HEC MONTRÉAL

Diversification Benefits of Real Estate. A Performance and Portfolio Integration Impact Analysis from a Canadian Pension Fund Perspective.

by

Yann Zaccour

HEC Montréal

Christian Dorion/Martin Boyer

MASTER OF SCIENCE
(Finance specialization)

A thesis submitted in fulfillment of the requirements for the degree of Master of Science in Finance
(M. Sc.)

January 2022

Abstract

This article examines real estate’s role in Canadian pension funds' mixed-asset portfolios using private directly held real estate properties, private equity real estate funds, and publicly traded equity REITs. The goal is to examine diverse real estate-related risk/return opportunities and to assess the impact on portfolio performance of integrating these real estate investment vehicles. The results suggest that Canadian pension funds would be better served by including some form of real estate in their asset allocation than not, based on their targeted risk/return objectives. For the investors preferring low-risk portfolios, the addition of unlevered private real estate properties or core style PERE funds forms the optimal investment vehicles to improve the performance of stocks and investment-grade bonds portfolios. For the investors who prefer high-risk portfolios, public REITs seem to be the right vehicle serving this allocation preference. This thesis also examines the effects of (1) adjusting private real estate data for the appraisal smoothing bias and the liquidity risk; and (2) varying the private real estate data inputs used in the MVO model to analyze how mixed-asset portfolio performance changes under various scenarios and different constraints.

Keywords: Portfolio Management, Real Estate, Diversification, Strategic Allocation

Research methods: Mean-Variance Optimization Framework

Table of Contents

List of Tables

List of Figures

1. Introduction

Diversification is a risk management strategy that mixes a variety of investments vehicles within a portfolio. The rationale behind this technique is that diversification benefits hold only if the chosen asset classes, or the assets within the classes, are not perfectly correlated with each other. A well-diversified portfolio contains a mix of distinct asset classes to limit the impact of unsystematic risk events, so the positive performance of some investments neutralizes the negative performance of others. This strategy allows the well-diversified portfolios to yield, on average, higher long-term returns with less volatility than any individual holding or security.

Since the emergence of research on portfolio optimization using real estate in the 1980s, real estate has the reputation of being an excellent portfolio diversifier by exhibiting lower volatility and very low correlation with the more traditional asset classes, like stocks and bonds. These characteristics, in turn, enable higher risk-return performance for the portfolios that allocate capital to real estate. The benefits of real estate spread across the industry as we witnessed a significant allocation increase to real estate in the Canadian pension funds asset mix. As measured by the RBC Investors & Treasury Services "Canadian Defined Benefit Pension Survey 2021", the allocation to real estate increased from 4% on average in 1999 to 12% in 2019 (see appendix A). However, research on portfolio optimization with real estate in the Canadian market is scarce and incomplete. The increased allocation for this asset class in the recent decades was not sufficiently backed by empirical analysis in Canada and was most likely inferred from research in other developed markets, like the United States. Considering the high heterogeneity of real estate assets, assuming research results from different geographic markets is plausibly less reliable than with more homogeneous asset classes and highlights the need for a local and more in-depth study.

This study aims to verify and analyze the extent of the diversification benefits that different real estate investment vehicles can bring to stocks and investment-grade bonds portfolios in the Canadian market from the perspective of pension funds. Is real estate as beneficial to Canadian pension funds portfolio performance as we commonly think it is? We answer this question by simulating hundred of thousands of portfolios with various asset combinations using data from 2005 through 2021. The tests are performed with and without adjusted private real estate data, which we detail further in the methodology section. These simulated portfolios are then optimized using a mean-variance optimization framework to yield an efficient frontier for each asset combination tested. We compare the resulting efficient frontiers' asset compositions with each other. As measured by the Sharpe ratio, the difference in the portfolio performance of the optimal portfolio of each efficient frontier demonstrates the impact and diversification benefits of integrating certain real estate investment vehicles or a mix of these vehicles. The results are as follows.

First, as a base case, only Canadian investment-grade bonds and Canadian equities are included in the portfolio simulation and optimization model. We call it the control portfolio since it is the only asset class combination without real estate investment. As expected, the results are large allocations to investment-grade bonds for the very conservative portfolios and large allocations to equities for the more aggressive portfolios. The optimal portfolio of this efficient frontier has a 0.91 Sharpe ratio with an 82% allocation to bonds and a 18% allocation to stocks. As the target portfolio return increases, stock allocation gets heavier, and the Sharpe ratio falls.

