 billion lived in developing countries (73%). By 2050, the urban population is expected to reach 7.3 billion people, of which 5.2 billion (83%) will be in developing countries. In other words, over these 40 years the urban population in the developing world will grow 15 times as fast as that in the developed world (2.6 billion versus 160 million). A quarter of this growth will be in Sub-Saharan Africa and another 10% in West and North Africa. A quarter will be in the Indian subcontinent with another 10% in Southeast Asia. Fifteen percent will be in China. India alone will add 500 million urban dwellers, China 400 million, and Nigeria and Pakistan rank third and fourth with about 125 million each. The U.S. ranks fifth with 100 million new city dwellers, followed by Indonesia and Bangladesh. With the exception of the United States, these are all countries with relatively low degrees of urbanization today, low GDP per capita, low government expenditures (relative to GDP), less effective government, less political stability, worse rule of law, and more corruption (Angel, 2012).

EM urbanization will give rise to huge infrastructure needs, especially given the poor state of the existing infrastructure. Every type of infrastructure will be needed: clean water and sanitation services, electricity, public transportation, road networks, telecommunication infrastructure, and social infrastructure (schools, hospitals, prisons, etc.). There will also be a huge need for modern real estate: affordable urban housing, office, retail, logistics, etc.

High infrastructure spending in developing countries can propel economic growth, through higher employment and lower trade costs, as well as promote social progress, for example by improving health outcomes. Infrastructure spending in other words generates many positive externalities. Therein lays a challenge for international investors: how can they share not only in the private returns of infrastructure projects but potentially also in the higher social returns from their investment? A potential strategy could be to view emerging markets infrastructure investments more broadly to include equity stakes in the businesses that directly benefit from the infrastructure investment, thereby capturing some of the positive externalities while at the same time increasing the likelihood that the infrastructure project succeeds.

The evolution of emerging markets to more energy-intensive and more urbanized economies has already had and will continue to have a major impact on climate change. Climate change has disproportionate adverse effects in developing countries because these countries lack an effective government response, the health care infrastructure, or the insurance markets to offset some of the impact. However, the awareness to do something about climate change in the EM economies has changed as well. As we pointed out in the clean energy discussion above, many developing economies are making large investments in renewable energy.

They can take advantage of the technological progress that has been made in the developed world. Their pool of cheaper labor has helped to reduce the cost of certain alternative energy technologies, like photovoltaic panels. Obviously, meeting the growing energy needs of the developing world through alternative energy further adds to the challenge that is emerging market infrastructure.

7.2.2 Market size

Over the last 20 years, 3.8 percent of world GDP has been spent on (economic) infrastructure. Annual infrastructure spending has been trending down in advanced economies, from 3.6 percent of GDP in 1980 to 2.8 percent in 2008, but has been rising in emerging market economies, from 3.5 percent to 5.7 percent. The latter figure is driven by particularly high fixed-capital investment in Asia, especially China (McKinsey, 2010). To respond to an increased demand for infrastructure services, Yepes (2008) calculates that developing economies must spend 6.6% of GDP, with much higher numbers for the lowest-income countries (12.5%) than for the upper-middle-income countries (2.3%).

Bhattacharyay (2012) predicts that 32 Asian economies will need $8.2tn in infrastructure spending in the 2011–2020 decade. Two thirds is needed for new capacity and one third for maintenance and replacement of existing assets. About half of this should go toward energy, about one third for transport and the rest for telecommunications (13 percent) and water (five percent). China needs more than half and India more than a quarter of these sums. Inderst and Stewart (2014) put Africa's annual infrastructure needs at $93 billion in 2009. Latin America has invested very little (2% of GDP from 1990–2010) and needs a lot of catch-up investment to stimulate growth. Overall, developing economies will need to spend 6–8% of GDP to meet their infrastructure needs, raising infrastructure spending from $800–900 billion a year in 2010 to $1.8–2.3tn a year by 2020. However the current funds may not be available to generate that kind of investment. WEF (2015) estimates the global infrastructure financing gap at about $1tn per year, while others believe $1tn is the financing gap just for the developing economies only.

7.2.3 Financing, institutional investors, and risk

Funding sources   In terms of funding, EM infrastructure is funded 70% by governments, 10% by development agencies, and 20% by private sources ($200 bn). Government funding is getting squeezed and development aid is unlikely to fill the gap. Further, Basel III is making bank financing of long-term EM infrastructure projects more costly. Many emerging market borrowers lack access to corporate bond markets. Originate-to-distribute loan-to-

bond financing deals fell out of favor after the financial crisis and seem an unlikely solution for EM infrastructure financing even today. This makes pension funds, insurance companies, sovereign wealth funds, and private equity funds an important source of private sector finance going forward. The long horizons of such funds, their need for yield and new sources of income more generally, and the possible benefits from further asset diversification make EM infrastructure investments a good match for these investors.

Institutional investors Notwithstanding the difficulties, several larger institutional investors have invested in large EM infrastructure projects, both in the form of listed equity and direct equity investments (in addition to project bonds, PPPs, and government bonds). Canada's AIMCo, for example, invested in a Chilean electricity transmission and distribution company, as well as in a Chilean toll road. The CPPIB invested more than \$1bn in five Chilean toll roads and in Peru's largest natural gas pipeline operator. Australian, Dutch, Danish, and UK pension funds have all made sizable investments in EM infrastructure projects in the last five years (see Inderst and Stewart, 2014, for more examples).

Dedicated infrastructure funds have been around since the 1990s in Australia and since the 2000s in North America and Europe. They were popular just before the financial crisis and have made a resurgence in the last few years. Between 2004 and 2013, 123 funds targeting emerging markets infrastructure were closed, with a volume of \$41 billion. As of June 2015, Preqin counted 61 fund managers seeking to raise \$27 billion outside of Europe and North America. While competition for new funds is fierce, existing infrastructure funds struggle to find good projects to deploy their capital. EM infrastructure funds are currently sitting on more than \$20bn of dry powder (Preqin, 2015b).

Risk As discussed in Section 4.3, direct infrastructure investment is subject to a number of risks not present with listed investment. These risks are especially salient in EM economies, and include political risks (e.g., changes in government or infrastructure policies, shifting popular sentiment towards privatized national services), regulatory risks (e.g., changing energy regulations, no viable PPP legal framework), and management and governance risk (e.g., corruption). The operational risks associated with emerging markets infrastructure projects are also elevated. Service disruptions or accidents associated with directly-owned infrastructure projects pose headline risk for the asset owners. Some of these risks can be mitigated at least to some extent. For example, sovereign default risk and currency risk are insurable through well-functioning financial markets. There is even political risk insurance available from private insurance companies that covers contract repudiation, asset expropriation, or political violence. Obviously, counter-party risk of the insurer must be assessed

carefully. See Section 4.3 for a more detailed discussion.

The substantial amount of available private capital in search of viable infrastructure projects on the one hand and the massive financing need for such projects on the other hand point to a mismatch between the desired and the available risk-return trade-offs currently available in emerging markets infrastructure. For example, to be induced to bear the many risks associated with say an investment in a Nigerian toll road or marine port, private investors may require so high a return as to make the project unprofitable or politically infeasible (maybe because users are unwilling to pay the required service charges). Buying insurance that mitigates some of the risk may result in a return that fails to meet the investor's hurdle rate. Governments and supra-national institutions like the IMF or Worldbank could play a useful role by offering (partial) guarantees or reasonably priced insurance for hard-to-insure risks associated with emerging markets infrastructure investment.

8   Investing in real estate and infrastructure: A manager perspective

GPFG's current investment approach is a very cost-efficient, passive implementation of a reference portfolio with a small overlay of relatively short-term active management strategies.[48] The reaction to negative deviations from market returns in 2008 indicates acceptance of passive reference-portfolio losses, but less tolerance for short-term deviations from passive portfolio returns. This very short-term orientation is limiting GPFG's long-term potential.

The experience from other larger investors is in support of an upper limit of 10% for real estate and infrastructure as recommended in this report. If Norges Bank were to use the full allocation and eventually acquire close to $100 billion of real estate and infrastructure assets, it would represent a fundamental shift towards a long-term strategy, with far-reaching implications for the culture and governance of the organization, and for reporting of short-term results.

The authors of this report have differing views on how Norges Bank should do the actual implementation of infrastructure investments (see Section 1). This reflects a fundamental difference of opinion on prospects for outperforming market returns. Some of the academic literature sees limited opportunities to outperform listed market indices. It sees unlisted investments as a second-best solution.

Most peers with large, long-horizon portfolios believe that markets tend towards long-run efficiency, but can be inefficient in the short run. Their challenge is not to be drawn into the drama of short-term asset price fluctuations, to take advantage when they happen if possible, to absorb whatever lessons can be learned from such episodes, and to stay focused on the goal of positioning for superior long-term wealth creation. They view direct investments in unlisted markets as a way to get beyond the zero-sum game of trading in listed securities, and to create wealth in good companies through greater control over governance and better alignment between owners and operators at the asset level, in part by exploiting operational efficiency and economies of scale.

Under the very demanding opportunity cost model, any difference in asset mix from the reference portfolio will give rise to larger deviations from passive return. Allocating more capital to real estate and infrastructure only makes sense if it increases return on risk. Passive risk can be structured to have the same systemic risk as the reference portfolio, but the annual return profile of the revised asset mix will by definition be different. Successful implementation implies that deviations from reference portfolio return will be larger in both directions, with an upside expected bias.

---

[48]This section reflects the views of one of the members of the committee (de Bever).

There will also be larger deviations from reference-portfolio returns from active management of the new assets. An organization supportive of active management should over time aspire to contribute a sustained 0.5%–1% to passive return. Calculations below suggest that this can be achieved by 1% additional portfolio volatility. The incremental return opportunity should come from long-term strategies that by definition cannot be replicated with a succession of short-term strategies. That typically implies a trade-off between lower short-term return and superior expected long-term returns. The practical consequence is that the benefits of this strategy may not be visible for a number of years.

One should be skeptical of GPFG's opportunity to enjoy the benefit of market liquidity, since it owns a material share of global stock markets, and a smaller share of a more constrained definition of global bond markets. Any attempt to sell more than a small part of the portfolio will move market prices, not just because of the size of that transaction, but because of the secondary reaction it is likely to trigger in other market participants. Moreover, the urge to sell never occurs to just one market player. Liquidity is no longer binary. For GPFG, listed markets are partially illiquid, and unlisted assets are now traded frequently enough to be partially liquid.

As for short-term pricing efficiency, it does not seem credible that a measure of "true" value of global listed companies like the MSCI ACWI index can drop 49% between February 2007 and November 2008 and then rise 56% by February 2010. It seems more plausible to view such price swings as mostly a reflection of human emotion ("fear and hope") and over-levered investors getting caught in a temporarily dysfunctional illiquid market. Long-term strategy is based on the idea that long-term return is the reward for periodically having to sustain painful short-term downside.

GPFG has some strong comparative advantages. It is at the upper extreme of global asset managers measured in both asset size and investment horizon. A fund with that kind of capital, which can be mostly deployed to grow for decades, and no need to transact in times of crisis, should not be overly worried by unrealized valuation changes in listed positions that are effectively mostly long-term listed holds.

The Fund should aspire to monetize these comparative advantages. Norges Bank fits the profile of the ultimate long-term asset manager. The bank should have no difficulty getting access to the best available investment talent and the best international investment platforms to innovate in the pursuit of inefficient market niches, and creating incremental value from direct investments.

This is a good time to build the internal capacity to fill the proposed maximum allocations when the opportunity arises, without having to be in a great hurry to fill them. The reference portfolio model is particularly unforgiving of filling allocations without carefully considering

the opportunity cost of moving away from stocks and bonds. Strong recent performance has increased allocations to the proposed assets faster than the supply of good investments. At the end of a declining interest rate regime, there is reason to be selective, although the interest sensitivity of these new assets can be tempered by high positive long-term correlation with economic growth. Odds are that in infrastructure, more supply will become available in the near future as public sources of funding fall short of social infrastructure needs.

Real estate and infrastructure are very heterogeneous and inefficient asset classes, characterized by less than perfect information flows. Listed investments in this space are accessible to investors with $80 thousand or $80 billion to allocate. The management mandate to Norges Bank opens up for active investments within a tracking error limit in listed assets, including real estate and infrastructure. If Norges Bank believes listed infrastructure offers superior returns on risk, this is presumably already reflected in active over-weights in the current portfolio.

The full unlisted real estate universe ranges from a $2 million strip mall asset to office towers of $1 billion or more. Unlisted infrastructure includes highly levered small court house PPPs and multi-billion-dollar toll roads. This covers assets with a wide range of risks and returns and operations costs.

Contrary to what the review of these listed and unlisted data bases recommends, one observes that large pension fund portfolios with internal management capacity hold few listed real estate and infrastructure assets beyond index exposure. They invest mostly in a subset of unlisted assets: long-term investments in large retail and office buildings in major markets with strong cash flow prospects and relatively modest risk. In infrastructure, they typically hold large quasi-monopoly assets like toll roads, transmission facilities, and power plants with strong concession contracts or a history of fair regulation. Direct investments allow participation in opportunities where the ability to commit stable capital for long periods provides opportunity not accessible to most investors.

Part of the literature attaches great importance to the lack of reporting transparency from stale dated asset values. Peer funds with a long investment horizon see this as a minor annoyance in the absence of both need and opportunity to transact in meaningful volumes at daily prices. They are more concerned about governance and alignment of interest at the asset operation level, and the ability of patient capital to create net wealth through greater direct control over conditions affecting the value creation of these assets.

Evidence from peer organizations indicates that my proposed strategy requires:

1. Strong governance support for a high-performance active management culture.
2. Concentration on long-term strategies and evaluation of outcomes over a long horizon.

The rest of this section documents innovation in peer fund strategies and the underlying

investment beliefs, their approach to real estate and infrastructure, risk management, governance, and performance measurement.

8.1 The challenges of being a long-term investor

Moving up to 10% of assets into real estate and up to 10% into infrastructure has far greater implications for governance and investment management culture than shifting percentage allocations to stocks and bonds from, say, 50-50 to 70-30. No matter how the strategy for real estate and infrastructure is implemented, it will commit GPFG to becoming a long-term investor. That only makes sense if Norges Bank is confident it can access sources of return not available to short-term investors. Since market structures are not static, that will require continuous innovation, not just imitation.

Long-term investing is like a 10 km speed-skating race: the goal is to get to the finish line first, but the exact path to getting there is not important. The current way of evaluating Fund investment outcomes treats the 10 km as a succession of 1 km races, with success measured as being first past the post every time.

To extend the analogy to desirable performance goals, the average skater will not win gold medals. One could conclude that winning is just too hard, so don't bother to try. But Norges Bank has the athletic build to be a champion. What it needs is determination, building up experience, the humility to recognize that there is a lot to be learned, and tolerance for failure along the way. Fear of failure is a barrier, but as Wayne Gretsky once observed: "I missed 100% of the shots I did not take."

Superior returns from long-term strategies by definition cannot be obtained from a sequence of short-term strategies. They sometimes involve giving up short-term return relative to a positive trend in reference portfolio return. The main practical obstacles to achieving long-term success are behavioral. The world consists of 95% second-guessers reluctant to deviate from the tried and true, and 5% "first-guessers", people always on the look-out for new and better ways to earn return. Being different can be lonely. This problem was stated eloquently (and is therefore often repeated) by Keynes (1936, Chapter 12):

> "[I]t is the long-term investor, he who most promotes the public interest, who will in practice come in for most criticism, wherever investment funds are managed by committees or boards or banks. For it is in the essence of his behaviour that he should be eccentric, unconventional and rash in the eyes of average opinion. If he is successful, that will only confirm the general belief in his rashness; and if in the short run he is unsuccessful, which is very likely, he will

not receive much mercy. Worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally."

Keynes was successful in the long run (Chambers and Dimson, 2013), he learned from setbacks in the short run, and he pursued his preferred strategy despite his own warning about potential impact on reputation, because it was the right thing to do for his clients. Long-term investing needs a supportive governance model to defuse the tension between long-term strategy and short-term results. There must be passionate support for the belief that it is possible to create an organization with a high performance culture that can do significantly better for asset owners in the long run.

8.2 The importance of supportive investment beliefs

Electing not to pursue an active investment strategy rests on the following beliefs:

- Markets are efficient. The average active asset manager earns market returns net of implementation cost. Doing better than average is too hard to bother with.
- Any persistent anomalies from market efficiency can be duplicated with cheap, liquid exposure to return factors: equity, bond, value vs growth, small cap, low volatility.
- Listed markets offer the same opportunities as unlisted markets, and there is a cost to giving up liquidity.

Successful long-term large pension and endowment organizations strongly support the potential value of active management and hold the following beliefs:[49]

- Markets tend to efficiency in the long run. Short-term listed asset pricing efficiency is imperfect because it requires both instantaneous incorporation of all relevant information and one unambiguous way to reflect new information into asset valuation.
- Market prices are best viewed as a reflection of a shifting balance in the activity of investors with different views of risk and return. Prices at times reflects mostly the exuberance of the marginal buyer (e.g., 1999), or despair of the last seller (e.g., 2008).
- If markets are not completely efficient, active management can add value. Large investors can exploit economies of scale, lower cost access to internal and external talent, and access to the best research and information to drive innovation.
- Access to stable capital that will not be needed for many decades allows for a long investment horizon, and provides opportunities not available to managers that may

---

[49]Most of the examples are drawn from Leo de Bever's own experience in Canada and the Netherlands, e.g., the investment principles of APG (https://www.apg.nl/en/article/investmentprinciples/478). A European perspective on culture and governance can be found in Clark and Urwin (2008).

have to return capital at short notice: long-term investors have the freedom not to transact in times of crisis.

The recent emphasis on "risk factors" as an easy way to capture known market inefficiencies likely overstates the opportunity to apply this approach consistently. Most institutional managers are fundamentally value investors, but they know that this predisposition does not pay off in some market environments. Momentum is famous for working until it stops working. Until someone develops a reliable recipe for switching these factors on and off, they represent just another way of framing the active management challenge.

8.3   Peer fund innovation efforts to maximize return on risk

Imitating what made others successful will not guarantee future success. Market structures and the state of investment and risk management technology are not static. The best funds have a history of pushing the technological and investment management boundaries to find new value added opportunities.

Historical data should be a partial guide to that process, but any innovator must by definition confront the fact that he is entering uncharted territory.[50] A UK institutional manager phrased it like this: "Why have we been so seduced into believing that insights into the future are all in the historical data, when the answers are all in the thinking and understanding?"

The first wave of large-scale institutional innovation and efficiency was driven by the rapidly declining cost of computing in the 1980s, which facilitated the development of cheap stock market index derivatives.[51] OTPP was an early user of index derivatives when it captured as much as half of the Canadian index market in 1991–1994 to create synthetic 60% synthetic equity exposure.[52] Because of the declining trend in interest rates, there were a few years where that was seen as a debatable decision, but it eventually proved to be very

---

[50]Leo de Bever built macro-models to help guide monetary policy at the Bank of Canada in his early career. He realized that the econometric tools he was using conditioned his perspective, because they required long-horizon, high-frequency data. Those data were of high quality for sectors that were in relative decline. The most interesting problems for forecasting the future economy were connected to the service sector, for which historical data were of poor quality.

[51]Bond markets lagged, but eventually also improved liquidity and cost-efficiency, at least for institutional investors. Retail investors still in do not have efficient access to bonds in many countries; on-line trading vehicles are closing that gap, but price and cost transparency is still imperfect (see, for example, Ontario Securities Commission, 2014).

[52]When OTPP was created in 1990, assets consisted nearly 100% of unlisted Ontario bonds. Under CIO Bob Bertram, the fund swapped the return on Canadian bonds for cash, and then swapped cash return for equity exposure. Bertram saw the indices as a cheap way to get broad exposure, without any strong views on efficient markets. He felt that indices were just someone else's actively managed portfolio, with its own governance issues. Adding and deleting firms and control blocks can still be a matter of judgment.

profitable.

Using these new tools, even institutions of modest size and some internal investment capacity could cut the annual cost of maintaining their stock-bond asset mix in half to 0.2% or less. The bigger funds started to resemble index funds with an active management long-short overlay. GPFG is a good example of how economies of scale can reduce costs even further.

Real estate was in most cases the first target for diversification into other asset classes, based on a long-term CAPM view that the assets of interest in this category had an attractive long-term return on long-term tail risk somewhere between stocks and bonds. There was no compelling desire to hold a representative slice of "the" market for listed or unlisted real estate because of its vast range of return and risk characteristics did not reflect what these funds were looking for. Real estate portfolios were constructed using direct holdings in high-quality office and retail assets with strong cash flow prospects, located in major markets based on the following considerations:

- The initial entry point was the domestic real estate market, and the percentage of listed instruments required to fill the allocation would have been large. GPFG will be facing the same issue on a global scale. Using round numbers, the Fund manages USD\$800 billion, and the listed real estate market is around USD\$1 Trillion. If GPFG were to build a 10% real estate allocation with listed instruments, it would own 8% of global supply. A holding of that size in small cap instruments is not tradable without considerable market impact, making listed liquidity largely illusory, so accessing the unlisted markets is the only realistic option.

- Management costs for internally managed direct assets were much lower than for external management. Internal management costs are typically 0.5%–1%, depending on the location, the type of asset and its stage of development. An indication of the cost difference of external vs. internal management can be seen in the 3% gap between NCREIF gross and net returns (NCREIF, 2015).

- Direct investments offer better control over asset governance, better alignment between operator and owner, and more control over development strategy.

To a long-term investor, short-term valuation lags are not an important enough governance issue to give up other advantages. The only prices that ultimately matter to their stakeholders are acquisition price and sale price. Everything in between is someone's estimate, of limited relevance without a necessity to transact.

Pension managers worry more about governance defined as alignment of asset manager motivations with their own objective to maximize net return on asset risk. External managers are (quite understandably) motivated to maximize franchise value and fees as a percent

of assets. Over and above the direct asset operating cost, they face incremental cost of marketing, fund reporting, and managing client relations. Standards of governance and business culture in some jurisdictions are not always aligned with expectations set by domestic and other key international regulators.[53]

Early real estate initiatives were typically implemented working with external managers. As experience was gained, emphasis shifted to cheaper, internally managed direct investments, co-investments, or acquisition of captive platforms, e.g., the acquisition of Cadillac Fairview by in 2000, and Oxford Properties by OMERS in 2003. In these "platform" acquisitions, the pension funds gained much better insight into asset operating governance and acquired broader operating expertise, while the real estate managers gained access to stable capital and eliminated their marketing and fundraising costs. More recently, the platform acquisition model seems to be repeating itself in other asset classes.

Cooperation with external organizations can still be mutually beneficial in unfamiliar jurisdictions. It helps to work alongside operating companies with local knowledge. For the external operator, having a reliable partner able to mobilize capital quickly for the right opportunity cuts down on the cost and uncertainty of marketing efforts, provides for more efficiently accessing larger transactions, and diversifying the resulting exposure.

The development and application of firm-wide quantitative risk management technology in the late 1990s motivated other asset mix diversification efforts, e.g., infrastructure timberland, and commodities. The assumption around infrastructure was again that the assets of interest in these categories had a return on risk somewhere between stocks and bonds. Entrance into these asset classes followed a similar investment and cost reduction trajectory.

The direct market was preferred, although a few pension plans invested in listed companies with a view to taking them private, which proved too difficult. Investing through externally managed infrastructure funds had mixed results because of poor governance: the layering of fees and charges caused clashes between managers and investors. Having learned from this experience, subsequent investments were largely direct.

Future sources of innovation in active management are likely to involve superior knowledge management. Again, this runs counter to the notion of efficient markets, and involves uncharted territory, but in a world of accelerating technological change, marginally better insight into its impact on portfolio performance could be very valuable.[54]

While most plans probably benefitted from pushing the boundaries of technology and diversification strategy, Dyck and Pomorski (2011) document that bigger plans with internal

---

[53]The long arm of US legislation and regulation in particular needs to be considered by anyone institution with business interests in the United States, no matter where any perceived weakness in governance is identified.

[54]Van Gelderen, E. and Monk, A., "Knowledge Management in Asset Management," 2015

management capacity and allocations to alternative assets benefitted most (see also Section 4.2):

> The largest plans outperform smaller ones by 43–50 basis points per year. Between a third and one half of these gains arise from cost savings related to internal management, where costs are at least three times lower than under external management. Most of the superior returns come from large plans' increased allocation to alternative investments and realizing greater returns in this asset class. In their private equity and real estate investments large plans have both lower costs and higher gross returns, yielding up to 6% per year improvement in returns. The ability to take advantages of scale depends on plan governance with better governed plans having higher scale economies.

The numbers may understate long-term outperformance, because some of these programs are still facing set-up and development costs, particularly in infrastructure.[55] CPPIB and AIMCo are more reflective of conditions GPFG is likely to confront: a gradual re-allocation from stocks and bonds into other asset classes, set-up costs, and a trade-off between lower short-term performance to earn higher long-term returns. That requires stakeholder patience. As the 2015 CPPIB annual report puts it:

> As the quintessential long-term investor, CPPIB's most important measure of success will always be its long-term performance. It is the Board's duty to ensure management is building a sustainable organization and strategy, and that it has implementation plans in place designed to deliver such results over the coming decades. This exceptionally long horizon presents a unique challenge to Directors and the [senior management team] alike as it requires us to make decisions today, the results of which will not be realized for years to come.

8.4 Real estate in a large institutional portfolio

Pension plan real estate managers never considered investing in a representative slice of the broad real estate market, listed or unlisted because it did not meet their requirements. They

---

[55]Leo de Bever is very familiar with the underlying CEM data base, and specifics for 5 global pension managers. It has its issues, primarily at the most granular data level, where it is often hard to establish accuracy of classification. He found net total return tends be accurate, but data submitted by some institutions have on occasion made it difficult to establish whether gross return includes all costs. CEM has been trying to deal with that issue. Some of this has to do with local cost accounting requirements. However, the broad conclusions drawn by Dyck and Pomorski corroborate his own experience.

are attracted to a subset of real estate and infrastructure opportunities that (1) diversify equity and fixed income exposure, (2) fit on a long-term CAPM return-risk frontier somewhere between stocks and bonds.

Liquidity and high frequency valuation were never top of mind. The expected holding period is 15 years or more. The large Canadian fund portfolios have grown, but their initial core holdings have not changed materially over time. The most significant modifications at some of the Canadian funds were the transition — within a growing portfolio — from joined ownership of a particular asset to exclusive ownership, mostly because one partner wanted more complete control over future development of specific properties.

The assumption has been that in the long run, real estate should earn an unlevered return equivalent to that on 60% stocks and 40% bonds, or about 4%–5% after inflation. The income component of about return is typically a fairly stable nominal 4%–5%.

Recent returns of most asset classes have been far higher than long-term expectations. Return on risk is time-varying. According to the Dimson-Marsh-Staunton Global Returns Data (DMS) data base,[56] the real 1900–2014 return on US equities was 6%, and the equivalent for bonds has been 1.6%. However, since 1982, real stock and bond returns were respectively 8.7% and 6.7%, reflecting in part the decline in interest rates.

The value of real estate depends on prospects for local growth, but it is also connected to interest rates and the cost of leverage, even for investors who prefer to invest without leverage. The decline in government bond interest rates has likely run its course. That should make bond returns fall below their long-term average. Real estate will suffer as well, as future cash flows will be discounted at higher rates. Empirically, that causality chain is far from direct. Capitalization rate spreads with government bonds tend to absorb some of the change in interest rates, particularly if interest rates are changing because of higher growth (Mouchakkaa, 2014).

The academic literature expresses a clear preference for REITs. However, the assets underlying REITs are in principle bought, sold, and managed in the same way as unlisted real estate. So, if managers of equal skill manage identical direct and REIT real estate, long-run unlevered asset returns before investment management fees, incentive fees and other charges should be the same. REITs provide opportunities to invest in real estate for a wide range of small and medium size investors, but do not provide a way to express GPFG's comparative advantage as a long-term investor. A REIT vehicle has marketing and oversight costs. It is managed in part to be competitive with other REITS and gives large investors less control and transparency over operational strategy.

---

[56]First published and described in Dimson, Marsh, and Staunton (2002), and updated in Dimson, Marsh, and Staunton (2013).

The literature also places great emphasis on the understatement of risk derived from stale-dated unlisted valuations. Most risk systems already compensate for the impact of stale-dating on risk by dialing up the assigned risk. Valuations of direct holdings are also updated more frequently than used to be the case. To an investor without a need to transact, the importance of high frequency portfolio valuations is primarily of interest to catch mistakes in the accurate recording of transactions, cash flows, and corporate actions. There is no urgent reason to invest in REITs just to get more current pricing.

The high frequency ups and downs of REIT prices may not be all that indicative of "true value" in any case, if one accepts that there is a strong behavioral and emotional component to short-term market movements, and a response to short-term capital inflows. The REIT price to NAV ratio tends to cycle around median NAV over a period of a few years in any case. And when REITS are part of an index trade, they behave like the stock market because they are part of the dynamics and sentiment around the stock market.

In the short run, there is significant impact from a rapid influx of capital from new allocations to the asset class. Increased allocations are often motivated by a rear-view mirror perspective on future returns. The reference portfolio model is particularly unforgiving of making a mistake in that environment. It asks management to evaluate whether buying more real estate has a higher expected return than a 60-40 stock bond combination. That is the exception.

In most cases, the Board representing the asset owner takes responsibility for the asset allocation, sets the asset class specific benchmark, and instructs the manager to "fill the bucket." Management warnings against mediocre future expected returns are often met with "let us worry about that." However, success has many fathers, and failure is an orphan: low absolute returns in practice still ends up reflecting on the manager.

The MSCI/IPD report on GPFG's real estate for 2014 shows a significant effect on return from currency recognition.[57] Managers located in small currency areas with global portfolios will be subject to transitory short-term noise from this source that can at times swamp relative performance.

8.5   Infrastructure in a large institutional portfolio

Infrastructure, as current large investors define the asset class, consists mostly of unlisted investments with moderate return and risk, intended to be held for the very long run. OECD (2014b) documents that most large institutional investments are unlisted and echoes the way this asset class is approached in this section:

---

[57]MSCI, Norwegian Government Pension Fund Global Real Estate Portfolio Report, 2014

The core infrastructure assets outlined above have the following characteristics: large, long-term assets providing essential services, limited or no competition and high barriers to entry, predictable and steady cash flows with a strong yield component, low volatility, and low correlation to the performance of other asset classes.

A private equity investor is already thinking of an exit, the moment he pays for an asset. In infrastructure, the purchase decision is considered to be the most important one, and hopefully the last, for a very long time. The objective, much like in real estate, is capturing incremental long-term return from investing in businesses with good long-term prospects.

Individual investments tend to be larger than most real estate holdings, and because they are often quasi-monopoly assets, investors put great reliance on fair regulation and enforceability of contracts. There is a greater frequency of "club deals", i.e., transactions involving multiple funds to pool due diligence expertise and diversify exposure to a single asset.

What is commonly called "infrastructure" and included in an infrastructure index is largely defined by what things look like on the outside. For example, investing in an electrical power plant may be attractive or not depending on how revenue is generated, and how the risks are shared. A "merchant" plant supplying electricity during peak periods is typically less attractive and riskier than a facility that will receive revenue on an "availability" of capacity, instead of how much is being utilized.

Listed infrastructure indices do not make that clear distinction. They include good enterprises that would still not be part of an infrastructure portfolio. Investors relax these investment criteria at their peril. Ports can be great infrastructure portfolio assets or pose unreasonable risk, depending on barriers to entry, i.e., the potential for construction of a competing facility nearby.

There is still considerable debate over the role private capital should play in financing social infrastructure, and preferences differ by jurisdiction, often for little reason other than local history. That debate is often confused by the notion that there is no need to offer investors a higher return than the cost of government borrowing. That is true if there is no real risk transfer to the private sector for construction cost overruns and operational risks. Experience suggests that publicly financed greenfield projects having to raise far more capital tend to take far longer to complete.

The first generation of Public Private Partnerships (PPPs), generally involves smaller and highly levered projects with high set-up costs, a small equity slice and therefore little real transfer of risk, and has not attracted much interest from larger funds because. PPPs could be far more efficient if they were structured like ISDA agreements with standard business

clauses, and an appendix outlining the specifics of the project. Unfortunately, there are often too many cooks in the kitchen to make that possible.

Significant institutional allocations to infrastructure are still rare. They are increasing rapidly, but without a corresponding increase in the availability of investable opportunities with the right return and risk, experienced investors believe that valuations have become stretched, applying prudent valuations and risk margins. A few managers are experimenting with proxy "placeholder" portfolios drawn from a listed subset most closely related to desirable unlisted assets.

The market for privately financed infrastructure is likely to expand in the not too distant future, as fiscal constraints collide with the pressing need to refurbish existing infrastructure, or build new to accommodate economic growth. Norges Bank therefore has time to build the required infrastructure expertise, initially working alongside peers, or by acquiring a culturally compatible platform. Relevant infrastructure opportunities and challenges mirroring my own experience are documented in Clark, Monk, Orr, and Scott (2012).

Australia and Canada were the investment pioneers in this asset class, and worked through the teething problems of the early investment vehicles. The dire condition of various Australian states led to privatization of large scale airport, toll road, energy, and communications assets. The asset class provides clear examples of external manager governance problems that are vastly more significant than any stale valuation issues.

Inderst and Della Croce (2013) describe how many of the privatized Australian assets ended up in listed and unlisted infrastructure funds. Canadian funds were some of the early fund buyers. "The model involved buying infrastructure assets that were highly leveraged, complicated with a variety of agency conflicts, and bundled into listed vehicles at high fees." These issues are described in more detail in Lawrence and Stapledon (2008). Disclosed management fees were not the main bone of contention. The real problem was the opaqueness of charges to the funds for "services" that were either unnecessary or overpriced, implying a brazen disrespect for investors.

That experience taught investors a great deal about agency issues associated with divergence between external manager and investor objectives: it is better to control governance directly by being on the inside, than be left wondering on the outside. The market is transitioning away from the early problematic merchant bank sponsored vehicles to structures that offer a clearer line of sight on the underlying assets.

Clear contracts and regulation are essential. The importance of social infrastructure (roads, water and sewer facilities, pipelines, electric transmission networks, airports and ports) means that there has to be a mechanism to maintain pricing fairness from both a customer and an investor perspective. Regulators of various networks (e.g., gas, electric, wa-

ter and sewage) and politicians supervising concession contracts (e.g., toll road concessions) are subject to asymmetrically greater pressures from users. When the choice is between a few (often non-local) investors and a lot of users who can vote or complain, the temptation can be to shade the outcome against the investors, with detrimental effects on the ability to attract new infrastructure capital. Listed assets do not provide immunity from that kind of issue.

The best infrastructure jurisdictions from an investor perspective have carefully built a reputation for regulatory fairness and contract reliability, because they understand that investor risk perception of an infrastructure investment is reflected in the required rate of return: the better the reputation, the lower the implied risk premium. The most difficult geographies are on either side of middle income economies: high income jurisdictions can take the view that they have the luxury to give in to local consumer pressure, and politically unstable low income areas may be unable or unwilling to pay the expected return or enforce concession contracts.

Despite the governance issues in early infrastructure investment vehicles, investors often earned higher than expected returns because of the subsequent influx of new capital. The push to fill allocations has recently meant that disciplined investors have lost by margins as high as 30%. In some cases, part of that gap can be explained by bidder specific synergies with existing holdings. Better opportunities will likely emerge over the next decade for those willing to innovate around the obstacles. A few examples may illustrate the point.

Fiscal challenges in most OECD countries should increase the supply of privately financed infrastructure projects. So far, that has not happened, despite $70 trillion of estimated "need" by 2030 (G20, 2014). In most cases the problem is one of decades of serious under-pricing of the total cost of providing infrastructure. Voters expect governments to deliver social infrastructure services, but often do not understand (or do not want to understand) that current pricing needs to recover capital cost, depreciation, and operating cost. Most cities and states now have accumulated more infrastructure than they can maintain with current revenue structure.

One acute example is water and sewage. Most urban jurisdictions charge developers for the capital cost of building water and sewage networks. They then pay for the operating cost by charging user fees. Attempts to keep costs low have resulted in a rising amount of deferred maintenance, which eventually turns into far higher costs of repair when things break down. Issues like this typically resolve themselves when these problems create a crisis.

Investors in this sector should favor jurisdictions where privatization of infrastructure has been preceded by pricing reform. When a private operator takes over a public network where pricing has not been modified, the need to recover both capital and operating costs

will typically require a dramatic increase in user charges.

As stated above, superior returns depend on continued innovation. Water again is an area where constructive cooperation between institutional investors and political decision makers could lead to profitable investment opportunities in accelerated identification and adoption of better water infrastructure technology, which also has the potential of being far more energy efficient.

For example, the state of "California needs to spend $39 billion to meet its drinking water infrastructure needs over the next 20 years and $29.9 billion in wastewater infrastructure over the next 20 years, according to the American Society of Civil Engineers."[58] On top of that, the State is being forced to reduce water use by 30% because of drought, and rapid aquifer exhaustion.[59] Planned price increases of between 5% and 31% cannot balance demand and supply Looking at the pattern of existing water use, there is ample opportunity to invest in facilities that will improve efficiency of water use and water treatment. Investing in new technology could therefore be both profitable and environmentally desirable.

New energy infrastructure is another area where innovation can shape new types of infrastructure investments. We will see greater efforts to make fossil fuels more capital, water, and energy efficient, and investments in application of new and more integrated approaches to solar, wind, and battery technology. The opportunity for innovation in this space comes in three forms: identifying and commercializing superior technologies, capturing the benefit by deploying that technology, and gauging the implications of these changes (positive and negative) for a fund's broader equity holdings.

The "need" for infrastructure in developing economies is even more acute than in the OECD, but this is an area where extreme caution is advisable. Need and ability to pay to cover return on investment do not always coincide. In practice, it will be difficult to invest in this area without having local partners, as well as supra-nationals providing both debt financing and political risk cover.

The emphasis in this report has been on equity investments in infrastructure. New capital requirements for banks are causing them to abandon certain lines of business, including infrastructure financing. This is not the usual equity infrastructure opportunity, but it could provide an attractive return on (lower) risk for organizations that keep their eye on risk and return and think beyond rigid asset class silos.

---

[58]Ky Trang Ho, Forbes, April 4, 2015, "Promising Investment Opportunities Gushing From California's Epic Water Crisis."

[59]http://www.reuters.com/article/2015/08/18/us-california-water-rates-idUSKCN0QN1PH20150818

8.6 Governance of long-term investment strategies

New real estate and infrastructure investments will have a much longer-term focus. Managing a long-term investment strategy without a clear separation of investment oversight from short-term political scrutiny is extremely difficult. Any unusual investment can be called into question, whether justified from an investment perspective or not. Infrastructure assets in particular have a higher risk of attracting public attention because they tend to be iconic "story investments."[60] An asset management organization at arm's length from policy and politics is better equipped to deal with such issues.

GPFG and CPPIB are similar in that the assets of Canada's CPPIB are owned by the people of Canada, but they have very different governance models. The operational management of GPFG is delegated to Norges Bank, with very detailed direct supervision and high frequency reporting obligations. The legislation defining the governance structure of OTPP and CPPIB provide for clear separation from the political process by setting up a standalone asset management organization with an independent Board, supervising a very high level investment policy, following a "prudent man" model.