Then comes the first model with real estate, which we call the unadjusted model. We add the three real estate investment vehicles to the control portfolios in the unadjusted model. In this case, the private direct property investments and the PERE funds are unadjusted for the appraisal smoothing effect and the liquidity risk. Interestingly, in this test, every efficient portfolio lying on the efficient frontier includes some allocation to at least one of the real estate investment vehicles. However, the allocations to real estate are pretty significant and perhaps too significant. The Sharpe ratio of the optimal portfolio in this model has more than doubled to 2.33, with a total summed allocation to real estate of 65%. From this result, the stochastic dominance and the diversification benefits of real estate seem overstated by the appraisal smoothing effect and are neglecting the liquidity risk, which we did not consider yet.

The next model, the adjusted model, presents the same asset class mix as the previous one, but it uses adjusted private real estate data. The adjustments increased the volatility of returns associated with the private real estate investment vehicles and caused a flattening of the efficient frontier compared with the unadjusted model. The optimal portfolio of this efficient frontier has a Sharpe ratio of 1.66 with a total summed allocation to real estate of 43%. So, even after the adjustments, the performance is still remarkably better with real estate than without any. This model is considered more representative of the impact on portfolio performance of integrating all three real estate investment vehicles to Canadian stocks and bonds portfolios. However, the dimensionality problem might overestimate the allocation to real estate, which we further explain in the result section. For reasonable measures, we also analyzed how robust the adjusted efficient frontier results are with respect to changes in the input value of the private real estate investment vehicles. The adjustments made to account for the common private real estate data biases were possibly not strong enough. Even in the strictest scenario tested (the correlation parity scenario with enhanced unsmoothing parameter), the optimal portfolio warrants a 1.13 Sharpe ratio with an allocation to real estate of 24%, demonstrating the resilience of Canadian real estate investments.

The optimization results indicate that Canadian pension funds would be better served by including some form of real estate in their portfolio allocation than not, based on their targeted risk/return objectives. Even after correcting for the risk understatement of private real estate caused by the smoothness of returns and the liquidity risk, real estate remains a desirable feature of a well-diversified portfolio, although in lower allocation than before the adjustments. By looking at how the allocation to real estate varies throughout the portfolio target returns, we observe that for more risk-averse investors, the addition of unlevered private real estate properties or core style PERE funds to stocks and investment-grade bonds portfolios can significantly improve the risk/return performance. On the other hand, for the investors who prefer high-risk portfolios, public REITs seem to be the right vehicle serving this risk-aversion preference as REITS provide more return but are much more volatile than their private counterparts.

This study adds to the existing literature in numerous ways. First, it completes and updates the current research on portfolio optimization with real estate in Canada by integrating three real estate investment vehicles: private directly held real estate properties, private equity real estate funds (e.g., PERE funds), and also, publicly traded equity REITs. The inclusion of PERE funds in an MVO framework in addition to the two other real estate investment vehicles is a first in the literature. Furthermore, since private real estate return data is subject to biases like the appraisal smoothing effect and higher liquidity risk, we adjust the private real estate data so that the results from the analysis account for these biases. These adjustment techniques lean on concepts and theories from Clayton, Geltner, and Hamilton (2001) for the smoothing effect appraisal and Lin and Vandell (2007) for liquidity risk adjustment. The results ought to be more representative of the actual impact on portfolio performance of integrating these real estate investment vehicles to Canadian stocks and bonds portfolios. As for the application and results from the MVO framework with real estate, this paper is closest to Bond, Hwang, and Richards (2006), who calculate the optimal allocation of real estate in a portfolio and account for the liquidity risk in the U.K. market. Also, this paper builds on Hudson-Wilson et al. (2003), Hudson-Wilson et al. (2005), and, Mutahi and, Othieno (2015). They analyze the risk-adjusted performance of REITs and direct real estate investments in the US market using a multi-constraint portfolio optimization approach.

The remainder of this paper is structured as follows: Section 2 reviews the relevant existing literature on the subject, section 3 describes the raw data used as inputs for the base model. Section 4 details the methodology employed to adjust the data, develop the model, and optimize the mixed-asset allocations. We also explain the nature of the robustness tests in section 4. Section 5 presents and interprets the primary results, while section 6 presents the additional and secondary outcomes from the robustness tests. Finally, section 7 concludes.