The role of the Board of a large asset manager connected to the public sector can be difficult. In the aftermath of 2008, concern about personal director liability has dramatically increased, creating greater reluctance to support promising but unusual investment decisions. Board members are selected in part because they have had successful careers. No Board member wants his or her appointment to become a blemish on that track record. The fear of reputation risk can cause directors to be far more cautious than they have to be to move the organization forward.

Yet, superior performance depends on continued innovation in response to changes in the investment and technological environment, and the emergence of better risk management and other analytical tools. The best (and toughest) Board members use their career experience to make management aware of all the questions to ask, but having received good answers, they support management in its efforts to maximize value.

The governing body of a standalone asset manager needs to have the freedom to establish a high-performance management culture, motivated to compare and contrast the long-term return and risk of all available opportunities that exploit its comparative advantage. Most pension plans are not reaching maximum potential for their clients because they cannot get past the obstacle of having to hire and reward the talent to execute a long-term strategy.

Attempts to implement a high performance investment culture using watered-down ver-

---

[60]For example, private ownership of the Toronto Highway 407 toll road concession became a high profile issue for equity owners when a new government tried to invalidate the contract. It was eventually upheld in court, which was positive for future infrastructure investments in Ontario.

sions of the independent asset management model have failed. Having high ranking civil servants or elected representatives involved in governance makes it harder for both the asset management organization and the asset owners to separate politics and policy from investing.[61] One clear sign of independence is the ability to set a budget, compensation model, and performance measurement structure appropriate for the organization.

8.7 Performance measurement

Ang et al. (2014a) state that "We caution that while the Opportunity Cost Model is conceptually relatively simple, it is one that is challenging to operationalize." The challenges are manageable as long as everyone remembers what the objective is: superior long-term net return on risk compared to the reference portfolio, and that the road to the long-term is full of twists and turns. Unfortunately, that consistency of thinking about goals does not always win the day.

Most pension and endowment funds reward internal management for delivering incremental return over market benchmarks. The reference portfolio model sets an even higher standard by making management accountable for redeploying risk away from a passive stock-bond allocation.

The traditional approach to benchmarking performance is bottom-up: the Board sets a benchmark for each portfolio, and the total fund benchmark is an asset-weighted composite. Under that model, if the decision to deploy capital in real estate and infrastructure was premature, senior management is held harmless as long as they do no worse than the real estate benchmark. Under the reference portfolio model, management is accountable for any deviation from reference portfolio return, in the long run.

How does all of this change with the introduction of more real estate and infrastructure? Listed strategies, including those involving REITs and listed infrastructure continue to fit well within the existing active management model. Such efforts to earn active returns typically have a relatively short-term horizon, i.e., up to 4 years, and total and active risk can be measured with conventional tools. One can set an active risk budget to control security selection risk at the portfolio level, and hold senior management accountable for total risk.

Fitting long-term unlisted strategies into this framework is also conceptually not that difficult: each asset needs to be evaluated in terms of the components of the reference portfolio it displaces from a risk and expected return perspective.

---

[61]Alberta Investment Management Corporation was established in 2008. Political interference was never an issue. However, one of AIMCo's first (and very profitable) decisions involved an equity and debt investment in an Alberta based company with large international operations. The initial public opinion response suggested this investment was made for political reasons, given the presence of the Deputy Finance Minister on the Board. The Finance Minister responded by removing the Deputy, and the issue never arose again.

Senior asset managers should be accountable for selecting investments, but separation of duties requires that someone else in the organization has to set the stock-bond mix it displaces. That can become a very granular exercise, since it should in principle be tied back to a stock-bond combination in the geographical location of the new investment.

The most logical place for evaluating what part of the reference portfolio is being displaced is the risk department. It typically already has a reporting line to the Board, and is used to setting standards for risk measurement and management in cooperation with but independent of senior management. There should be a policy describing the general approach to these decisions, which can be applied to smaller investments, with periodic summary reporting back to the Board. For investment decisions that require Board approval, the investment proposal should include a discussion documenting the risk and return characteristics of the investment.

The critical performance measurement issue, particularly around greenfield infrastructure and real estate, is the transition period between short-term results and long-term expected return.[62] In many cases, assets may be held at cost for the first few years by accounting convention. If a project involves construction, recorded value only increases by value put in place.

If unlisted assets require some period to become fully productive, they should be excluded from short-term performance comparisons and put into a J-curve pool, keeping track of the amount invested at benchmark returns. Performance targets should be qualitative during that period, e.g., using comparisons relative to the original business plan. When the asset matures, compound annual value added since inception of the investment can from then on be measured as compound actual return minus compound return on the stock-bond combination being replaced.

8.8   Management compensation

Few pension and endowment funds have been able to shift to internal asset management. Many see the merits of doing so, but have been unable to get approval for a compensation arrangement that could attract and retain the right talent.

Managers are not solely motivated by compensation, but whatever the arrangements, they have to be seen as fair compared to the way others with comparable skills are compensated. There is optionality to any compensation arrangement: good people will have alternatives. That requires making pragmatic compromises between theoretical purity, mathematical and

---

[62]Some organizations avoid infrastructure and development real estate because they were unable to solve this time horizon problem. That is an unnecessary agency problem.

conceptual rigor, and the need to achieving the principal goal: attracting and retaining staff, and motivating them to make good long-term decisions.

Setting targets and determining rewards often consumes an excessive mount of governance time. The fear is that performance payments, however well structured, may be publicly criticized in an organization close to the public sector, because it is operating in an industry with pay ranges far above the national average.

For organizations that manage a large diversified portfolio predominantly internally, internal investment staff compensation typically makes up about a small part of the total budget. Internal management is up to 75% cheaper, depending on the asset class. Despite this cost advantage, it is far less controversial to pay a much higher amount to an anonymous external manager than a smaller amount to a very visible member of internal staff.

Even in largely internally managed funds, is not unusual for 15% of a portfolio that is externally managed to account for over 60% of the total budget. This to some extent reflects the pragmatic realization that some investment products are attractive, but internalizing the required expertise is not cost effective.

AIMCo is representative of the cost structure outlined above. It has more investments outside of stocks and bonds than the Fund is currently contemplating. So, while its asset weighted unit cost is below peer average, its total cost will be higher than should be the case for GPFG. The AIMCo reward system was calibrated to be "0.5 and 5": less than 0.5% operating cost as a percentage of assets, and a 5% variable compensation for active value added, with the other 95% going to the client.

Compensation systems have to be fairly granular to achieve their goal. Good incentives shape desirable behavior; poorly designed ones shape dysfunctional behavior. Top level management should be mostly rewarded for contributing to overall performance, with the remainder based performance of their asset class. At the portfolio level (e.g., real estate and infrastructure), a lower share tied to total results is more appropriate, as managers respond best to performance standards that are within their span of control.

The opportunity cost benchmark for real estate and infrastructure, i.e., some weighted combination of stock and bond returns in the reference portfolio the new real estate and infrastructure investments is intended to replace, does not work well with lower-level managers, unless results are truly measured over a long horizon.

In the case of real estate, a lower level IPD benchmark may not be ideal, but it is both reasonably effective and acceptable to portfolio managers. For infrastructure, the lower level benchmark at AIMCo was 50% equities and 50% index linked bonds. That is consistent with the reference portfolio model in the long run, but a hard sell for short-term performance measurement in some years. In AIMCo, the Board provided sufficient freedom to

address situations where the objective to set precise long-run targets runs into the reality that numbers and benchmarks don't always do justice to what was achieved on the way to that long run.

Short-term value added is inherently volatile. Even a top quartile manager will have a negative quarter once every four years. That leads to Board agonizing over rewarding potentially spurious performance. That can be resolved through multi-year averaging. Most reward systems work in good times, but the real test is whether they can motivate a management team in bad times.

Leo de Bever helped design one of the first variable compensation systems in Canada around 1997. It had a four-year horizon, largely because of income tax issues that eventually turned out to be incorrect. But the 4 years stuck. There are good reasons to make the averaging period longer, particularly under the circumstances under discussion here, where it may take sometime for superior results to materialize.

There is a view that long-horizon rewards are heavily discounted by those it tries to motivate. While that may be true, it is very useful to be able to reward extraordinary sustained four year-performance with the possibility of earning incremental rewards if that performance is to be sustained for the next four years.

The best constructed performance systems can still generate awkward optics. Negative absolute returns but positive relative returns are a case in point. For that reason, OTPP and AIMCo moderate payouts by an index of cumulative 4-year returns: i.e., payouts are lower in a low absolute return environment. CPPIB faced a somewhat different issue a few years ago: even though performance pay reflect a four-year average, having a negative fourth year created public perception problems.

Most compensation systems designed to attract and retain end up punishing staff for leaving, through forfeit of long-term deferred compensation. That creates a conflict with the goal of making good long-term decisions that will come to fruition after the responsible manager leaves the firm. Private managers have found solutions to that issue by structuring a retained interest for employees that leave the company to recognize value produced by decisions made during time spent at the firm.

8.9   Risk measurement and management

In asset management, the scarce resource is the asset owner's tolerance for downside risk. Assets are just envelopes for risk. All asset management is fundamentally risk management. All investments can in principle be compared on their return on risk. It is not meaningful to compare return on asset categories without factoring in risk.

To long-term asset owners, risk is not annual volatility, but the probability of having to crystalize a long-term capital loss. The experience from large institutional investors is in support of the following:

1. To measure risk on lower probability, more painful outcomes, e.g., 1% Value at Risk or 1% Extreme Tail Loss.

   - Whenever the net distribution of return is non-normal, this is not just a linear transformation on volatility.
   - Mature real estate and infrastructure investments have relatively modest long-term tail risk because of high and stable annual cash returns

2. To measure risk over a longer horizon, because that can significantly change the relative return on risk of different asset classes.

   - Campbell and Viceira (2002) show, for instance, that as one lengthens the investment time horizon, the relative riskiness of stocks declines, while that of bonds increases, a particularly important consideration today, given the likelihood of regime shift in interest rates.
   - Funds like GPFG have the ability to analyze and manage at all the pieces of the long-term return and risk puzzle in one coherent asset management framework. They can go beyond thinking in terms of asset class silos, and explore inefficient niches within asset classes and between traditional silos.

3. To not use tracking error for investments in real estate and infrastructure.

Measuring risk to control downside is important, but the more valuable contribution of good risk measurement is providing a common language to debate the trade-off between total long-term expected return and risk. In that context, to make risk measurement useful in the daily operation of an asset management organization, it has to be rigorous enough to stand up to scrutiny, yet not so complicated that it cannot be easily interpreted and applied.

There is a noticeable difference in how managers of various asset classes perceive and manage risk. Stock and bond managers are used to short-term analytics, and tend to manage their strategies over a 1-4 year horizon. That in itself is already much longer than is common in organizations that cannot rely on stability of capital. Real estate and infrastructure managers are focused on acquiring and managing good long-term assets and are likely to look at short-term risk measures as necessary to fit into the over-all system, but not immediately relevant to how they can affect outcomes.

This is one of those cases where doing things approximately right is better than getting it precisely wrong. For example, the under-measurement of risk because of stale dating in unlisted assets can be fairly easily corrected for by adding an extra risk margin.

Judging by the large variations in tracking error and portfolio standard deviations, the Fund is using a very short time horizon. That is common in organizations with short-term strategies, but leads to counterintuitive results: a 100–200 day risk measure will rise sharply as market values drop. Yet, making new investments is clearly more attractive at the new lower valuations.

For a long-term manager, downside loss should be built up from individual security positions over as long a history as is available (some funds use 25 years of daily data), to capture the worst history has seen, and to avoid making assumptions about the empirical distribution of returns and changing correlations.

Management should be motivated to maximize return on total risk. A given loss suffered because of poor markets is financially no less painful than the same loss from a poor active management decision.

The goal should be to allocate risk to equalize the expected return on risk across all investments. The practical issue is that over any horizon, expected return estimates are by definition uncertain, and relatively small changes in expected return can produce unreasonably large swings in asset allocation when evaluated by mean-variance optimizers.

There are times when the direction of expected returns is more certain than others. After nearly forty years of interest declines, it is difficult to argue that recent high historical bond returns can persist over the next ten years. So investments in any asset class that is sensitive to interest rates (e.g., bonds, real estate, and infrastructure) currently are at greater than normal risk. Equities may suffer as well, but if the rise in interest rates occurs in reasonably strong growth and profit margin environment, it should do well in a relative sense.

In an ideal world, total portfolio risk should be allocated to equalize the incremental return on risk across all investment categories and strategies, i.e., shifting a bit of risk from one place to another should not change total return on risk. That is theoretically superior but becomes very difficult to motivate for managers, since it makes the risk contribution of one asset class a function of every other asset class.

8.9.1   Tracking error as a resource for active managers to earn return

The appropriate risk control at the Fund level is total risk, measured over a long historical period, Ex ante tracking error, again estimated over a long history, still serves a purpose as the first line of defense in an active risk budgeting framework at the portfolio level. In an active risk budgeting framework I first deployed at OTPP in 1997 and later at AIMCo, tracking error or tail risk is the resource available to active managers to earn active return. The goal is never to take risk, the objective is to deploy risk when it has an attractive return.

The total active risk budget is allocated annually to all the active programs based on

an assessment of available opportunities, and an assumption that manager return on risk or information ratio was 0.25, i.e., between median (information ratio of 0) and top quartile (information ratio of 0.5). Empirically, a collection of active managers with an average information ratio of around 0.25 will collectively perform at a higher information ratio because of diversification across active strategies.

The diversification of active and passive risk depends very much on the nature of the active positions. Simply overweighting equities or results in active and passive risk being essentially additive. Negative correlation between the two sources of risk can result from positions that overweight securities with a higher expected return on risk than the index, and underweight those for which the opposite is true. In unlisted asset classes, interaction of active risk and systemic risk is part art and part science, but reasonable proxies served their purpose.

AIMCo used a green-yellow-red system to monitor risk usage: managers had full discretion if they were within their risk limits (green), had to get their manager's permission if they were above that limit (yellow), and were stopped out when they reached 125% of their allotment.

The preferred use of total risk as a Fund risk measure, the limitations of tracking error, and the use of risk measurement in the debate over return on total risk can be illustrated by the conundrum faced by Ontario Teachers' Pension Plan's senior management at the end of 1999. Equity valuations had become stretched by any measure, so management felt that the return on equity risk going forward was likely to be small or negative. They wanted to cut equity exposure significantly. However, they had set a tracking error limit with the Board. Even though the proposed move reduced total risk significantly, it created far more tracking error than was provided for in the active risk budget. Tracking error only signals potential differences from passive return. It does not say anything about expected return on active risk, or impact of active risk on return on total risk .

The Board agreed that total risk was the better control on risk and let management proceed as long as it did not violate the plan's total risk constraint, and despite some doubts, reflecting the euphoria of the time, i.e., "there is something different about the 'New Economy'." There were some nervous moments in the middle of 2000, but this framework for managing to a total risk constraint turned out to be both effective and profitable for OTPP's clients in the next few years.

8.9.2 Managing total risk for maximum total return

As mentioned earlier, the reported standard deviation and tracking error for the Fund appears to be very short-term. Assuming long-term passive reference portfolio risk is 9%, a

total risk budget of passive risk plus 1%, or 10% should be reasonable compared to similar funds. It will motivate management to manage its total risk budget for the most effective way to earn active return. Total risk is calculated as:

$$\left[(\text{Active risk})^2 + (\text{Pass. risk})^2 + 2 \times (\text{Active-Pass. Correl.}) \times (\text{Active risk}) \times (\text{Pass. risk})\right]^{\frac{1}{2}}.$$

Given the active-passive correlation and a 9% passive risk, one can solve for the active risk room. Assuming a 0.5% top quartile information ratio, that provides achievable value added targets in the table below:

Active-Passive Risk Correlation	Implied Active Risk Budget	Top quartile value-added potential
1.0	1.0%	0.5%
0.5	1.7%	0.9%
0.0	4.4%	2.2%

GPFG's recorded diversification between active and passive risk implies a correlation of active and passive return higher than 0.5.

8.10 Setting value-added targets

Having built the required internal management capacity and links to external product platforms, GPFG should over time be able to set a long-term stretch goal of adding net active returns of 1% per year or more over the reference portfolio. That objective should fit within the 1% incremental total risk limit over reference portfolio risk outlined above.

Annual value added will be volatile. Even a top quartile manager should expect a negative active return about 1 year in 4. The experiences of the three large Canadian funds, which are indicative of both successes and challenges, are detailed in the following. Other relevant experience can be drawn from funds in the UK and the Netherlands.

Ontario Teachers' Pension Plan was founded in 1990. It adopted a formal risk budgeting strategy and value added targets around 1997. Initial value added targets were a relatively modest 0.5%. When it was suggested to aim higher around 2000, there were fears that increasing the targets over time would be great for clients but hard on management performance measures. That fear was largely unfounded. Despite a challenging 2008, active return over market benchmarks has averaged 1.8%.

If OTPP had used the reference portfolio model, that outperformance would have been somewhat lower, because of the use of absolute return benchmarks for some asset classes

with returns appropriate in the long run, but less challenging over the recent measurement period.

CPPIB value added relative to the reference portfolio model established in 2006 has been about 0.4% per year. The buoyant recent performance of bond and stock markets set a very tough hurdle. The large influx of capital and the rapid expansion of allocations to unlisted portfolios did create serious headwinds due to accounting for set-up costs, J-curve effects, and conservative accounting conventions for unlisted assets. These kinds of effects should taper off as these portfolios mature.

AIMCo set its first modest top quartile active return "stretch" targets in 2009, in part reflecting the Board's view that one had to learn to crawl before trying to run. Portfolio benchmarks were similar to the reference portfolio approach. Value added targets were gradually increased from 2009–2014 and realized value added has averaged about 0.5%–0.7%. Some legacy issues, and the rapid expansion of allocation to asset classes not included in the reference portfolio posed some of the same set-up challenges encountered at CPPIB.

Long-term stock and bond returns suggest that the reference portfolio will have an expected return after inflation of around 4%. Setting an aspirational target of earning an extra 1% is not easy or without risk, but a 25% faster real rate of accumulation represents a material improvement from an asset-owner perspective. We are not held accountable for the opportunity cost of not trying, although perhaps we should be. To paraphrase John F. Kennedy, we should not set ourselves targets because they are easy, but because they are hard. Even Keynes disregarded his own observation about the perils of trying to succeed unconventionally.

Conclusion

In summary, the findings from the academic literature on real estate and infrastructure are fundamentally at variance with the experience reflected in the observed investment strategies of GPFG's peers: they are almost exclusively focused on unlisted assets. One of us (de Bever) believes GPFG should create the internal capacity to emulate these strategies since they reflect comparative advantages, opportunities, and governance concerns specific to large funds:

1. Being able to deploy stable capital long term provides opportunities to create value in relatively inefficient markets beyond trading in a zero sum game listed markets
2. The main governance issue is seen to be lack of alignment with external manager objectives. Direct investments provide better control over asset governance. Stale valuations are a minor annoyance for very long term investors.

3. Any liquidity advantage of listed assets is limited, given the large size of the holding for GPFG peer institutions.

4. No peer fund seems to target a proportionate slice of either the real estate or the infrastructure market, listed or unlisted, as discussed in the academic literature. All seem to concentrate on large, high quality assets in real estate and stable return infrastructure assets.

5. Internal management is far more cost-effective than going through listed vehicles.

References

Almond, N., 2015, Money into property 2015, Technical report, DTZ Research.

Amihud, Y., and H. Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 33, 263–291.

Amihud, Y., H. Mendelson, and L. H. Pedersen, 2005, Liquidity and asset prices, *Foundations and Trends in Finance* 1, 269–364.

Andonov, A., R. Bauer, and M. Cremers, 2012, Can large pension funds beat the market? Asset allocation, market timing, security selection, and the limits of liquidity, Working paper, Maastricht University.

Andonov, A., P. Eichholtz, and N. Kok, 2015, Intermediated investment management in private markets: Evidence from pension fund investments in real estate, *Journal of Financial Markets* 22, 73–103.

Andonov, A., N. Kok, and P. Eichholtz, 2013, A global perspective on pension fund investments in real estate, *Journal of Portfolio Management* 39, 32–42.

Ang, A., M. W. Brandt, and D. F. Denison, 2014a, Review of the active management of the Norwegian Government Pension Fund Global, Technical report, Columbia University.

Ang, A., B. Chen, W. N. Goetzmann, and L. Phalippou, 2014b, Estimating private equity returns from limited partner cash flows, Working paper, Columbia University.

Ang, A., N. Nabar, and S. J. Wald, 2013, Searching for a common factor in public and private real estate returns, *Journal of Portfolio Management* 39, 120–133.

Angel, S., 2012, *Planet of Cities* (Lincoln Institute of Land Policy, Cambridge, MA).

Bauer, R. M. M. J., K. J. M. Cremers, and R. G. P. Frehen, 2010, Pension fund performance and costs: Small is beautiful, Working paper, Maastricht University.

Berk, J. B., and R. C. Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, 1269–1295.

Berk, J. B., and R. Stanton, 2007, Managerial ability, compensation, and the closed-end fund discount, *Journal of Finance* 62, 529–556.

Bhattacharyay, B. N., 2012, Estimating demand for infrastructure, 2010–2020, in B. N. Bhattacharyay, M. Kawai, and R. M. Nag, eds., *Infrastructure for Asian Connectivity*, chapter 2, 19–79 (Edward Elgar).

Black, F., 1989, Universal hedging: Optimizing currency risk and reward in international equity portfolios, *Financial Analysts Journal* 45, 16–22.

Black, F., 1990, Equilibrium exchange rate hedging, *Journal of Finance* 45, 28–30.

Black, F., and R. Litterman, 1992, Global portfolio optimization, *Financial Analysts Journal* 48, 28–43.

Bond, S. A., and Q. Chang, 2013, REITs and the private real estate market, in H. K. Baker, and G. Filbeck, eds., *Alternative Investments: Instruments, Performance, Benchmarks, and Strategies*, chapter 5, 79–97 (John Wiley).

Boudry, W. I., N. E. Coulson, J. G. Kallberg, and C. H. Liu, 2012, On the hybrid nature of REITs, *Journal of Real Estate Finance and Economics* 44, 230–249.

Brown, K. C., L. Garlappi, and C. Tiu, 2010, Asset allocation and portfolio performance: Evidence from university endowment funds, *Journal of Financial Markets* 13, 268–294.

Campbell, J. Y., C. Pflueger, and L. M. Viceira, 2014, Monetary policy drivers of bond and equity risks, Working paper, Harvard University.

Campbell, J. Y., and R. J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

Campbell, J. Y., and L. M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press).

Case, B., 2015, Does private real estate actually have a low correlation with public real estate?, `http://allaboutalpha.com/blog/2015/03/15/does-private-real-estate-actually-have-a-low-correlation-with-public-real-estate/`.

Chambers, D., and E. Dimson, 2013, Retrospectives: John Maynard Keynes, investment innovator, *Journal of Economic Perspectives* 27, 213–228.

Chen, J. S., H. G. Hong, M. Huang, and J. D. Kubik, 2004, Does fund size erode mutual fund performance? The role of liquidity and organization, *American Economic Review* 95, 1276–1302.

Chen, L., 2009, On the reversal of return and dividend growth predictability: A tale of two periods, *Journal of Financial Economics* 92, 128–151.

Cherkes, M., J. Sagi, and R. Stanton, 2009, A liquidity-based theory of closed-end funds, *Review of Financial Studies* 22, 257–297.

Clacy-Jones, M., and B. Teuben, 2014, Real estate investment market size: A better understanding of the professionally managed real estate investment universe, Technical report, IPD/MSCI.

Clark, G. L., A. H. B. Monk, R. Orr, and W. Scott, 2012, The new era of infrastructure investing, *Pensions* 17, 103–111.

Clark, G. L., and R. Urwin, 2008, Best-practice pension fund governance, *Journal of Asset Management* 9, 2–21.

CPP Investment Board, 2015, 2015 Annual Report.

Cremers, M., and C. Lizieri, 2015, How active is your real estate fund manager?, *Journal of Alternative Investments* 18, 22–36.

Da Rin, M., and L. Phalippou, 2014, There is something special about large investors: Evidence from a survey of private equity limited partners, Working paper, Tilburg University.

David, A., and P. Veronesi, 2013, What ties return volatilities to fundamentals and price valuations?, *Journal of Political Economy* 121, 682–746.

De Francesco, A., S. Doole, P. Hobbs, B. McElreath, and A. Sharma, 2015, Sectoral aspects of global infrastructure investment, Technical report, MSCI.

Dimson, E., P. Marsh, and M. Staunton, 2002, *Triumph of the Optimists: 101 Years of Global Investment Returns* (Princeton University Press, Princeton, NJ).

Dimson, E., P. Marsh, and M. Staunton, 2013, Global investment returns sourcebook 2013, Technical report, Credit Suisse/London Business School.

Doeswijk, R., T. Lam, and L. Swinkels, 2014, The global multi-asset market portfolio, 1959–2012, *Financial Analysts Journal* 70, 26–41.

Dyck, I. J. A., and L. Pomorski, 2011, Is bigger better? Size and performance in pension plan management, Working paper, University of Toronto.

151

Eggertsson, G., and N. Mehrotra, 2014, A model of secular stagnation, Working Paper 20574, NBER.

Eichholtz, P., N. Kok, and J. Quigley, 2010, Doing well by doing good? Green office buildings, *American Economic Review* 100, 2492–2509.

Eicholtz, P., N. Kok, and E. Yonder, 2012, Portfolio greenness and the financial performance of REITs, *Journal of International Money and Finance* 31, 1911–1929.

Engle, R. F., and C. W. J. Granger, 1987, Cointegration and error correction: Representation, estimation and testing, *Econometrica* 55, 251–276.

Fama, E. F., and K. R. French, 2012, Size, value, and momentum in international stock returns, *Journal of Financial Economics* 105, 457–472.

Fisher, L. M., and D. J. Hartzell, 2013, Real estate private equity performance: A new look, Working paper, University of North Carolina at Chapel Hill.

Fuerst, F., 2015, The financial rewards of sustainability: A global performance study of real estate investment trusts, Working paper, Cambridge University.

Fuerst, F., and P. McAllister, 2011, Green noise or green value? Measuring the effects of environmental certification on office values, *Real Estate Economics* 39, 45–69.

Funk, D., D. Weill, and D. Hodes, 2014, 2014 institutional real estate allocations monitor, Technical report, Cornell University/Hodes Weill & Associates.

Funke, C., T. Gebken, G. Michel, and L. Johanning, 2010, Real estate risk in equity returns, Working paper, WHU.

G20, 2014, Increasing investment in infrastructure, Policy note.

Geltner, D., 1991, Smoothing in appraisal-based returns, *Journal of Real Estate Finance and Economics* 4, 327–345.

Geltner, D., 1993, Estimating market values from appraisal values without assuming an efficient market, *Journal of Real Estate Research* 8, 352–345.

Geltner, D., and B. Kluger, 1998, REIT-based pure-play portfolios: The case of property types, *Real Estate Economics* 26, 581–612.

Giliberto, S., 1990, Equity real estate investment trusts and real estate returns, *Journal of Real Estate Research* 5, 259–263.

Gyourko, J., and D. Keim, 1992, What does the stock market tell us about real estate returns?, *Real Estate Economics* 20, 457–486.

Hasseltoft, H., 2009, The "Fed model" and the changing correlation of stock and bond returns: An equilibrium approach, Working paper, University of Zurich.

Hewitt EnnisKnupp, 2014, Global invested capital market, `http://www.hekblog.com/wp-content/uploads/2014/06/2014.06.30-Global-Invested-Capital-Market-FINAL-US.pdf`.

Hoesli, M., and E. Oikarinen, 2012, Are REITs real estate? Evidence from international sector level data, *Journal of International Money and Finance* 31, 1823–1850.

Hoesli, M., and E. Oikarinen, 2013, Are public and private asset returns and risks the same? Evidence from real estate data, Working paper, Swiss Finance Institute.

Inderst, G., 2013, Private infrastructure finance and investment in Europe, Working Paper 2013/02, EIB.

Inderst, G., and R. Della Croce, 2013, Pension fund investment in infrastructure: A comparison between Australia and Canada, OECD Working Papers on Finance, Insurance and Private Pensions 32, OECD Publishing.

Inderst, G., and F. Stewart, 2014, Institutional investment in infrastructure in emerging markets and developing economies, Technical report, Public-Private Infrastructure Advisory Facility, World Bank Group.

IPD/MSCI, 2014, Global infrastructure investment: An overview of the institutional landscape.

Jaffee, D., R. Stanton, and N. Wallace, 2013, Energy factors, leasing structure and the market price of office buildings in the U.S., Working paper, U. C. Berkeley.

Johansen, S., 1988, Statistical analysis of cointegration vectors, *Journal of Economic Dynamics and Control* 12, 231–255.

Keynes, J. M., 1936, *The General Theory of Employment, Interest and Money* (Macmillan, London).

Koijen, R., H. Lustig, and S. van Nieuwerburgh, 2015, The cross-section and time-series of stock and bond returns, Working paper 15688, NBER.

Koijen, R., and S. van Nieuwerburgh, 2011, Predictability of stock returns and cash flows, *Annual Review of Financial Economics* 3, 467–491.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran, 2015, The tail that wags the economy: Belief-driven business cycles and persistent stagnation, Working paper, NYU Stern.

Kutlu, V., 2010, *The Long Term Relation Between Indirect and Direct Real Estate*, Msc thesis, Tilburg University.

Lawrence, M., and G. P. Stapledon, 2008, Infrastructure funds: Creative use of corporate structure and law — but in whose interests?, Working paper, University of Melbourne.

Lerner, J., A. Schoar, and J. Wang, 2008, Secrets of the academy: The drivers of university endowment success, *Journal of Economic Perspectives* 22, 207–222.

Lettau, M., and S. Van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607–1652.

Li, J., R. M. Mooradian, and S. X. Yang, 2009, The information content of the NCREIF index, *Journal of Real Estate Research* 31, 93–116.

Ling, D., and A. Naranjo, 2015, Returns and information transmission dynamics in public and private real estate markets, *Real Estate Economics* 43, 163–208.

Lintner, J., 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 47, 13–37.

McKinsey, 2010, Farewell to cheap capital? The implications of long-term shifts in global investment and saving, Technical report, McKinsey Global Institute.

Merton, R. C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.

Mouchakkaa, P., 2014, Frozen on the rates: Impact of interest rates on capitalization rates, Technical report, Morgan Stanley.

MSCI, 2014, MSCI core real estate indexes methodology, index construction and maintenance methodology for the MSCI core real estate indexes.

Myer, F. C. N., and J. R. Webb, 1993, Return properties of equity REITs, common stocks, and commercial real estate: A comparison, *Journal of Real Estate Research* 8, 87–106.

NCREIF, 2015, NCREIF Fund Index closed end value add (NFI-NEVA), second quarter 2015, `https://www.ncreif.org/documents/ceva2015q2.pdf`.

OECD, 2014a, Annual survey of large pension funds and public pension reserve funds: Report on pension funds' long-term investments.

OECD, 2014b, Pooling of institutional investors capital — selected case studies in unlisted equity infrastructure.

Oikarinen, E., M. Hoesli, and C. Serrano, 2011, The long-run dynamics between direct and securitized real estate, *Journal of Real Estate Research* 33, 73–103.

Ontario Securities Commission, 2014, The Canadian fixed income market, `http://www.osc.gov.on.ca/documents/en/Securities-Category2/20150423-fixed-income-report-2014.pdf`.

Pagliari, J. L., Jr., K. A. Scherer, and R. T. Monopoli, 2005, Public versus private real estate equities: A more refined, long-term comparison, *Real Estate Economics* 33, 147–187.

Pástor, Ľ., and R. F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.

Pedersen, N., F. He, A. Tiwari, and A. Hoffmann, 2012, Modeling the risk characteristics of real estate investments, Technical report, PIMCO.

Preqin, 2015a, Preqin investor outlook: Alternative assets, H2 2015.

Preqin, 2015b, The Q2 2015 Preqin quarterly update: Real estate.

RARE, 2013, RARE guide to listed vs. unlisted infrastructure.

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

Roll, R., 1977, A critique of the asset pricing theory's tests: Part I: On past and potential testability of the theory, *Journal of Financial Economics* 4, 129–176.

Ross, S. A., and R. C. Zisler, 1991, Risk and return in real estate, *Journal of Real Estate Finance and Economics* 4, 175–190.

Sharpe, W. F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425–442.

Shepard, P., Y. Liu, and Y. Dai, 2014, The Barra private real estate model (PRE2), Research notes, MSCI.

Song, D., 2014, Bond market exposures to macroeconomic and monetary policy risks, Working paper, University of Pennsylvania.

Stanton, R., and N. Wallace, 2015, A new dynamic house-price index for mortgage valuation and stress testing, Working paper, U.C. Berkeley.

Stefek, D., and R. Suryanarayanan, 2012, Private and public real estate: What is the link?, *Journal of Alternative Investments* 14, 66–75.

Tsai, J. P., 2007, *A Successive Effort on Performance Comparison between Public and Private Real Estate Equity Investment*, Master's thesis, MIT Center for Real Estate.

WEF, 2014, Infrastructure investment policy blueprint, World Economic Forum.

WEF, 2015, Strategic infrastructure: Mitigation of political and regulatory risk in infrastructure projects, World Economic Forum.

Yepes, T., 2008, Investment needs for infrastructure in developing countries 2008–15, Technical report, The World Bank.

Yunus, N., J. A. Hansz, and P. J. Kennedy, 2012, Dynamic interactions between private and public real estate markets: Some international evidence, *Journal of Real Estate Finance and Economics* 45, 1021–1040.

A   Mandate

The group shall prepare a public report and give a presentation of the report. The theme for the report is how investments in real estate and infrastructure may improve the trade-off between risk and return of GPFG and how the Ministry should regulate such investments in the mandate to Norges Bank. The report should address the following issues:

1. **Should the Ministry open up for private infrastructure investments and increase the share of the Fund's capital that may be invested in real estate?**
   The group's assessments and recommendations should be based on a review of how increased investments in private real estate and infrastructure may improve the trade-off between risk and return of the Fund, and include a discussion of:

   - Value added from diversification, factor exposure, time variation in risk premia and security selection relative to investments in public equity and fixed income.
   - Risk exposure, including exposure to inflation and economic growth, on short and long horizons relative to investments in public equity and fixed income
   - GPFG's comparative advantages or disadvantages in this area

2. **How should the Ministry regulate real estate and possible infrastructure investments in the management mandate to Norges Bank and how should Norges Bank's performance be reviewed?**
   The group's assessments should be based on a review of the opportunity cost model as recommended by Ang, Brandt, and Denison (2014a) compared to the current model of delegation, and give recommendations on:

   - Benchmarking, including a discussion of the strength and weaknesses with the use of a public benchmark versus indices of public and private real estate and infrastructure assets.
   - Risk regulation, including a discussion of how absolute and relative risk should be measured, managed and regulated.
   - Reporting requirements, cf. Ang, Brandt, and Denison's (2014a) recommendation of a more detailed reporting.
   - Performance evaluation, including a discussion of how the performance should be adjusted for risk and what expectations the Ministry should have to excess return.

To the extent that the Ministry decide to open up for infrastructure investments, the mandate will also include a discussion of opening up for unlisted infrastructure for renewable energy and infrastructure in emerging markets. These investments should be subject to the same financial requirements as the other investments of the GPFG. The report should describe these markets, including size and share of the total market for infrastructure as well as discuss in general the trade-off between risk and return for such investments.

The group's assessments should be based on acknowledged finance theory and empirical work and similar funds' management, results and experience from investments in real estate and infrastructure. The assessments should further be based on a requirement that the Fund's absolute market risk level should be approximately as today.

B   Detailed results on U.S. REITs

This appendix analyzes U.S. REIT data, in parallel with the analysis for global real estate stock in the main text.

B.1   Background: REIT rules and history

In the United States, a REIT is a company that owns, and in most cases operates, income-producing real estate. In contrast, mortgage REITs finance real estate and own mortgages and mortgage-backed securities. To be a REIT, a company must distribute at least 90 percent of its taxable income to shareholders annually in the form of dividends. Because taxable income subtracts out depreciation, there is typically substantial cash flow left over after dividend distribution. The key advantage of this organizational form, in addition to access to capital markets, is that dividends paid by REITs are deductible from corporate taxes, so that REITs typically pay no or very little tax. Other requirements to qualify as a REIT are (2) diffuse ownership (at least 100 shareholders, not permitted for five or fewer individuals (and certain trusts) to together own more than 50% of a REIT's stock), (3) asset concentration in real estate (at least 75% of the REIT's total assets must consist of real estate, mortgages, cash, or federal government securities, and at least 75% of the REIT's yearly gross income must be derived directly or indirectly from real property including mortgages, partnerships and other REITs), (4) buy-and-hold investment (REITs must derive income from primarily passive sources like rents and mortgage interest, as distinct from short-term trading or sale of property assets. REITs are subject to penalty tax of 100% on net income from "prohibited transactions" (e.g., sale of property held primarily for sale rather than long term investment). Non-prohibited transaction arise if: (i) property held for more than 4 years, and (ii) the aggregate adjusted basis of the property sold is less than 10% of the aggregate basis of all REIT assets.

The REIT business has been around for a long time in the U.S. Real estate investment trusts originated in the 1880s at a time when investors could avoid double taxation, or a tax at corporate and individual level. In the 1930s, this tax benefit was removed, causing investors to pay "double tax." In an effort to stimulate more public ownership of real estate assets, President Eisenhower signed the REIT tax provision contained in the Cigar Tax Excise Tax Extension in 1960. It allowed the tax-free pass-through of income from real estate, but only passive portfolio management of real estate properties was permitted. In the 1970s, mortgage REITs dominated the REIT space. The Tax Reform Act of 1986 allows active property management and encouraged a shift away from mortgage REITs towards property or equity REITs. The Revenue Reconciliation Act of 1993 modified the "five or

fewer" rule to make it easier for institutional investors to invest in REITs. Many analysts consider this the start of the modern REIT era. Hence, much of our quantitative analysis will start in January 1994. The REIT Modernization Act of 2001 made some further changes to the requirements of REITs, including a reduction in the minimum income payout ratio from 95% to 90%.

B.2 Market size

As of April 30, 2015, there were 221 publicly traded REITs with a combined market capitalization of $926 billion. Of that, $855 billion was equity REITs and the remaining $74.1 billion mortgage REITs. Of these, 192 REITs trade on the NYSE; their combined market capitalization is $872 billion. Average daily trading volume in April 2015 was $6.2 billion, up from $3.9 billion in April 2010, and $1.4 billion in April 2005. Total equity REIT market capitalization, plotted against the left axis on Figure 31, grew from $1bn in 1982 to $10bn in 1991. Growth accelerated in the early-to-mid 1990s and market capitalization exceeded $100bn by the middle of 1996. REITs continued their steady growth to reach a pre-crisis market capitalization peak of $400bn at the end of 2006. By the end of 2008, equity REITs had lost more than half their market capitalization and ended the year at $175 billion. The number of listed equity REITs fell from 152 at the end of 2005 to 113 at the end of 2008. Since then, market capitalization has rebounded strongly, as has the number of listed REITs. The pre-crisis peak in market capitalization was exceeded by the end of 2011, and the number of listed REITs surpassed the pre-crisis number by 2013.

U.S. equity REITs typically own and manage high-quality properties in and around major metropolitan areas. REITs own approximately $1.7 trillion of commercial real estate assets; this number includes public but non-listed REITs. Based on balance sheet data from 2014.Q4, equity REITs have a debt-to-asset ratio of 31%. Listed REITs paid out $42 billion in dividends in 2014; non-listed REITs paid an additional $4 billion.