2. Literature Review

In the past decades, real estate gained much popularity as investors searched for alternative investment vehicles to the point that it successfully established itself as an institutional asset class. Even if real estate and real estate portfolio management have been thoroughly studied throughout the years, there are still many holes in the literature. Real estate assets are heterogenic by nature. Their use tends to change across time, making inference across different geographic markets or time periods less reliable than other, more homogenous asset classes. To properly model real estate portfolio optimization in Canadian pension funds portfolios, we must first dig into the relevant literature to improve our understanding of the asset class and the inefficiencies that it is subject to.

2.1. Direct Real Estate Investments

In this first sub-section, I explore the financial literature related to direct real estate investments. Since the data on the direct real estate portion used in this research is an index valued by an appraisal approach, I'll emphasize parts of the literature that studies the common biases that this method is subject to and explore the different factors that can impact real estate valuations.

2.1.1. Real Estate and Interest Rates

Interest rates and interest rates variations are essential factors that can impact real estate valuations and the demand for real estate products. Generally, when interest rates decline, financing real estate projects becomes more affordable, and the availability of capital for investment increases. Professionals of the real estate industry believe that interest rate variations can also impact real estate valuations by affecting the capitalization rate (cap rates). A capitalization rate is simply the ratio of a property's net operating income (NOI) to its market value. Most investors associate movements in cap rates with changes in asset values. For example, falling cap rates signal rising property values. Intuitively, it seems evident that real estate cap rates and interest rates should be related. The real estate industry is capital intensive and relies heavily on debt, so the cost of debt capital should be important. Conner and Liang (2004) examined the complex interaction between real estate cap rates and interest rates. They state that:

While many factors have contributed to the downward trend in cap rates in recent years, low-interest rates have played a key role in the recent decline. Falling interest rates have enabled real estate investors to pay a premium for assets by using leverage to enhance their returns and have made real estate yields relatively attractive compared with other asset classes.

2.1.2. Real Estate Risk Premium and the Liquidity Risk

Directly owning real estate properties is considered by many to be a safer form of investment than owning stocks because private real estate markets are typically less volatile. Even though the interest of this paper is on commercial real estate, some interesting insights can be drawn from extensive studies on housing. It is the case with Jordà et al. (2019), who computed annual returns for equity, housing, bonds, and bills in 16 advanced economies from 1870 to 2015. Their main finding on risky returns, including equity and housing, states that housing returns are similar to equity returns but much less volatile. Like Shiller (2000), they found that the long-term capital growth return is low (around 1% per annum in real terms) and significantly lower than capital gains in the stock market. On the other hand, the income return component of real estate is typically considerably higher and more stable than the dividend yield of equities. Total returns are of comparable magnitude, and the high and steady income partially explains the lower volatility of total returns. Using the same database, Jordà, Schularick, and Taylor (2019) attempted to explain the puzzling housing and equity risk premium. By adding housing to the representative-agent model, they found that the total risk premium puzzle gets even more significant and that standard and even more exotic models cannot resolve these giant puzzles. As a result, an uncharacteristically large, implied risk aversion parameter is needed to balance the equation, meaning that for housing wealth and total wealth, the risk aversion would be twice as significant as for the equities taken alone, which doesn't seem rational.

As for the low volatility measures of real estate, many researchers state that the risk associated with real estate investing is often underestimated. The early empirical studies in the commercial real estate literature comparing risk and return with other asset classes concluded that real estate had extremely high risk-adjusted returns. It's hypothesized that real estate returns could be less volatile because real estate indices are estimated by using both appraisal data with transaction prices. Also, appraisal valuations tend to smooth values and make them appear more stable than they are. Geltner and Goetzmann (2000) assembled a transaction-based NCREIF index to validate this hypothesis. They find that the volatility, measured by the standard deviation of returns, increased from 3.66% to 4.26% in the 1978 to 1998 period, marking an increase in real estate risk. However, the resulting standard deviation is still very low compared to traditional asset classes.

Since the risk-premium puzzle in real estate could not be entirely explained by the appraisal smoothing, Lin and Vandell (2007) made a deep dive into the literature surrounding the liquidity bias associated with real estate. They determined that the most critical aspects defining real estate illiquidity in residential and commercial markets are the time required for sale and the uncertainty of the marketing period. Lin and Vandell (2007) formally develop a framework that defines and measures *marketing period risk* by proposing a new measure of *ex-ante* return and variance, which replaces the transaction-based return and variance measure to capture price risk and one component of liquidity risk. Their results demonstrate that the risks associated with *ex-ante* returns are significantly higher than calculated with the standard return and variance methodology. For example, they find for the commercial real estate market with an expected marketing period of eight months, from a long-term investor perspective, that the effective risk is 30% higher than what is assumed under more standard calculation methodologies. Therefore, “traditional real estate valuation methodologies, which assume that real estate can be sold immediately and ignore the uncertainty of marketing period, can seriously underestimate real estate risk.” (Lin and Vandell, 2007, p. 321).