B.3 Mortgage REITs

Mortgage REITs, plotted against the right axis on Figure 31, are an order of magnitude smaller than equity REITs. The combined market capitalization of mortgage REITs grew quickly over the past 6 years from $14 billion in 2008 to $61 billion in 2014. It too suffered a 50% drop in market capitalization between 2006 and 2008. The number of mortgage REITs fell from 38 in 2006 to 20 in 2008 as many non-agency REITs went out of business, mostly driven by large defaults on non-agency residential and commercial mortgages during the crisis. An earlier crisis occurred in the late 1990s, when mortgage REIT market capitalization
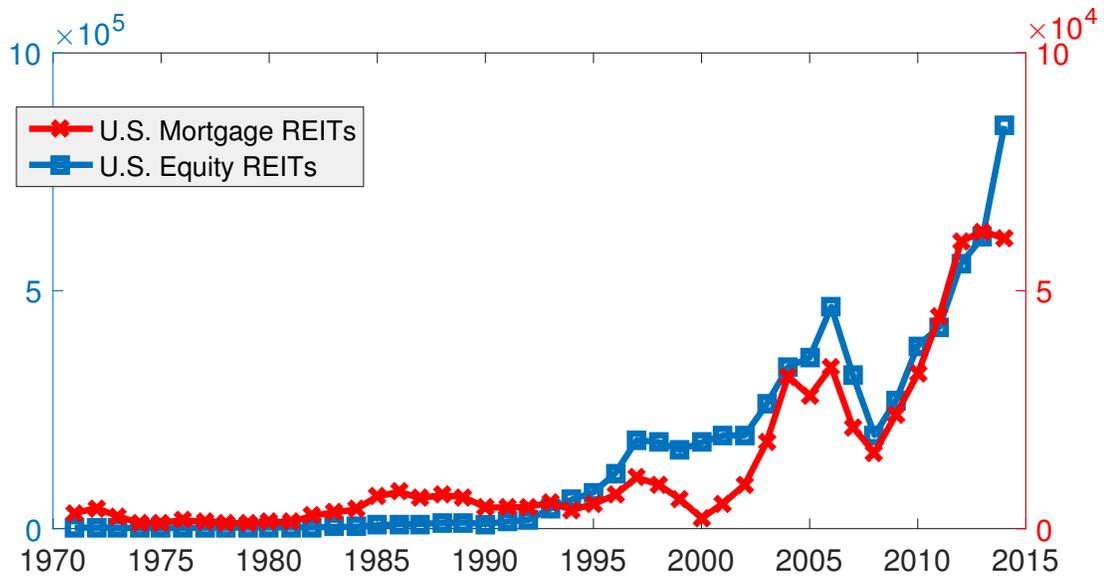
Figure 31: Market Capitalization of U.S. REITs

Aggregate REIT market capitalization, expressed in millions of December 2014 real dollars using the Consumption Price Index of the BLS. REIT data are from NAREIT for 1971–2014. Equity REITs are plotted against the left axis, while mortgage REITs are plotted against the right axis.

fell from $10.5bn at the end of 1997 to $2.2bn at the end 2000. An earlier crisis in 1993-94 was caused by a rise in interest rates. Similarly, between the middle of May and the middle of August 2013, mortgage REITs lost 25% of their market capitalization during the "taper tantrum."

Mortgage REITs come in two flavors. Currently, agency mortgage REITs dominate the market. They invest in long-term mortgage-backed securities guaranteed by Fannie Mae, Freddie Mac, and Ginnie Mae. These securities have no credit risk. They fund themselves with short-term debt, typically 90-day repo. They mostly make money off the net interest margin, the spread between the interest on the long-term mortgage assets and the short-term repo liabilities. It follows that they are subject to interest rate risk, akin to the duration risk that a bank engaging in maturity transformation takes on. They hedge interest rate risk with swaps and sometimes swaptions, but it is unclear just how successfully they are hedged against changes in rates, changes in the slope of the term structure, and changes in interest rate volatility. The taper tantrum episode suggests that they are not as well hedged as they may make believe. Agency mortgage REITs used to have asset-equity leverage ratios of 10/12-to-1. Over the past five years, they have reduced these ratios to 4/7-to-1. However, (1) these ratios are still much higher than for equity REITs, and (2) over the same period they have increased the duration gap between their assets and liabilities so that their overall portfolio risk may not have fallen much.

161

The second type of mortgage REIT is the non-agency REIT, which invests in non-guaranteed residential mortgage-backed securities or commercial mortgage-backed securities, or directly in whole mortgage loans. This business was an important component of the mortgage REIT landscape between 2003 and 2007, but collapsed spectacularly during the financial crisis. Given the tight underwriting standards that ensued after the crisis, the last couple of years have seen the renaissance of the non-agency REIT space. While still small, this group of REITs is growing both because new non-agency REITs come to the market and because existing agency or hybrid mortgage REITs start to shift their portfolios towards non-agency paper.

B.3.1 Return analysis

Data   We use the NAREIT All Equity REIT index return, the industry benchmark. The NAREIT time series start in January 1972. We end our sample in December 2014. We define the modern REIT era as the period starting in January 1994, following the passage of changes to REIT rules in 1993. The 1994–2014 sample will also prove to be useful for comparison with global real estate, for which this will be the longest available sample. NAREIT also reports return indices by property type (Office, Industrial, Retail, Apartments, Hotels, Health Care, Diversified, and Self-storage), which start in January 1994. We shall use those later in the report.

We also use the equity market index return from CRSP, the value-weighted return on all publicly listed stocks in the U.S.; the size, value and momentum return factors (SMB, HML, and MOM, respectively) from Ken French's web site; the return on a constant-maturity 10-year U.S. Treasury; and the traded Pastor-Stambaugh illiquidity factor (Pástor and Stambaugh, 2003) (from Lubos Pastor's web site). The risk-free rate is the one-month T-bill rate from Ibbotson.[63]

Returns: means and volatilities   Table 9 reports annualized means and standard deviations of monthly returns for the full sample (516 months) and for the modern REIT sample (252 months).

Full sample   Panel A shows that average returns on equity REITs were 13.1% per annum over the full sample, with a volatility of 17.1%. The excess return over the one month T-bill, which averaged 4.9% over the full sample, was 8.2% (U.S. real estate risk premium). The Sharpe ratio for equity REITs was 0.48. Mortgage REITs returned 7.2%

---

[63]We are aware of the non-traded nature of the Fama-French factors. A follow-up analysis could use a tradable small value index, for example from Russell, instead.

per annum with a volatility of 20.2% and a Sharpe ratio of 0.11. For comparison, the stock market generated returns of 11.3% and excess returns (over the one month T-bill) of 6.4% (equity risk premium). The volatility of stock returns was 15.8% per annum and the Sharpe ratio 0.41. Finally, ten-year Treasuries returned 7.7% over the same period, a premium of 2.8% over T-bills, with a volatility of 7.8% and a Sharpe ratio of 0.36. The table also reports return skewness which is a measure of how asymmetric returns are. Both equity REITs (−0.72) and stocks (−0.54) have negative skewness, implying that large negative returns are more frequent than large positive returns. Mortgage REITs (−0.26) have a small negative skewness and bonds (0.36) have positive skewness. Based on these numbers, U.S. equity REITs have outperformed U.S. stocks and bonds, while U.S. mortgage REITs under-performed both. We return to finer performance analysis below.

Modern REIT period Panel B shows results for the 1994–2014 period. The performance of equity REITs in this period was stronger than in the first half of the sample. Average returns were 12.6% per year, or 9.9% above the risk-free T-bill rate. However, the volatility was also higher at 19.8%, arguably because the financial crisis is a more influential observation in the shorter sample. The Sharpe ratio of equity REITs is a respectable 0.50. Skewness is more negative at −0.82, again mostly because of the financial crisis. Mortgage REITs had return of 8.3%, a volatility of 21%, and a Sharpe ratio of 0.26, all substantially higher than in the first half of the sample. Mortgage REITs became more dominated by Agency REITs, which performed better than non-agency REITs, which dominated the previous period. The stock market as a whole had a Sharpe ratio of 0.49 and the bond market one of 0.47 over this period. Both bonds and stocks were marginally less volatile over this period than over the full sample. Over this 1994–2014 period, then, listed real estate was substantially more volatile than stocks as a whole. This should not come as a surprise given that real estate was at the heart of the Great Recession and the financial crisis. The Sharpe ratio on equity REITs is in line with that of stocks and long-term bonds, and roughly double that of mortgage REITs.

Correlations Next we turn to the correlation analysis on monthly returns. Over the full sample, we see that equity REITs have a 60% correlation with stock returns at the monthly frequency. While this is obviously a non-trivial positive correlation, it leaves open the possibility of substantial gains from diversification. Equity REITs have a 54% correlation with mortgage REITs due to the fact that both have assets whose cash flows ultimately derive from the performance of real-estate-linked assets. Equity REITs have only a 9% correlation with bond returns, similar to the 11% correlation of all stocks with bonds. Listed real estate

Table 9: **Return Summary Statistics REITs — Monthly Horizon**

Means, standard deviations, and Sharpe ratios are annualized. Skewness is the skewness of monthly returns.

	Equity REITs	Mortgage REITs	Stocks	T-Bond	T-bill
		Panel A: 1972–2014			
Mean	13.09	7.20	11.36	7.72	4.93
Standard Deviation	17.12	20.28	15.79	7.82	0.97
Sharpe Ratio	0.48	0.11	0.41	0.36	0.00
Skewness	−0.72	−0.26	−0.54	0.32	0.53
		Correlations			
Equity REITs	1.00	0.54	0.60	0.09	
Mortgage REITs	0.54	1.00	0.49	0.24	
Stocks	0.60	0.49	1.00	0.11	
Bonds	0.09	0.24	0.11	1.00	
		Panel B: 1994–2014			
Mean	12.60	8.28	10.29	5.99	2.68
Standard Deviation	19.79	21.25	15.39	7.02	0.63
Sharpe Ratio	0.50	0.26	0.49	0.47	0.00
Skewness	−0.82	−1.17	−0.76	−0.01	0.04
		Correlations			
Equity REITs	1.00	0.47	0.57	−0.02	
Mortgage REITs	0.47	1.00	0.39	0.15	
Stocks	0.57	0.39	1.00	−0.17	
Bonds	−0.02	0.15	−0.17	1.00	

is therefore not bond-like, contrary to some "folk wisdom." Mortgage REITs, which are akin to long-short bond portfolios, nevertheless only have a 24% correlation with Treasury bond returns. This is due to credit risk in non-agency REITs and prepayment risk embedded in Agency REITs. Mortgage REITs actually have a higher correlation (49%) with the stock market than with the bond market, despite being bond-like plays.

Modern REIT period   Panel B shows that, over the 1994–2014 period, the correlation of equity REITs with stocks was marginally lower (57%) and the correlation with bond returns was substantially lower (−2%). This reflects a broader and bigger shift in the correlation between stocks and bonds, which was −17% over this period. Economists have documented this change in the sign of the correlation between stocks and bonds from the pre-1994 to the post-1994 period (see Campbell, Pflueger, and Viceira, 2014; David and Veronesi, 2013; Hasseltoft, 2009; Song, 2014). They have ascribed it to a different monetary-policy regime, in which investors changed their perception about positive inflation shocks from being harbingers of bad news to being good news. Rather than being inflation bets, stocks became deflation hedges. Over this 21 year period, equity REITs displayed essentially no interest-rate risk. We return to this in the factor analysis below.

Effect of investment horizon   It is important to assess how the risk-return tradeoff changes with the investment horizon. Longer-term risk measures are arguably more appropriate for the long-term investment horizon of the GPFG. Also, correlation properties may differ by investment horizon, and longer-term correlations may be more reflective of the true risk, especially for real estate, which is a long-lived asset. We form annual returns by reinvesting the monthly returns over the twelve months in each year, and then reproduce the return statistics.

Means and variances   As Panel A of Table 10 shows, the mean returns are a little bit higher due to compounding within the year. Compounding may also explain the higher volatility of returns. The increase in mean return and volatility is especially dramatic for Mortgage REITs. Annual Sharpe ratios are similar to (annualized) monthly Sharpe ratios. The skewness of annual returns is similar to that of monthly returns for equity REITs but turns positive for Mortgage REITs. We caution that skewness is hard to measure reliably in small samples (43 annual observations). Turning to the 1994–2014 sample (21 annual observations) in Panel B, we see that annual mean returns and return volatilities are again substantially higher than the (annualized) monthly ones. Compounding strongly increases the average return on mortgage REITs but also dramatically increases their volatility. The

changes for equity REITs are modest by comparison. Contrary to some "folk wisdom" returns on REITs do not become less volatile as the horizon increases.

Table 10: **Return Summary Statistics REITs — Annual Horizon**

Means, standard deviations, and Sharpe ratios are annual. Skewness is the skewness of annual returns.

	Equity REITs	Mortgage REITs	Stocks	T-Bond	T-bill
		Panel A: 1972–2014			
Mean	13.84	9.27	12.23	8.07	5.09
Standard Deviation	18.27	29.25	18.35	9.49	3.44
Sharpe Ratio	0.48	0.14	0.39	0.31	0.00
Skewness	−0.74	0.12	−0.73	0.75	0.42
		Correlations			
Equity REITs	1.00	0.75	0.59	0.09	
Mortgage REITs	0.75	1.00	0.40	0.26	
Stocks	0.59	0.40	1.00	0.05	
Bonds	0.09	0.26	0.05	1.00	
		Panel B: 1994–2014			
Mean	13.01	11.17	11.33	6.25	2.73
Standard Deviation	19.76	33.85	19.47	8.87	2.24
Sharpe Ratio	0.52	0.25	0.44	0.40	0.00
Skewness	−0.90	0.18	−0.80	0.01	−0.03
		Correlations			
Equity REITs	1.00	0.69	0.44	−0.18	
Mortgage REITs	0.69	1.00	0.18	0.12	
Stocks	0.44	0.18	1.00	−0.35	
Bonds	−0.18	0.12	−0.35	1.00	

Correlations   The key correlations change very little at the annual horizon for the full sample. The correlation between equity REITs and stocks is 59% (versus 60% at monthly frequency) and that between equity REITs and bonds remains 9%. The correlation between stocks and bonds is 5% (vs. 11%). The biggest change is in the correlation of mortgage REITs with equity REITs (75% vs. 54%). The changes are more substantial over the 1994–2014 sample, where correlations between equity REITs and stocks shrink to 44% at the annual frequency (from 57% at the monthly frequency). Those with bonds are more negative at −18% (versus −2% at monthly frequency). Equity REITs were good hedges against interest rate risk over this 21-year period, again contrary to the widespread belief that real estate is bond-like. The effect of annual frequency is even larger for stocks who had a −35%

correlation with bonds at the annual frequency versus −17% at monthly horizon. Equity and mortgage REITs become more positively correlated (69% at the annual versus 47% at the monthly frequency).

Correlations at longer horizons   We also computed correlations for 24-month and 36-month excess returns over the period 1973–2014. We find a correlation between equity REITs and stocks of 61.4% at the monthly horizon, 59.4% at the 12-month horizon, 74.1% at the 24-month horizon, and 52.6% at the 36-month horizon. The correlation between equity REITs and bonds is 9.3% at the monthly horizon, 8.7% at the 12-month horizon, −8.5% at the 24-month horizon, and −3.7% at the 36-month horizon. These correlations should be interpreted cautiously since we only have 21 two-year periods and 14 three-year periods in the full sample. Obviously, going to even lower frequency correlations would aggravate the problem of data availability. This evidence is not particularly supportive of lower correlations between stocks and real estate at lower frequencies, but we do find some evidence that REITs become better interest rate hedges at lower frequencies. This analysis, which indicates instability in longer-term correlations prompts us to investigate time variation in volatilities and correlations in more depth.

Time-variation in volatility and correlations   The previous analysis calculated volatilities from a long sample of 43 years and a shorter sample of 21 years of data. To investigate the issue of time-variation in volatilities and correlations further, we compute 120-month rolling window volatilities and correlations.

Volatility   Figure 32 shows that equity REIT returns became substantially more volatile in the last decade, which includes the Great Recession, compared with the decades (120-month windows) that ended before the year 2000. An important question is whether the volatility of equity REITs will return to its pre-2000 levels once the crisis period is no longer inside the 120-month window. All else equal, the elevated volatility makes equity REITs less attractive and will reduce their role in the optimal portfolio. The volatility of mortgage REITs in the top right panel shows different dynamics. Earlier mortgage REIT crises led to elevated volatility in the 1970s and 1990s. Volatility peaked in the decade between 1999–2009 and has been falling modestly since then, as larger Agency Mortgage REITs started to dominate the sample. Equity return volatility, plotted on the same axis in the bottom left panel, looks muted by comparison. Equity REIT volatility seems to decouple from broader stock market volatility over the past 15 years. Finally, bond return volatility, plotted against a much smaller vertical axis range (6-12%) fell after the Volcker disinflation through the mid

1990s, and has hovered around 7% since.
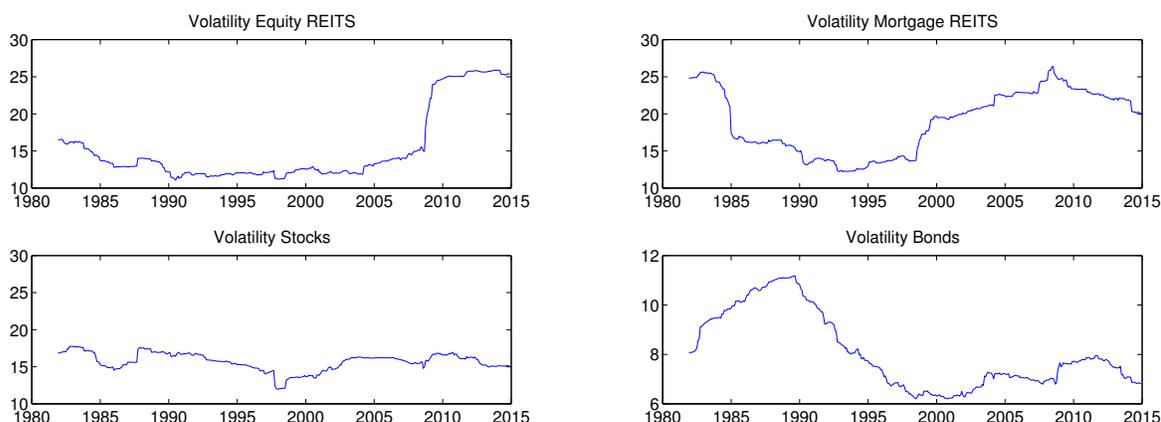
Figure 32: Time-Varying Volatilities

The figure plots annualized standard deviations of equity REIT, mortgage REIT, stock market, and 10-year Treasury bond returns. Each point reflects the standard deviation of the last 120 months of returns, multiplied by $\sqrt{12}$. The sample period is January 1972 until December 2014.

Correlation: Equity REITs with stocks Figure 33 shows the rolling-window correlations. The left panel shows that equity REITs had a correlation of around 70% with the overall stock market in the 1970s and 1980s. The correlation then fell precipitously in the 1990s, in part because of the technology sector boom and bust which had little effect on equity REITs. Over the last decade, the correlation between stocks and equity REITs rose dramatically as the overall decline and rebound in the stock market was very much connected to the performance of commercial (and residential) real estate markets. The latest observation, which reflects the correlation between January 1995 and December 2014 is 77.8%, an all-time high. Gains from diversification into equity REITs turned out be modest over this period. A key question is whether this correlation will revert to its lower long-term mean. We expect that it will once the financial crisis period is no longer part of the rolling windows. However, real estate capital markets are much more integrated with overall capital markets than they were in past decades. Hence the scope for a reduction in correlation seems limited.

Correlation: Equity REITs with bonds The middle panel of Figure 33 shows the correlation between equity REITs and bonds. This correlation has changed from mildly positive (+25%) in the first 30 years of the sample to mildly negative in the most recent 15

168

years (-10%). These dynamics mirror the correlation between the overall stock market and bond market in the right panel. Equity REITs saw less of a fall in correlation with bonds (to -10%) than the overall stock market did (to -30%), possibly a reflection of higher interest rate sensitivity of equity REITs compared to the universe of stocks. Stocks and bonds have been great complements in a portfolio over the past 15 years. The flight-to-safety feature of U.S. Treasuries and the continuing decline in bond yields over the past 5 years (in part due to Quantitative Easing policies) alongside a stock market rally account for the fall in correlation. With bond yields still at historically low values, and bound to go up, and with stock markets trading at high multiples, it is conceivable that both stock and bond prices may fall in unison. This would increase the correlation between stocks and bonds, possibly pushing it back into positive territory.
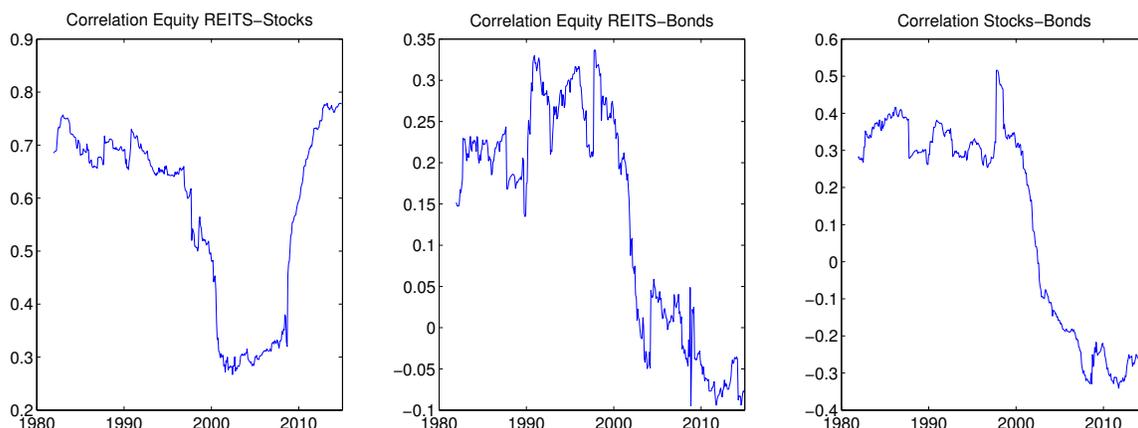
Figure 33: Time-Varying Correlations

The figure plots correlations between equity REIT, stock market, and 10-year Treasury bond returns. Each point reflects the correlations of the last 120 months of returns. The sample is January 1972 until December 2014.

Our conclusion from this section is that U.S. equity REITs have substantial correlations with stocks, but low correlations with bonds. In addition, correlations display substantial time variation, making it important to supplement unconditional mean-variance portfolio analysis with analysis that considers changes in risk.

B.4 Factor analysis

Next, we investigate the performance of publicly traded commercial real estate in the U.S. using standard asset pricing factor models. The analysis serves to advance the univariate

correlation analysis of the previous section. In addition to understanding what risks commercial real estate is exposed to, we can investigate whether equity REITs have displayed abnormal performance (alpha) relative to the factors considered.

Full-sample results We start with an analysis for the full 1972–2014 sample of monthly returns (516 months). Table 11 shows the results, reporting both point estimates and Newey-West t-statistics.

Table 11: **Analyzing equity REIT Performance 1972–2014**

The dependent variable is the excess return on the equity REIT index. The independent variables are listed in the main text. The first row reports the intercept α, the other rows report risk factor exposures β. The t-statistics are computed using Newey-West standard errors with one lag. The last but one row reports the R^2 of the regression. The last row reports the expected return according to the regression model. It includes the risk-free rate and excludes the alpha. The data are monthly from 1972–2014.

	(1)	(2)	(3)	(4)	(5)	(6)
α	0.33	0.32	−0.03	−0.06	0.06	0.03
t-stat	1.84	1.77	−0.21	−0.40	0.36	0.19
β^s	0.65	0.65	0.71	0.70	0.69	0.68
t-stat	10.06	9.76	12.83	12.17	13.29	12.72
β^b	——	0.06	——	0.14	——	0.16
t-stat	——	0.56	——	1.73	——	2.05
β^{smb}	——	——	0.41	0.43	0.41	0.44
t-stat	——	——	7.72	7.99	8.11	8.37
β^{hml}	——	——	0.67	0.67	0.64	0.64
t-stat	——	——	8.31	8.22	8.86	8.73
β^{mom}	——	——	——	——	−0.09	−0.10
t-stat	——	——	——	——	−1.57	−1.72
β^{liq}	——	——	——	——	−0.03	−0.02
t-stat	——	——	——	——	−0.56	−0.41
R^2	36.09	36.16	54.11	54.48	54.73	55.21
Exp. ret.	9.12	9.26	13.49	13.85	12.33	12.68

CAPM The first column shows the standard CAPM. Equity REIT excess returns have a stock market beta of $\beta^s = 0.65$, which is estimated very precisely with a t-stat of 10. The monthly outperformance of equity REITs relative to the CAPM is 0.33% per month or about 4% per year. Covariation with the stock market alone explains 36.1% of the variation in equity REIT returns.

Two-factor model with stocks and bonds Column (2) adds a 10-year Treasury bond excess return as the second factor. Equity REITs have essentially zero exposure to the bond market, consistent with the correlation analysis. The resulting 2-factor α is essentially unchanged at 0.32% per month. The two factor model explains 36.2% of monthly return variation in equity REITs, with the bond factor adding only marginally to the R^2. In other words, 2/3 of the variation in publicly listed commercial real estate index returns in the U.S. is unaccounted for by a stock and a bond index. This is an important observation in light of our discussion of the Opportunity Cost model below.

Fama-French Column (3) reports results on the standard three-factor Fama-French model. It shows that equity REITs have substantial loadings on the size (small-minus-big) factor and on the value (high-minus-low book-to-market) factors. In other words, equity REITs behave like small value stocks. The R^2 of the three-factor model increases to 54.1%, a substantial jump compared to the 2-factor model in column (2). Given that the SMB and HML factor are not easily (and cheaply) tradable, it would be good to replicate this result with a Russell 2000 value index, a liquid small value stock index.

Fama-French + bonds Column (4) adds the bond factor to the three-factor FF model. The bond beta increases somewhat to 0.14 once smb and hml factors are included in the regression. Its t-statistic rises to 1.7. The four-factor model's R^2 is 54.5%.

Momentum and liquidity In Columns (5) and (6), we add two more equity risk factors which have become standard in the asset pricing literature: the momentum factor and the Pastor-Stambaugh illiquidity risk factor. The latter is the return on a tradable strategy that goes long illiquid stocks and short liquid stocks. Equity REITs have a small negative exposure to both the momentum factor and the illiquidity factor, but neither exposure is different from zero in Column (5). Once the bond factor is included, the momentum exposure becomes marginally significant at the 10–15% level. More interestingly, the bond factor exposure of equity REITs now becomes significantly positive: the bond beta is 0.16 with a t-statistic of 2.05. The market, smb and hml factors retain their importance. The model in Column (6) explains 55.2% of the return variation. While this is substantially higher than the 36% of the CAPM, there remains a large component of equity REIT returns, call it a real estate factor, that is not captured by standard stock and bond portfolios.

Finally, abnormal returns, which were 32 basis points per month in the 2-factor model, disappear completely once the size and value factors are included in the analysis. There is no indication that equity REITs as an asset class have outperformed a portfolio of stocks and bonds — once a small-stock and value-stock portfolio are included in the analysis.

Cost of equity capital   The last row of Table 11 reports the model-implied expected return, or equivalently the cost of equity capital. It is constructed as the estimated beta on each factor multiplied by the average excess return on the corresponding factor (ex-post average returns over 1972–2014), summed across factors, plus a risk-free rate (set equal to the historical risk-free rate over the same sample). The six-factor model in column (5) implies an expected return of 12.68%. The latter consists of a compensation of 4.9% for time value of money (the average annualized one-month T-bill rate), compensation for overall stock market risk of 4.4%, compensation for bond market risk of 0.4%, compensation for small stock risk of 1.0% and value risk of 2.9%, and risk discounts of −0.8% for momentum and −0.1% for illiquidity risk. The two-factor model with a stock and a bond index (column 2) implies an expected return of only 9.26%, or 3.4% lower than that in column (6). This 3.4% is abnormal return according to the two-factor model but compensation for additional sources of risk (especially small value risk) according to the six-factor model, where there is no residual alpha. Put differently, investing in commercial real estate allows one to capture this small value premium.

Modern REIT sample: 1994–2014   Table 12 reports the same set of results for the period 1994–2014. We highlight only a few salient points.

- The best-fitting model continues to be the 6-factor model in column (6). It explains 59.5% of equity REIT return variation, slightly more than in the full sample (55%). The stock beta is higher (0.75), in line with the increased importance of the financial crisis in this shorter sample.
- The bond beta is also much higher (0.39) and more significant. This suggests that equity REITs are indeed exposed to interest-rate risk, but a univariate correlation analysis like that above misses this exposure completely (recall the −18% univariate correlation for this sample). In fact, even the two-factor model substantially understates the bond exposure (by a factor of almost two). This underscores the importance of a rich model of risk.
- The *smb* beta is a bit higher at 0.54. The *hml* beta is substantially higher at 0.88. The momentum beta is now significantly negative with a t-statistic of −2.4, suggesting that real estate stocks provide a hedge against momentum risk. Finally, the Pastor-Stambaugh illiquidity beta remains zero.
- As in the full sample, we find that equity REITs are fairly priced: they display a zero alpha relative to all models and conventional statistical levels of significance, but especially in the last four models. The cost of capital for equity REITs in the shorter sample is 12.2%, comprising a 2.7% time value of money compensation and 9.5% risk

172

compensation: a 5.7% equity risk premium, a 1.3% bond risk premium, a 1.1% small-stock risk premium, a 2.2% value risk discount, a -0.9% momentum risk premium, and a 0.1% illiquidity risk premium. The decomposition is quite similar to that in the full sample.

Table 12: **Analyzing equity REIT Performance 1994–2014**

The dependent variable is the excess return on the equity REIT index. The independent variables are listed in the main text. The first row reports the intercept α, the other rows report risk factor exposures β. The data are monthly from 1994–2014.

	(1)	(2)	(3)	(4)	(5)	(6)
α	0.36	0.29	0.05	-0.06	0.17	0.03
t-stat	1.15	0.89	0.18	-0.25	0.59	0.12
β^s	0.73	0.75	0.79	0.81	0.74	0.75
t-stat	5.90	6.11	9.15	9.76	8.85	9.44
β^b	$--$	0.21	$--$	0.33	$--$	0.39
t-stat	$--$	1.04	$--$	2.24	$--$	2.61
β^{smb}	$--$	$--$	0.48	0.51	0.50	0.54
t-stat	$--$	$--$	6.28	6.75	6.67	6.99
β^{hml}	$--$	$--$	0.92	0.93	0.87	0.88
t-stat	$--$	$--$	8.47	8.55	8.90	9.11
β^{mom}	$--$	$--$	$--$	$--$	-0.14	-0.17
t-stat	$--$	$--$	$--$	$--$	-2.17	-2.42
β^{liq}	$--$	$--$	$--$	$--$	-0.02	0.01
t-stat	$--$	$--$	$--$	$--$	-0.34	0.12
R^2	32.13	32.69	56.37	57.63	57.85	59.53
Exp. ret.	8.23	9.07	12.05	13.38	10.61	12.21

Time-variation in factor risk exposure To investigate how the risk exposure of equity REITs to the various risk factors has changed over time, we estimate the two-factor and six-factor risk models and estimate them over rolling 120-month windows. Figure 34 plots the resulting multivariate risk factor exposures. The first two panels display the stock and bond market betas. They are the multivariate counterparts to the first two panels in Figure 33, except that the latter measure correlations while these graphs measure betas.[64] Several interesting patterns emerge:

---

[64]The beta equals the correlation times the ratio of the standard deviation of the equity REIT to the standard deviation of the risk factor.
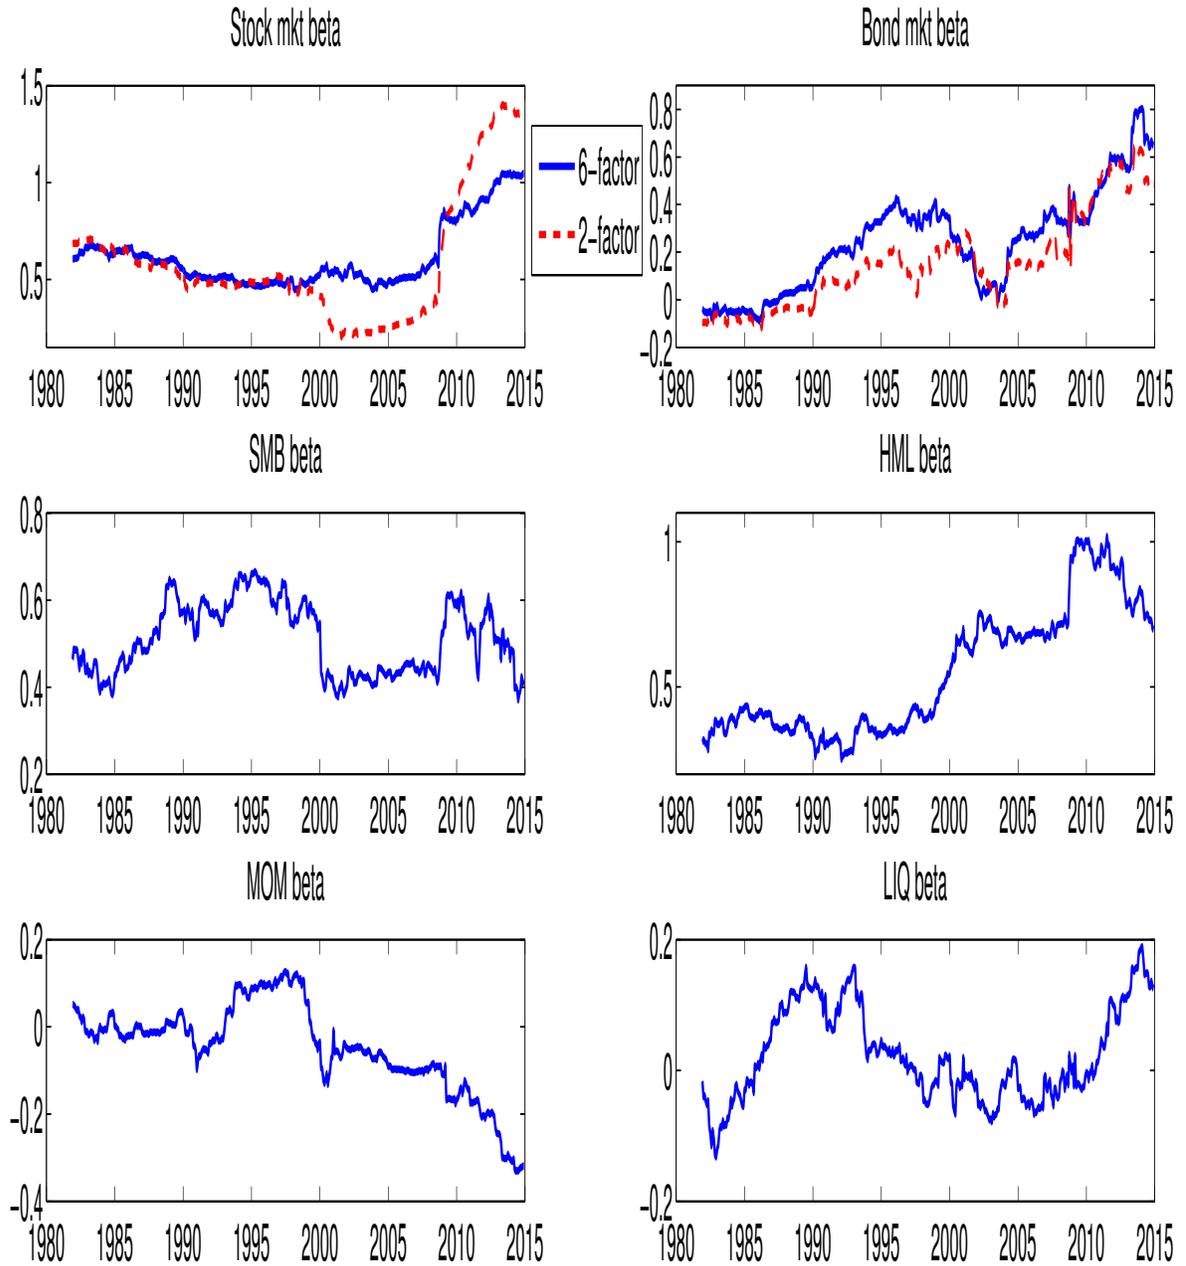
Figure 34: Time-Varying Betas for equity REITs

The figure plots the exposures (betas) of equity REITs to six risk factors: the stock market, the 10-year bond market return, the size (SMB) factor, the value (HML) factor, the momentum (MOM) factor, and the Pastor-Stambaugh illiquidity factor. Each set of risk-factor exposures is estimated via a multivariate regression using the most recent 120 months of data. The sample period is January 1972–December 2014.

174

1. Equity REITs have seen a strongly increasing stock market beta in the last twenty years, from around 0.5 to 1. This is a combination of a rise in the correlation between the two return series and a rise in the relative volatility of equity REITs relative to stocks (recall Figures 33 and 32). Also plotted in the figure is the stock market beta obtained from the two-factor model with only stock and bond factors (red dashed line). It shows a much lower equity market beta in the samples ending in 2000–2010 and a much higher equity beta in the samples ending in 2010–2015. Omitted risk factors bias the stock market beta for equity REITs. The difference between the solid and dashed lines serves as a cautionary tale for omitting important risk factors.

2. The bond market beta of equity REITs remains mostly positive and looks dramatically different from the correlation between equity REITs and bonds. In fact, the signs are different in the second half of the sample. That difference underlines the importance of multi-variate analysis. Further, the decline in the correlation between equity REITs and bonds is offset by the strong rise in the relative volatility of equity REITs versus bonds. The result is a strongly increasing bond beta, which reaches values of around 0.7 at the end of the sample. Also plotted in the figure is the bond market beta obtained from the two-factor model with only stock and bond factors (red dashed line). The dashed line tracks the solid line quite well, so the two models have similar implications for the interest-rate risk of equity REITs.

3. The exposure of equity REITs to the small stock factor is relatively stable over time (middle left panel). In contrast, the exposure to value stocks more than doubles over time (middle right panel). The momentum risk of equity REITs turns from zero to substantially negative towards the end of the sample (bottom left panel).

4. The illiquidity risk exposure is small and relatively stable.

5. The corresponding six-factor alphas, shown in the solid line in the left panel of Figure 35 fluctuate between +0.3% and −0.3% per month. The last reading for the ten-year period ending in December 2014 is −0.18%. The dashed red line is for the 2-factor model and shows that one would dramatically overstate the abnormal returns on commercial real estate stocks if one omitted size, value, momentum and illiquidity factors in the analysis. A simple passive strategy into an equity REIT sector index would generate an annual alpha of 6% or more for much of the sample. That outperformance is spurious once additional risk factors are taken into account. We caution that estimated alphas over short horizons (120 months) are to be interpreted with caution.