Up next, Cheng, Lin, and Liu (2013) developed a concept called the liquidity risk factor (LRF), which provides an analytical tool for quantifying the liquidity risk of private assets. Using commercial real estate data as a testing ground, they find that under the independent and identically distributed condition, and for typical institutional investors holding assets for 5 to 8 years, liquidity risk is of such magnitude that the conventional return volatility should be adjusted upwards by 21% to 75% in poor market conditions (standard selling time of 12 to 16 months), by 8% to 26% in normal market conditions (standard selling time of 6 to 9 months) and by less than 10% in good market conditions (standard selling time of 2 to 5 months). While such adjustment varies over market conditions and investment horizons, suffice to say that liquidity risk should not be ignored by rational investors, especially in down markets when liquidity is even more scarce.

2.1.3. Appraisal Smoothing in Real Estate Indices

While historical data on bonds or equity returns is readily available, it's an entirely different story for privately-owned direct real estate assets. Since transaction volume and frequency are much lower than bonds and equities, most indices measuring real estate return over time are appraisal-based. Therefore, these indices are not based on actual market values but instead on the best estimates of market values.

Appraisal-based indices are widely known by researchers and professionals in the industry to "smooth" the growth of real estate value over time. Hence the expression "appraisal smoothing." The impact of this smoothing takes the form of an underestimation of the volatility of the underlying assets and an understatement of the correlation with other asset classes. This bias comes from the nature of the appraisal valuation process itself. The appraiser starts by looking at the last appraisal, which he adjusts with updated market data and comparable sales (Ibbotson and Siegel, 1984). The appraiser uses recent market data as much as possible, but it is still backward-looking, and comparable building data might be on properties with slightly different attributes. The results are that the appraisal of a property's value tends to demonstrate autocorrelation with the past assessment since appraisers cannot be sure of the accuracy of the inputs used to value the property. Quan and Quigley (1991) develop a model of appraiser's behavior on valuations that can be interpreted as "a weighted average of a previous appraisal, as proposed by previous researchers, as well as the most recently observed transaction price." The key to the model is the updating rule that the appraiser employs to extract the price signal from the "noisy" transactions made by imperfectly informed actors in the market.

To "unsmooth" the returns, a parameter α of confidence appraisers have in the current market information they observe must be specified. Firstenberg, Ross, and Zisler (1988) use a regression model to estimate the serial relationship between the present appraisal and the previous one. The inconvenience of this method is that the parameter α cannot differentiate the momentum caused by the smoothing effect from the natural momentum of the asset class. Geltner (1993) solves the momentum problem by developing a model that corrects the lagging and smoothing problems while allowing any true momentum in the return series to remain. The issue with this new method is that an assumption about the volatility of real estate returns is necessary. However, Clayton, Geltner, and Hamilton (2001) successfully estimate a reasonable level for the parameter α by regressing the purely contemporaneous transaction-based market value of properties, with the appraised value at time t-1 and a constant, on the appraised values at time t. The parameters estimated from the regression are the weights α and (α - 1), which represent the confidence in the market information at time t and residual confidence in past market data from t-1, respectively. The results are that appraisers, on average, put an 81.5% weight on current market information and (1 - 0.815) = 18.5% weight on past appraisals.

2.2. Indirect Equity Real Estate Investments Vehicles

While direct real estate investment involves directly acquiring properties, indirect real estate equity investment often involves buying shares of companies that invest in real estate. This type of property investment includes shares, private or public funds, and derivatives. In the following sub-section, I explore the literature on the two most popular indirect real estate investment vehicles for Canadian pension funds, first with public REITs (Real Estate Investment Trusts) and then with PERE funds (Private Equity Real Estate).