6. The six-factor model explains a strongly fluctuating fraction of equity REIT return variation, as shown in the right panel of Figure 35. The last reading for the ten-year period ending in December 2014 is 72.7%, up from only 35.7% in the decade ending

in 2002. The model also worked well in the 1970s and 1980s, but not that well during the 1990s and 2000s. The two-factor model performs much worse in this dimension, with the gap being especially large in the 1990s and 2000s.
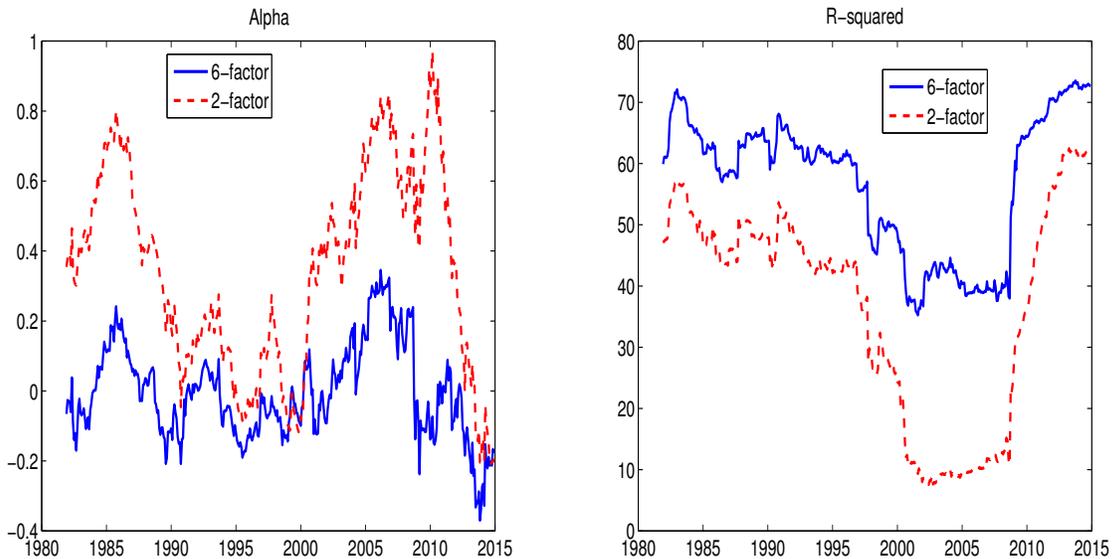

Figure 35: Time-Varying alpha and R^2 for equity REITs
The left panel plots the abnormal monthly returns (alphas) of equity REITs from a six risk factor model that includes: the stock market, the 10-year bond market return, the size (SMB) factor, the value (HML) factor, the momentum (MOM) factor, and the Pastor-Stambaugh illiquidity factor. Each alpha value is estimated using a multivariate regression using the most recent 120 months. The right panel plots the corresponding R^2 values of these rolling-window regressions. The sample is January 1972 until December 2014.

Cost of capital The time-varying beta model can be applied to calculate the fair expected return on equity REITs. Figure 36 plots that cost of capital over time. The left panel is for the six-factor model, while the right panel is for the two-factor model. To make the figure, we hold the average risk premium on each factor constant at its full-sample average (given the difficulty in estimating average returns) and only allow the betas to fluctuate over time. This avoids misinterpreting negative realized returns on factors as periods with low risk premia. The six-factor model implies an annual risk premium that is on average 3.8% higher than that in the two-factor model. The gap is as large as 6.4% for the decade ending in 2002 and as low at 0.5% for the decade ending in December 2014. In both panels, the bulk of the risk premium is earned for exposure to regular stock market risk (about 4%), with another 2.5% on average for exposure to value stock risk and 1.2% for exposure to small

176

stock risk. Bond market risk only contributes about 0.7% to the annual risk premium on average, while momentum exposure reduces it by 0.4%.
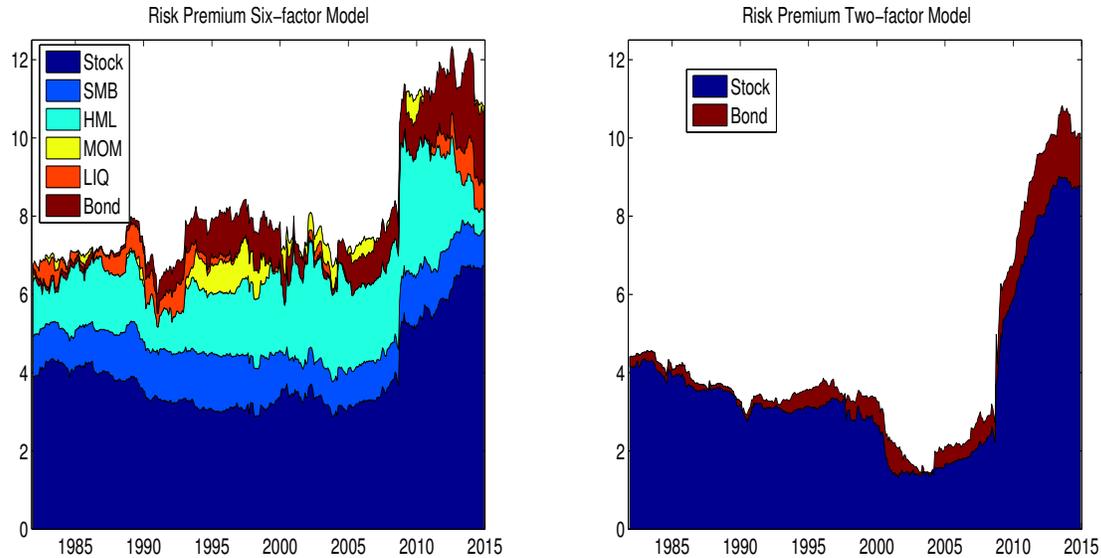

Figure 36: Risk Premium Decomposition for equity REITs

The left panel plots the expected return on equity REITs as implied by the six-factor model. The risk factors are: the stock market, the 10-year bond market return, the size (SMB) factor, the value (HML) factor, the momentum (MOM) factor, and the Pastor-Stambaugh illiquidity factor. The right panel plots the expected return on equity REITs as implied by the two-factor model, which contains only the stock market and the 10-year bond market return. The betas on the factors are estimated on 120-month rolling windows. To calculate the risk premium, we multiply each beta with the average return on each factor, where the averages are computed over the full 1972–2014 sample.

Risk premium   The estimated risk premium at the end of the sample in December 2014 is 10.6% per annum. It contains a 6.7% equity risk premium, a 3.2% value risk premium, a 0.9% small stock premium, a 1.8% bond risk premium, a -2.7% momentum risk discount, and a 0.7% illiquidity risk premium. To get to the cost of capital we must add in a risk-free rate, which was 1.4% over the past 10 years. Combining delivers a cost of capital for U.S. equity REITs at the end of the sample of exactly 12%. Below we explore what this cost of capital implies for the valuation of REITs today.

177

B.5 Portfolio analysis

Despite the lack of a positive alpha associated with investing in U.S. commercial real estate stocks, the 0.6 correlation between REITs and stocks, and their near-zero correlation with bonds, suggest that they may still have a prominent place in a portfolio because of gains from diversification.

REITs, stocks and bonds: full sample We start by considering a portfolio of U.S. equity REITs, U.S. stocks, and U.S. Treasury bonds over the full 1972–2014 sample. As inputs we use the observed mean returns and return covariance matrix. The optimal portfolio calls for 32% equity REITs, 23% stocks, and 46% bonds. It has an annualized mean return of 10.2%, a volatility of 9.2%, and a Sharpe ratio of 0.58. The prominence of REITs is due not only to their favorable diversification benefits but also to their higher returns (13.1% versus 11.4% for stocks and 7.7% for bonds).

We investigate sensitivity of the optimal portfolio to the average return on equity REITs, an object which is notoriously hard to estimate precisely. If we lower the mean return on equity REITs to 11.4% per year, the same as for stocks, keeping everything else the same, the portfolio weight of REITs falls from 32% to 21%, while the weight on stocks rises from 23% to 29% and that on bonds from 46% to 50%. This portfolio has a mean return of 9.9%, a volatility of 8.6%, and a Sharpe ratio of 0.57.

If we further lower the expected return on REITs to 9.26% (the expected return according to the two-factor model) and hold all other inputs fixed, we get a REIT portfolio share of 6.4% (stocks 39.7%, bonds 53.9%). This weight is consistent with observed practice. It requires one to assume that over the long-run REITs will not earn any additional return beyond their compensation for stock and bond market risk. However, if we assume that REITs will continue to earn the extra 3.8% return they have earned over the past 40 years, their share in a portfolio of stocks and bonds ought to be much larger.

Seven-asset portfolios What if we add the other portfolios that feature in the six-factor model? The tangency portfolio with 7 assets (equity REITs, stock market, SMB, HML, MOM, LIQ, and bonds) produces an in-sample Sharpe ratio of 1.16, twice as large as that with 3 assets (equity REITs, stock market, and bonds). Its annual return is 10.3%, which is the same as that for the 3 asset portfolio, but the standard deviation is only half as large at 4.6%. All seven assets have positive portfolio weights. The weight on equity REITs is 2.5%, that on stocks 15.4%, and that on bonds 13.4%. The four stock portfolios that were added have the following weights: SMB 7.3%, HML 27.9%, MOM 18.9%, and LIQ 14.8%. Clearly, equity REITs are displaced to a large extent by other stock portfolios, notably those

in the value portfolio.Once the investment opportunity set is expanded and REITs no longer have abnormal performance relative to the expanded opportunity set, their weight in the portfolio drops dramatically from 32% to 2.5%.

Modern REIT sample: 1994–2014   Over the modern REIT sample, the position of equity REITs in the portfolio is similarly large.   REITs represent 26.5% of the optimal portfolio, compared to 33% for stocks and 40.5% for bonds. The tangency portfolio has an average return of 9.2% per year, with an annualized volatility of 9.3% and a Sharpe ratio of 0.70. The higher Sharpe ratio is largely accounted for by the much lower risk-free rate over the last 21 years. Just as in the full sample, the weight on real estate is very large, and far beyond typical institutional portfolio weights.

Reducing the expected return on REITs to that of stocks, 10.3%, lowers the weight on REITs to 19%, while increasing the weight on stocks to 37% and that on bonds to 44%. Lowering the expected return on REITs to 9.07% (its expected return according to the two-factor model over this period), we get a REIT portfolio share of 15% (stocks 39%, bonds 46%).

The tangency portfolio with 7 assets (equity REITs, stock market, SMB, HML, MOM, LIQ, and bonds) produces an in-sample Sharpe ratio of 1.14, 50% higher than the 0.70 Sharpe ratio with 3 assets (equity REITs, stock market, and bonds). Its annual return is 7.45% and the standard deviation is only 4.2%. The weight on equity REITs is 1.5%, that on stocks 16.3%, and that on bonds 33%. The four stock portfolios that were added have collectively a weight of 50%: SMB 6.6%, HML 17.3%, MOM 8.8%, and LIQ 16.5%. As in the full sample, equity REITs are almost entirely displaced by value and illiquidity stock portfolios, even under historical mean-return assumptions.

We conclude that static portfolio analysis indicates a prominent place for real estate in the portfolio alongside stocks and bonds. The results are consistent whether we use the full 43-year sample or only the last 21 years of returns. Both the high average returns on REITs and their diversification benefits contribute to this finding.

Dynamic portfolio choice   The static portfolio analysis ignores time variation in risk. Next, we investigate the sensitivity of the portfolio choice to having volatilities and correlations that change over time. Each month, starting in January 1982, we estimate the covariance matrix of equity REITs, stock, and bond returns using the past 120 months of observations, and compute the tangency portfolio weights. We hold this portfolio for one month and compute the return in that month. The next month, we recompute the covariance matrix and rebalance. Because average returns are hard to estimate, we set the average

returns on the three assets equal to their full-sample counterparts. We set the risk-free rate equal to 4.2%, its average value after 1982. The left panel of Figure 37 shows the evolution of the portfolio shares over time. It shows a clear downward trend in the weight of equity REITs in the tangency portfolio, from a high of 50% in the 1980s to just above 10% at the end of the sample. By construction, the change in weights is solely due to the changes in the covariance matrix of returns. In particular, three important changes documented above contribute. First, the volatility of equity REITs has been rising over the past twenty years, in absolute terms and relative to that of stocks and bonds. Second, the correlation of equity REITs with stocks has been rising over the past 10–15 years. Third, the correlation of stocks and bonds has been falling more than that between equity REITs and bonds, increasing the benefit of combining stocks and bonds (at the expense of equity REITs). In conclusion, the diversification benefits from including equity REITs in a portfolio of stocks and bonds have declined over time.

Seven-asset portfolios We repeat the time-varying portfolio analysis for the 7-asset portfolios. The bottom panel of Figure 37 shows the evolution of the portfolio shares over time. Consistent with our earlier observation, equity REITs are a lot less prominent in the 7 asset portfolio. And consistent with the three-asset portfolio, their importance in the tangency portfolio shrinks over time. In fact, starting in early 2008 (based on returns measured over the previous decade), the weight in equity REITs becomes negative. The other noteworthy change is a growing share in bonds. At the end of the sample, the weight in equity REITs is -4.35%. The weight in stocks is 13%, that in bonds is 32.5%. SMB gets 6.4%, HML 31.2%, MOM 9.7% and LIQ 11.5%. As shown above, these results are sensitive to which average returns are used.

B.6   U.S. REIT sector analysis

The analysis above was for the equity REIT index. We now analyze separately, the various commercial real estate sectors using REIT index data that are sector-specific. We start by showing return summary statistics and then go straight to the multivariate risk analysis, skipping the univariate correlation analysis.

B.6.1   Return properties

NAREIT provides monthly returns from January 1994 onwards for the following REIT sectors: residential (apartments), office, industrial, retail (shopping malls), lodging (hotels), health care, diversified, and self-storage. There is also a short return time series available
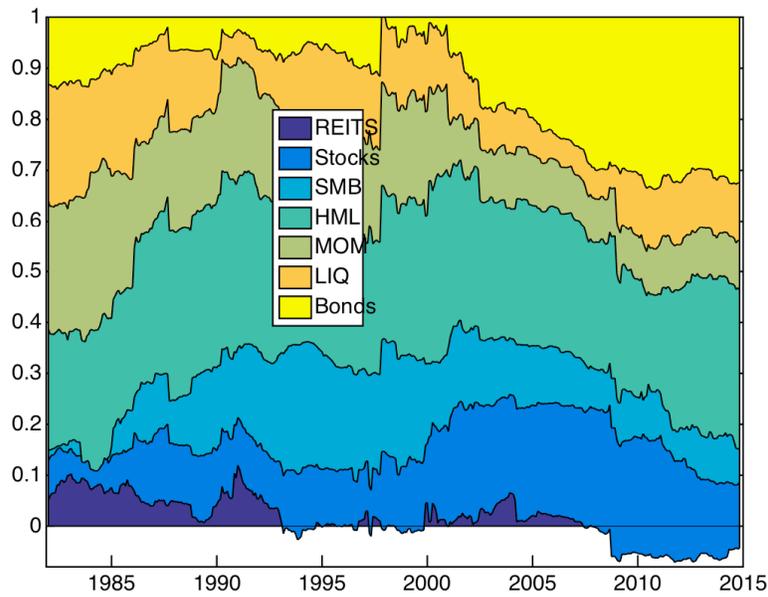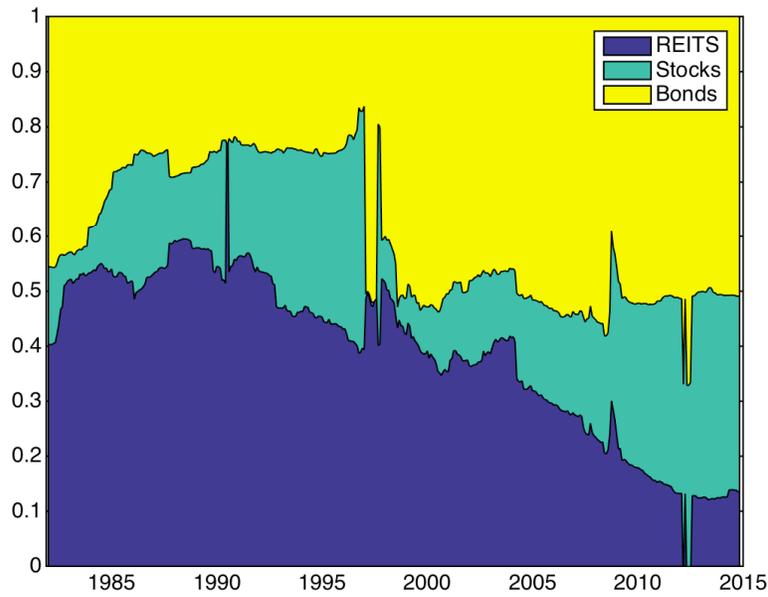
Figure 37: Dynamic Tangency Portfolio

The top panel plots the weights in the tangency portfolio of a three-asset portfolio: U.S. equity REITs, U.S. stocks, and U.S. bonds. The bottom panel plots the weights in the tangency portfolio of a seven-asset portfolio: U.S. equity REITs, U.S. stocks, and U.S. bonds, the size (SMB) factor, the value (HML) factor, the momentum (MOM) factor, and the Pastor-Stambaugh illiquidity factor (LIQ). Each month, starting in January 1982, the covariance matrix of returns is recomputed using the last 120 months of returns. Average returns on the risky assets are computed over the full 1972–2014 sample.

for timber (from December 2010) and for infrastructure (from January 2012) REITs. As of December 2014, the All Equity REIT index contained 177 REIT stocks with a market capitalization of $846 billion. Of those, 37 were retail (25% of the market capitalization), 34 diversified (11%), 22 office (10.5%), 21 residential (13%), 19 lodging (6.5%), 17 health care (11.4%), 14 industrial (6%), 4 self-storage (5.3%), 5 timber (3.8%), and 4 infrastructure (7.8%). We do not include timber and infrastructure REITs in this part of the analysis but focus on the eight main sectors with 256 months of data available.

Table 13 displays the return summary statistics in the top panel and it displays correlations with the All Equity REIT index return, with the overall stock market return, and with the bond market return. All correlations are calculated based on excess returns, as before. The first column repeats the moments for the All Equity REIT return, the same series we used in the previous section. They serve as a point of comparison for the sector indices, which collectively make up the All equity REIT index. We see that sector index returns range from 10.8% per year for lodging to 17.4% for self storage. Annualized volatilities range from 19.65% for residential to 30.94% for industrial. Sharpe ratios range from 0.26 for hotels to 0.75 for self-storage. Skewness ranges from −0.78 for residential to +0.94 for hotels. In other words, there is quite a big difference in the return moments, even though each sector is itself a diversified portfolio of multiple REITs. The range of Sharpe ratios in especially pronounced.

Table 13: **Return Summary Statistics REIT Sectors — Monthly Horizon**

Means, standard deviations, and Sharpe ratios are annualized. Skewness is the skewness of monthly returns. The second panel reports correlations with equity REIT index returns, stock returns, bond market returns, and inflation. Data are monthly from January 1994 until December 2014.

	All	Resi	Office	Indus	Retail	Hotel	Health	Divers	Selfstore
Mean	12.60	13.48	13.52	13.91	14.00	10.83	14.79	11.06	17.41
St Dev	19.79	19.65	21.68	30.94	22.49	31.07	20.81	21.17	19.71
Sharpe	0.50	0.55	0.50	0.36	0.50	0.26	0.58	0.40	0.75
Skewness	−0.82	−0.78	−0.49	0.28	−0.35	0.94	−0.34	−0.25	−0.47
					Correlations				
Eq REITs	1.00	0.90	0.95	0.86	0.96	0.83	0.84	0.93	0.84
Stocks	0.57	0.50	0.56	0.50	0.50	0.59	0.40	0.55	0.41
Bonds	−0.02	−0.05	−0.05	0.04	−0.00	−0.17	0.08	−0.05	0.09
Inflation	0.05	0.05	0.09	0.04	0.06	−0.00	0.01	0.06	−0.02

The second panel of Table 13 shows univariate correlations. The correlation with the overall equity REIT index is as low as 0.83–0.84 for hotels, health care, and self-storage REITs, and as high as 0.95–0.96 for the traditional large sectors, retail and office. Clearly, a

portfolio mostly comprised of retail and office REITs will behave very much like the overall equity REIT index. the correlations with the overall stock market range little, from 0.40 for health care to 0.59 for hotels. Next, the univariate correlations with bond returns range from -0.17 for hotels to 0.08 for health care and 0.09 for self storage. Finally, the correlations with inflation vary in a tight range between -0.02 for self-storage and 0.09 for office. No sector appears to be a good inflation hedge based on these statistics.

B.6.2 Factor analysis

We now perform the same factor analysis as the one we considered for the equity REIT index. In the interest of space, we confine attention to the two-factor model, with a stock and a bond market factor, and the six-factor model, which contains four additional stock portfolios. Table 14 presents the results. For comparison with our previous results, we repeat the analysis for the All Equity REIT index in the first column. The next eight columns report the eight-sector index results.

Panel A shows the results for the two factor model. Two factor alphas are positive for all sectors, though only significantly so for the self storage sector, where they are 75 basis points per month or 9% per year. The health care sector also saw abnormal two-factor returns of 6% per year. Stock market betas range from 0.56 for self storage and 0.58 for health care on the low end to 1.04 for industrial and 1.17 for hotels on the high end. The former are the "defensive" REIT sectors whereas the latter are the "cyclical" REIT sectors.

Bond betas in the two-factor model range from -0.33 for hotel to $+0.45$–0.46 for health care and self storage. The two largest sectors, Retail and Office look quite similar to the sector as a whole. The two factor model explains 35.7% of the variation in hotel returns but only 18–19% in the variation of health care and self-storage returns.

Panel B reports the results for our favorite 6-factor model. The first main observation is that all abnormal returns are substantially lower, about 25 basis points, to the point that most are economically and all are statistically indistinguishable from zero. Only self-storage has a 6-factor α of 51 basis points per month (6% per year), which is economically large but only statistically significant at the 10–15% level.

Six-factor stock betas are very close to the two-factor stock betas. However, six-factor bond betas are substantially higher than 2-factor bond betas, as we pointed out for the index before. The least interest rate sensitive sectors are hotels followed by residential with bond betas of 0.03 and 0.22, respectively. The most interest rate sensitive sectors are health care, self-storage and especially industrial. While this is an incomplete explanation, the ranking of bond betas is consistent with the typical lease duration for the various kinds of real estate. Residential and hotels have typically short lease durations, while health care and industrial

Table 14: **Analyzing Sector REIT Performance 1994–2014**

The dependent variable is the excess return on the equity REIT index in Column (1) and the 8 REIT subsectors in Columns (2)–(9). The independent variables are listed in the main text. The first row reports the intercept α, the other rows report risk factor exposures β. Panel A is the Two-factor model; Panel B is the six-factor model. The data are monthly from 1994–2014.

	(1) All	(2) Resi	(3) Office	(4) Indus	(5) Retail	(6) Hotel	(7) Health	(8) Divers	(9) Selfstore
				Panel A: Two-factor model					
α	0.29	0.46	0.35	0.11	0.39	0.03	0.52	0.17	0.75
t-stat	0.89	1.35	1.00	0.23	0.98	0.05	1.44	0.48	2.23
β^s	0.75	0.65	0.80	1.04	0.75	1.17	0.58	0.76	0.56
t-stat	6.11	5.74	6.42	4.77	4.99	6.71	5.23	6.33	5.46
β^b	0.21	0.11	0.15	0.58	0.27	-0.33	0.46	0.15	0.45
t-stat	1.04	0.61	0.67	1.43	1.01	-1.12	1.93	0.68	2.07
R^2	32.69	25.60	31.77	26.31	25.62	35.68	18.41	30.09	19.00
				Panel B: Six-factor model					
α	0.03	0.21	0.08	-0.20	0.14	-0.24	0.31	-0.05	0.51
t-stat	0.12	0.76	0.25	-0.38	0.37	-0.60	0.96	-0.18	1.73
β^s	0.75	0.69	0.83	1.05	0.72	1.08	0.57	0.76	0.54
t-stat	9.44	8.79	9.49	6.38	7.02	12.66	7.04	10.88	6.44
β^b	0.39	0.22	0.30	0.78	0.50	0.03	0.63	0.33	0.61
t-stat	2.61	1.53	1.77	1.97	2.59	0.18	3.11	2.28	3.64
β^{smb}	0.54	0.45	0.51	0.64	0.62	0.88	0.40	0.62	0.54
t-stat	6.99	6.27	5.48	3.28	6.00	7.10	4.00	7.69	6.38
β^{hml}	0.88	0.85	0.92	1.06	0.93	1.27	0.74	0.94	0.67
t-stat	9.11	9.04	8.79	5.52	7.48	10.65	6.02	10.40	6.19
β^{mom}	-0.17	-0.10	-0.13	-0.18	-0.26	-0.50	-0.22	-0.21	-0.10
t-stat	-2.42	-1.76	-1.60	-1.50	-2.51	-3.12	-3.67	-2.34	-1.73
β^{liq}	0.01	-0.04	-0.01	-0.00	0.04	0.04	0.06	-0.06	0.02
t-stat	0.12	-0.54	-0.17	-0.02	0.42	0.36	0.78	-0.86	0.20
R^2	59.53	49.07	54.64	41.97	52.41	66.57	37.45	59.39	36.20

facilities have typically much longer lease durations, with retail and office somewhere in between. Shorter lease durations allow real estate owners to raise rents when the economy improves and short-term interest rates rise. The rise in cash flows helps to offset the rise in discount rates so that prices are less sensitive to the rate increase.

SMB betas and HML betas are large and positive for all sectors. Hotels are most sensitive to HML while self-storage is the most growth-stock like. All sectors load negatively on momentum, about half of the sectors significantly so. Hotel stocks behave most like loser stocks while self-storage REITs behave more like winner stocks. Finally, the illiquidity factor does not help price any of the REIT portfolios.

The six factor model accounts for a substantially higher fraction of variation in sector returns, with the lowest R^2 being 36% (self storage) and the highest 66.6% (hotels). Still, the model misses between 1/3 and 2/3 of all return variation.

The six-factor model allows us to compute a required rate of return as the product of the average factor returns, averaged over the 1994–2014 period, and the exposures to the factors. It adds in a risk-free rate (averaged over 1994–2014). The expected return on the various commercial real estate sectors is: 10.9% for residential, 12.6% for office, 16.3% for industrial, 12.4% for retail, 13.7% for hotel, 11.1% for health care, 11.7% for diversified, and 11.3% for self-storage.

In unreported analysis where we allow the betas to move over time (estimated as before with 120-month rolling windows), we find that in December 2014, the required rate of return is higher still. For equity REITs as a whole it is 14.6%. Across sectors it ranges from 12.6% for self storage and 13.3% for health care to 14.4% for office, 16.0% for retail, and even 22.5% for industrial whose stock and bond betas reach all time highs at the end of the sample.

B.6.3   REIT sector portfolio analysis

The portfolio approach so far assumed that the investor invested in equity REITs in proportion of the market cap of the REITs in the overall index. In reality, investors can over- or underweight various REIT subsectors. Data on REIT sector returns are available for the modern REIT era (1994–2014). We start by forming the tangency portfolio of only REIT sectors (eight sectors). The tangency portfolio has an annualized average return on 20.2% with a volatility of 20.5%, and a Sharpe ratio of 0.85. It goes long residential (40.5%), office (25.1%), retail (20.3%), healthcare (20.2%), and self-storage (106%) and it goes short industrial (−21.6%), hotel (−21.1%), and diversified (−69.1%).

Next, we ask how this portfolio changes once stocks and bonds are included. Stocks receive a weight of 28.5% while bonds have a weight of 48.5%. A long-short portfolios of the 8 REIT sectors combines for the remaining 23%. The positions are long residential (12.7%),

office (4.7%), retail (12.7%), healthcare (6.1%), and self-storage (31.4%) and short industrial (−11.6%), hotel (−8.8%), and diversified (−24.1%). This portfolio achieves an annual return of 11.1% with a volatility of 8.1% and a Sharpe ratio of 1.03. We note that this performance is substantially better than a portfolio that has stocks, bonds, and the equity REIT index, which had a Sharpe ratio of 0.70 over the same period.

Finally, we investigate the sensitivity of this portfolio to average returns. We entertain two scenarios. In the first one, REIT returns all have the same return as stocks over this sample (10.3% per year). The portfolio now features 30.3% stocks, 61.5% bonds, and only 8.2% real estate. Interestingly, the portfolio now shorts office (−2%) and is long diversified (3.2%). the weight on self storage is much reduced (2.9%). This indicates that the long self-storage short diversified nature of the previous real estate portfolio was driven solely by differences in average returns. In the second counterfactual, we assume that average returns on all REIT sectors are given by the 2-factor model. The tangency portfolio is comprised of 32% stock, 68% bonds, and essentially zero real estate (0.16%). All positions in real estate are very small, with the largest position in absolute value being a −1.4% position in the office sector.

B.7   REIT return predictability

U.S. Analysis   Inspired by the present-value relationship (1), we ask whether the dividend-price ratio on equity REITs (cap rate) predicts future returns and/or future dividend growth rates. That is, we estimate regressions of returns and dividend growth rates on the lagged dividend-price ratio, as in equations (4) and (5):

$$r_{t+1} = \bar{r} + \kappa_r(dp_t - \overline{dp}) + \tau^r_{t+1}, \tag{4}$$

$$\Delta d_{t+1} = \bar{d} + \kappa_d(dp_t - \overline{dp}) + \tau^d_{t+1}, \tag{5}$$

We note that the dividend-price ratio is persistent, with annual autocorrelation of 0.74. However, this is not nearly as persistent as for stocks, which historically have had autocorrelations of 0.95 at annual frequency. The lower autocorrelation mitigates statistical problems with the return predictability regressions we present below.

Using all available data (43 years of annual observations of U.S. REITs), we estimate κ_r to be 0.115 when we use nominal returns, 0.055 when we use real returns, and 0.076 when we use returns in excess of the risk-free rate. However, the statistical evidence for return predictability is weak. The point estimate is not significant, and the R^2 is only 4.5% for nominal returns and 1% for excess returns, which is the more relevant object since we are not interested in predicting risk-free interest rates.

Interestingly, we find much stronger evidence for dividend growth predictability by the dp ratio. The coefficient κ_d is -0.114 for nominal and -0.152 for real dividend growth rates. These coefficients are estimated precisely, with t-statistics of -2.9 and -3.8, which is impressive for a relatively short sample (of 42 annual observations). Similarly, the R^2 are much larger at 8.6% and 16.2% for nominal and real growth rates, respectively.

Table 15: **REIT Return and Dividend Growth Predictability**

The table reports the slope coefficient, its standard error, t-statistic and the R^2 of a regression of the return on the NAREIT All Equity REIT index (top panel) and the corresponding dividend growth rate (bottom panel) on the lagged dividend-price ratio of that REIT index. The returns are expressed in nominal terms (row 1), in excess of a one-month T-bill rate (row 2), or in real terms (in excess of CPI inflation, row 3). Dividend growth rates are either in nominal (row 4) or in real terms (row 5). The underlying data are monthly from December 1972 to June 2015, but the regressions are run on annual returns and dividend growth rates (43 observations from 1972–2014).

returns	κ_r	std. err.	t-stat	R^2
1. nominal	0.115	0.088	1.309	4.5
2. excess	0.055	0.093	0.584	1.0
3. real	0.076	0.090	0.845	1.9
div. growth	κ_d	std. err.	t-stat	R^2
4. nominal	-0.114	0.039	-2.929	8.6
5. real	-0.152	0.040	-3.842	16.2

Based on these estimates, we investigate what fraction of the overall variance in the dividend yield (cap rate) reflects movements in discount rates versus movements in future growth rates. We estimate the present-value system and find that 25% the overall variance in dp is accounted for by the variance of the discount rates (the first term in (1)), 25% by the variance in growth rates (second term), and 50% by their covariance.[65] Alternatively, we can decompose the variance of dp into the covariance with future returns minus the covariance with future growth rates. We find that each term accounts for 50% of the variance.

We can use the simple present-value model in equations (4) and (5), the estimates for κ_r and κ_d, and the latest value of the dp ratio (as of December 2014), to predict returns and dividend growth rates for 2015. The model predicts an expected return of 3.3% (nominal), 2.5% (excess), and 1.9% (real) for 2015. This is a very low forecast, due to the ultra-low dividend-price ratio in December 2014, and a very uncertain forecast, given the low R^2 of the return predictability regression. The model implies a high expected dividend growth rate of 12.7% (nominal) and 11.3% (real) for 2015. Since the R^2 is higher, this estimate has a lower

---

[65]Note that the present-value model implies that $\kappa_r - \kappa_d = 1 - \rho\phi$, where ϕ is the autocorrelation of dp (in logs). This restriction is satisfied since $\phi = 0.823$ and $\rho = 0.938$.

standard error around it.

Global Analysis   Next, we study return and dividend growth predictability by the lagged dp ratio. Table 16 contains the results. For the CREI in Panel A, we find strong evidence for both return and dividend growth predictability by the dp ratio. The coefficients are very large and highly significant, despite having only 20 annual observations. The R^2 are also very large: 21% for returns and 33% for dividend growth rates. This is not just an artefact of the particular sample since the return and dividend growth predictability results are much weaker for the broad equity market (not reported). The model forecasts 9% nominal dividend growth on Core Real Estate in 2015 and nominal returns of -1%. The Campbell-Schiller decomposition shows that 59% of variation in the pd ratio comes from variation in discount rates and 41% from variation in cash flows. Like in the U.S. analysis, we find that that cash flows are quite responsive to changes in the valuation ratio for global real estate, more so than for global stocks as a whole.

Table 16: Predictability Core Real Estate and Core Infrastructure

Panel A: Global CREI (1995–2014)					
Returns	κ_r	std. err.	t-stat	R^2	pred2015
nominal	0.56	0.14	3.92	20.55	-0.01
excess	0.59	0.15	4.06	21.94	-0.04
real	0.56	0.15	3.77	21.23	-0.03
Div. gr.	κ_d	std. err.	t-stat	R^2	pred2015
nominal	-0.39	0.19	-2.07	32.81	0.09
real	-0.39	0.20	-2.01	33.67	0.07
	Var(DR)	Var(CF)	Cov(CF,DR)	Cov(DR,dp)	Cov(CF,dp)
CS decomp.	34.7	16.9	48.4	58.9	41.1

Panel B: Global CII (2005–2014)					
Returns	κ_r	std. err.	t-stat	R^2	pred2015
nominal	0.37	0.31	1.19	11.96	0.07
excess	0.45	0.29	1.56	18.72	0.05
real	0.38	0.29	1.32	13.79	0.05
Div. gr.	κ_d	std. err.	t-stat	R^2	pred2015
nominal	-0.49	0.10	-5.02	58.39	0.10
real	-0.48	0.09	-5.28	61.38	0.08
	Var(DR)	Var(CF)	Cov(CF,DR)	Cov(DR,dp)	Cov(CF,dp)
CS decomp.	18.3	32.8	48.9	42.7	57.3

For the CII, we only have 10 annual observations, so results must be interpreted with caution. Nevertheless, the data show very strong predictability of dividend growth rates

on infrastructure by the pd ratio with very large coefficients that are measured precisely. Return predictability coefficients are also very large but measured with too much noise in a small sample to be statistically significant. The model predicts a return on the CII of 7% in 2015 (5% in real terms) and nominal dividend growth of 10%. This suggests that strong returns and especially cash flow growth will continue this year.[66]

---

[66]The Campbell-Schiller decomposition shows that now the majority of price fluctuations in the CII are coming from cash-flow fluctuations (57%). In unreported results, we find similarly strong cash flow growth predictability for CREI over the 2005–2014 period. Not only are dividend growth rates highly predictable, but so are returns, with t-stats around 4.8 and R^2 around 17%. The majority of CREI price-dividend fluctuations over this period are attributable to cash-flow dynamics (54%).

C   Simulation analysis: Portfolio risk

In this appendix, we analyze the total risk of the real estate (or infrastructure) portfolio in order to understand how much of that total risk is systematic versus idiosyncratic.

We consider an equally-weighted portfolio of N individual real estate or infrastructure assets. Each asset has an alpha, a stock beta, and a bond beta. We assume that these three objects are randomly drawn from a normal distribution centered around the point estimates for the SL Green ROA betas reported in Column (1) of Panel B, Table 6. The standard deviations of the two betas are set to 0.05. We center the alpha distribution around 0.05% per month with a standard deviation of 0.25%. This distribution would make the estimated 0.54% SL Green alpha an unlikely but still possible outcome. The two betas and alpha for the N assets are mutually uncorrelated with one another, as well as across assets. Each asset also has idiosyncratic risk e, which is uncorrelated across assets, and has a monthly standard deviation of 5% (17.3% per year). We then create a time series of returns for each asset, using the observed excess factor returns for the 137 months between December 2003 and June 2015.

The resulting assets have average returns that are centered around 7.1% per year but range between -7% and 21% across assets. The systematic component of this return averages to 6.5% per year while the idiosyncratic component averages to the assumed average alpha of 0.6% per year. The bulk of the variation in average returns across assets comes from the idiosyncratic return component rather than from variation in the systematic component. Individual assets' total return volatility is centered around 20% per year and ranges from 17% to 23% per year across assets. The volatility of the systematic component of returns averages 10.1% per year, with a range from 8.3% to 12%, while the average volatility of the idiosyncratic return component averages 17.3% per year, with a range from 14.9% to 19.8%. (The statistics in this paragraph are based on a simulation with 100,000 assets. The reported ranges are based on the 1st and 99th percentiles.)

An equally-weighted portfolio of 50 assets with these characteristics has a mean return of 7.1% and a volatility of 10.4%. The Sharpe ratio is 0.48. These statistics are based on 10,000 simulations with 50 assets. In each simulation, the betas and alphas as well as the panel of idiosyncratic return shocks, $e_{t,j}$, are redrawn. The same observed factor returns are used in each simulation. On average, only 5.5% of the real estate (or infrastructure) portfolio variance comes from the idiosyncratic component. Clearly, a portfolio of 50 assets goes a long way towards diversifying the substantial idiosyncratic risk in individual real estate projects.

However, despite the small fraction of portfolio variance that is idiosyncratic, the portfolio has a large tracking error relative to the benchmark portfolio, R_p^{sys}. The annualized tracking

error from the 50 asset portfolio is 245 basis points per year. The effect on tracking error from non-zero alpha or cross-sectional variation in alpha turn out to be minor. Rather, the tracking error stems from the presence of idiosyncratic risk. To make this point, we vary the idiosyncratic volatility of a typical asset, call it σ_e. If σ_e is zero, tracking error is zero. If σ_e is 2% per month or 6.9% per year, tracking error is 1%. If σ_e is 4% per month or 13.8% per year, tracking error is about 2%. Studying more data points confirms that tracking error rises linearly in σ_e beyond this point.

The way to reduce tracking error is to increase the number of assets. The tracking error of an equally-weighted portfolio would be 0.69% with 100 assets (compared to 1% with 50 assets) if idiosyncratic volatility was 2% per month. Tracking error would be 1.38% with 100 assets (compared to 2% with 50 assets) if idiosyncratic volatility was 4% per month. If we believe the typical idiosyncratic volatility of a real estate or infrastructure project is around 6% per month (almost 21% per year), then tracking error would be 2.1% with a portfolio of 100 assets but 2.9% with a portfolio of 50 assets.

More analysis on individual unlevered stock returns could be done along the lines of our analysis for SL Green to arrive at a more accurate estimate of the empirical cross-sectional distribution of betas and alphas. This would result in a more appropriate range of volatilities for the real estate and infrastructure portfolios of a given number of assets.