2.2.1. Link Between Private and Public Real Estate

A real estate investment trust (REIT) is a company that owns, operates, or finances income-generating real estate. REITs pool the capital of multiple investors, offering an easy way for investors to gain exposure to the real estate sector without having to buy, manage, or finance any properties themselves. Stefek and Suryanarayanan (2012) studied the link between private and public real estate. They find that REITs and private real estate offer surprisingly different risk profiles. Private real estate appears much less volatile than its public counterpart even after accounting for the difference in leverage. Also, private and public real estate returns show little correlation over quarterly and annual horizons. If anything, public real estate returns seem to lead private market returns. But after correcting for appraisal smoothing and the lead-lag relationship between public and private values, they demonstrate a strong link between private and public real estate returns. The correlation between private and public returns strengthens as the investment horizon increases.

Stefek and Suryanarayanan believe that a different clientele could explain why public and private real estate markets may behave differently in the short run. Large or institutional investors with longer time horizons are more likely to invest in private real estate. In contrast, REITs are accessible to a broader range of investors, many with shorter time horizons and greater liquidity needs. Also, REITs seem to be more closely related to the broad equity market than private real estate.

2.2.2. REITs and Liquidity

Arguably the most important advantage of public equity REITs versus private real estate is the liquidity of the investment vehicle. REITs benefit from similar levels of liquidity as stocks do, permitting investors to get in and out of positions almost instantly at much lower transaction costs than directly buying or selling properties. Also, this high liquidity makes it easy to rebalance the allocation of a portfolio. Ciochetti, Craft, and Shilling (2002) researched why some institutional investors invest in real estate investment trusts stocks and others invest directly in private real estate. They find evidence that liquidity-constrained institutional investors strongly prefer liquid REIT shares compared with private real estate. Some institutions are simply too small to devote the significant resources necessary to invest directly in private real estate. Their results also show that institutional investors prefer larger and more liquid REIT stocks.

Clayton and McKinnon (2002) examined the short-run relationship between REIT prices and the value of direct real estate owned by REITs (i.e., Net Asset Value (NAV)) by developing a model in which fluctuations in the average REIT sector price premium to NAV is a function of time variation in REIT growth opportunities, the value investors place on REIT liquidity and sentiment-based trading by retail investors. They find a significant liquidity premium in REIT prices relative to property NAV that varies systematically with the liquidity of the private real estate sector. The liquidity benefit of REITs is valued more in a down private market than in an up one where NAVs are rising and investors are more focused on growth. All else equal, when NAVs are low, investors are more concerned with liquidity in the private market and place a higher value on public market liquidity.

2.2.3. REITs and Interest Rates

Interest rate is one of the most closely followed macroeconomic variables by REITs investors because of its impact on performance. Many practicians keep a close eye on the 10-year government yield and its spread with cap rates to evaluate performance and risk appetite. Chmiel and Rodriguez (2019) found a better predictor of REIT return than the 10-year government yield by introducing the spread between BBB-rated bonds and the 10-year GoC. They believe that the Canadian-listed property sector is closer in risk profile to a BBB rating. When the spread widens between 10-year GoC and BBB-rated bonds, investors have less appetite for risk and tend to move away from REITs, with the inverse relation being also true. They find very strong co-movement by plotting the TSX capped REIT and the inverse of the spread with an R^2 of 0.8. Ling, Webb, and Myer (2003) had previously found a similar relationship in the U.S. REIT market from 1972 through 1998. Their overall OLS result suggests that equity REITs are significantly affected by changes in yields on long-term U.S. government bonds and high-yield corporate bonds. Results indicate that the changes in yields on high-yield corporate bonds (Baa) have the strongest explanatory power for equity REIT returns.

2.2.4. PERE Funds

The importance of private equity real estate (“PERE”) funds grew dramatically in recent years, and this change is particularly evident amongst the largest Canadian pension funds, most of which have substantial exposure to PERE funds (Kuzmicki and Simunac 2008). According to the information manager Preqin, the aggregate PERE fund capitalization grew from $101 billion at the end of 2003 to $992 billion as of year-end 2019 (Kieran and Stevenson 2019). The U.S. is the largest single investable market.

PERE funds, as explained by Kuzmicki and Simunac (2008), are, as the name implies, a hybrid between “private equity” and “real estate” asset classes. While traditional private equity consists of direct ownership by investors, private equity funds relate to indirect ownership by investors through a third-party fund manager. In most cases, the investors themselves are private entities (i.e., pension funds). Strategies of PERE funds are diverse and unique for each fund, allowing their investors to expose their portfolios to specific geographic regions, real estate sectors, and strategies without dedicating the necessary resources to do it themselves. One of the strengths PERE funds have over other investment vehicles is their malleability and the ability of fund managers to create new funds in response to emerging opportunities. Each fund is also classified in terms of its risk and return profile from “core” funds which are well diversified and aim to invest in stabilized income-producing properties, to “opportunistic” funds that employ more risky strategies like re-positioning obsolete assets and new-build development projects. See Appendix C for a complete breakdown of the PERE funds classification and risk-return profiles.

2.2.5. PERE Funds Risk and Return Characteristics

There was little to no research on PERE funds in Canada or elsewhere until recently. It might reflect a lack of available data because, as privately held entities, PERE funds are not required to publicly disclose details regarding their activities or financial performance. PERE fund data is also subject to similar biases as direct real estate regarding appraisal smoothing and illiquidity, which we already covered previously.

Tomperi (2010) studies the realized returns, measured by the internal rate of return (IRR) and its relationship with the PERE fund size and sequence number (consecutive funds by the same manager). The analyses are performed on a large global sample of value-added and opportunistic private real estate funds. Different model specifications are used to study the fund and sponsor-related factors’ correlation with fund performance. He finds evidence that performance increases with fund size but declines with the sequence number. In addition to the fund-specific factors, funds that raised money during slow economic activity are likely to have better performance. The lower the GDP growth at the time of the fund launch, the higher the realized IRR, mainly because of the strong growth that typically occurs after an economic downturn. Arnold, Ling, and Naranjo (2019) confirm these results and add to the existent literature by investigating the factors that drive the performance of closed-end PERE funds. Using performance data through the end of 2017 on 467 PERE funds that came to market between 2000 and 2013, they regress performance metrics to fund-level characteristics, market risks, and macroeconomic variables. They found evidence that fund characteristics, market risks, and macroeconomic risk factors significantly predict PERE funds' performance. Performance is positively related to fund size, GDP growth, returns in the private real estate market, long-term interest rate changes, and changes in the credit risk spread.

Arnold, Ling, and Naranjo (2021) compare the performance of PERE funds to REITs. Their dataset contains 950 distinct PERE funds from 2000 through 2019. Relative to equity REITs, the typical PERE fund is significantly less liquid, uses more financial leverage, and requires investors to maintain liquid assets for capital calls of indeterminate size and timing, so they arbitrarily adjust the performance metrics of PERE funds to be comparable with REITs. Overall, they find that PERE funds have underperformed listed REITs, even before adjusting for risk, leverage, illiquidity, and the uncertain investment timing associated with unfunded capital commitments.

2.3. Real Estate Allocation in Optimized Portfolios

In the United States, the optimal portfolio allocation to real estate has been studied on numerous occasions, mainly since the 1980s. Firstenberg, Ross, and Zisler (1988) analyzed this matter to show how pension funds and other large investors could use modern portfolio theory to allocate funds to real estate. To conduct their research, an appraisal-based index is used in conjunction with stocks and bonds to construct an efficient frontier. Using data from 1978 to 1985, they conclude that, from a portfolio perspective, the attractive feature of real estate is its lack of correlation with other assets classes, even with the appraisal-based valuation process understating the risks associated with real estate. The lack of correlation makes real estate a desirable feature of a well-diversified portfolio, and they suggested that pension funds should allocate between 15 to 20% of their portfolios to real estate assets, which was significantly higher than the 3.6% allocation for the top 200 public and private funds in 1986. Later studies by Hudson-Wilson et al. (2003) also suggest that the optimal real estate allocation is between 15 to 20%. This latest study was revisited by the same authors and more in 2005 (Hudson-Wilson et al. (2005)). In order to confirm their previous findings, they create a capitalization-weighted index of the U.S. real estate investment universe, which includes public and private real estate equity and also public and private real estate debt vehicles. To calculate the optimal allocation to real estate in a mixed-asset class portfolio, they include the Lehman Corporate/Government bond index, the S&P 500, and the Treasury bill rate as a proxy for the other asset classes. Using quarterly returns from 1987 through year-end 2004, they find that, as before, the correlations between real estate and stocks, real estate and bonds, and real estate and cash suggest that real estate can play a significant role in a mixed-asset portfolio. In another study overseas this time, Bond, Hwang, and Richards (2006) calculate the optimal allocation to real estate in a portfolio accounting for the liquidity risk associated with the uncertain marketing period using U.K. asset return data. Their results indicate that after incorporating the liquidity risk, the allocations to real estate in a portfolio with a holding period of one-year fall. However, for a more extended holding period of five years, the impact of the liquidity risk on portfolio allocation is less significant. It cannot entirely explain the large discrepancy between observed portfolio allocations to real estate and the allocations suggested by the mean-variance optimization models.

In a more recent study, as an attempt to determine which real estate investment vehicle should be included in institutional investor portfolios, Mutahi and, Othieno (2015) analyze the risk-adjusted performance of REITs and direct real estate investments. Using U.S. data between 1980 and 2014 with a multi-constraint portfolio optimization approach, they find that direct real estate investments outperform REITs, but REITs outperform direct real estate when incorporating a minimum return constraint. They attribute these results to the higher risk-return characteristic of REITs. However, this study has a significant limitation. It assumes that investors are only willing to invest in one of the two investment vehicles. Recently, Pagliari Jr. (2017) examined the role of real estate in institutional mixed-asset portfolios using both private and public real estate indices. He used annual return data from the NCREIF Index, the NAREIT Index, and multiple stocks, bonds, and bills indices from 1978 to 2013 to construct and optimize portfolio allocations with different investment horizons and different maximal allocation constraints and varying risk preferences. The empirical results from all the efficient frontier tests suggest that investors preferring low-risk portfolios are better off with a heavier allocation on private (unlevered) real estate rather than public REITs. For those investors choosing high-risk portfolios, public real estate is the vehicle serving this allocation preference (with their embedded leverage of 40–50%). It also seems that one’s risk preference translates to leverage preference. Their overall allocation recommendation to real estate is approximately 10–15% of the mixed-asset portfolio. However, using one-year returns as the investment horizon yields an average allocation to real estate of 23.5%, whereas using annualized four years returns only yields an average real estate allocation of 11.5%. As the investment horizons lengthen, the correlation between NAREIT and NCREIF returns and other asset classes increases and, not surprisingly, suggests that diversification benefits may be less helpful to reduce portfolio volatility.

In Canada, real estate allocation in optimized portfolios remains an understudied topic. Vachini (2011) investigated the effects of adding REITs and real estate mutual funds to Canadian equity and fixed income portfolios. Using data from 2010 to 2018, he concludes that diversifying a portfolio with real estate can significantly improve performance. His results suggest allocating 25% to 34% of the portfolio's assets to the Bloomberg Canadian REIT Index and none to the S&P TSX REIT Index for an optimal risk-return relationship.

From another perspective, the negative skewness and high kurtosis of real estate, more specifically of U.S. REITs in the past decades, could explain the institutional investor's historical under-allocation to real estate. Xiong and Idzorek (2011) factored in the third and fourth moments of the return distribution of multiple asset classes from 1990 to 2010. They find that global high yield, U.S. REITs, U.S. TIPS, and value stocks have significant negative skewness and high kurtosis, resulting in a lower allocation than a standard mean-variance optimization model would typically suggest.

2.4. Tax Advantages of Canadian Pension Funds as Investors in Real Estate

From a fiscal standpoint, Canadian pension funds have considerable advantages regarding real estate investments. MacNevin (2013) analyses the competitive situation of pension funds vis-à-vis conventional taxable investors investing in real estate. To start, overall contributions to a pension fund are tax-deductible, and investment earnings accumulate in the fund tax-free. Taxes are paid when withdrawals from the plan occur. Both principal and investment returns are fully taxable at the taxpayer's prevailing marginal tax rate at the time of the withdrawal, which means that the fund itself is not taxable and serves as a deferral vehicle for the personal tax that will be paid ultimately by the beneficiaries.

In Canada, pension funds have many options for investing directly in real estate through tax flow-through investment vehicles. The most popular vehicle is by far the REIC (real estate investment corporation) described in paragraph 149(1) (o.2) of the Act. A pension fund may use a REIC to invest in real estate, either alone or with other participating funds, with complete corporate liability protection.

In the case of an indirect real estate investment, the REITs that meet income payout rules can flow earnings tax-free to the pension fund and provide a relatively high degree of liability protection through the declaration of trust. Tax effects arise only when earnings are withdrawn from the pension fund. In other words, pension funds have a significant advantage over taxable investors, which translates into higher rates of return.

3. Data

In order to conduct this research, data from five different datasets, spanning from 2005 through 2021, are assembled. Three of those represent the three prominent real estate investment vehicles used by pension funds in Canada. The first is the MSCI/REALPAC Canada Property Index published quarterly as the proxy for a diversified portfolio of direct property investments. The second is the MSCI/REALPAC Canada Quarterly Property Fund Index, representing investments made in PERE funds. The third is the iShares S&P/TSX Capped REIT Index ETF (XRE) for the Canadian REIT exposure. For the other asset classes, Canadian stocks, and bonds, we selected the iShares S&P/TSX 60 Index ETF (XIU) and the iShares Core Canadian Universe Bond Index ETF (XBB). All of these are measured or traded in Canadian dollars.

This section presents the two indices and the three ETFs, followed by a statistical analysis of their serial returns.

3.1. Data Sources

MSCI/REALPAC Canada Quarterly Property Index

The Canadian Property Index has quarterly data from the fourth quarter of 1999 to September 30th, 2021. As of September 2021, the Property Index tracked the unlevered performance of 2 401 Canadian commercial property investments, with a total capital value of CAD 167.5 billion, putting the average property value at CAD 69.8 million. MSCI estimates that the Canadian commercial real estate total market size is CAD 434 billion. So, the index market coverage ratio is pretty high at 39%. This index is the primary performance benchmark used by real estate portfolio managers in Canada. Although the index itself is not tradable, large institutional investors can track its performance by reproducing its geographical and sectorial allocation. As 80% of the index is allocated in four cities (Toronto, Vancouver, Montréal, and Calgary) and 97% in the four main asset types (retail, industrial, office, and multi-residential).

The main objective of the MSCI Canadian Property Index is to measure the performance of direct private commercial real estate investments accurately and objectively in the Canadian market. The constituents of the MSCI Property Index are real estate investments held in professionally managed portfolios. They, therefore, may include properties held in insurance and pension funds, listed property companies, unlisted pooled funds, charitable trusts, and other large private property owners. Property data are generally provided to MSCI by or on behalf of the managers of the real estate investment portfolios concerned. On occasion, MSCI supplements this data with public or third-party information sources, such as published financial reports.

The MSCI Property Index is predominantly based on professionally sourced market valuations carried out by independent valuers and, when possible, on property transacted prices. MSCI only uses asset valuations that are theoretically achievable estimates of actual market transaction prices to set a standard. Therefore, it provides a clear and precise definition of the valuation that investors and managers should provide. At the aggregate level, market valuations are unlikely to differ systematically from transacted prices.

The most widely used index performance measure to evaluate the investment performance of commercial real estate is the total return and its income and capital components, i.e., the total return (TR), capital growth (CG), and income return (IR). MSCI calculates these measures monthly and compounds them (time-weights/chain-links) over quarterly periods. Each month's measures are value-weighted, meaning that the contribution of each asset is in proportion to its monetary weight.

The direct real estate performance measure for this study is the total return. It is the most widely recognized performance figure and the most critical measure of overall investment performance used to compare different assets across time periods. It incorporates both capital and income elements and is calculated as the percentage value change plus net income accrual relative to the capital employed. It is recognized by GIPS (the Global Investment Performance Standard set out by the Chartered Financial Analyst Institute) as the standard composite measure of investment performance. With respect to a single month *t*, total return is defined as:

$$TR_t = \left(\frac{CV_t - CV_{t-1} - CExp_t + CRpt_t + NI_t}{CV_{t-1} + CExp_t}\right) \times 100$$

Where:

- TR_t is the total return in month *t*;
- CV_t is the capital value at the end of month *t*;
- $CExp_t$ is the total capital expenditure including purchases and developments in month *t*;
- $CRpt_t$ are the total capital receipts reflecting changes in the owner's interest in a property and other payments such as surrender premiums that can be capitalized in month *t*;
- NI_t is the rent receivable during the month *t* net of property management costs, ground rent, and other irrecoverable expenditures.

MSCI/REALPAC Canada Quarterly Property Fund Index

Contrarily to the Property Index (direct real estate), which measures the performance of aggregates of individual properties held within investment portfolios, the Property Fund Index (PFI) measures the performance of PERE fund vehicles in their entirety. The index data is available quarterly from 2005 to September 30th, 2021. It tracks the performance of 9 mainly core Canadian PERE funds totaling 1 035 assets with a Gross Asset Value (GAV) of CAD 42.2 billion as of September 2021. Like the Property Index, the PFI is not tradable but can still be tracked by investing in a few PERE funds with a core risk profile that are similarly allocated in the primary Canadian cities (Toronto, Vancouver, Montréal, and Calgary) and the four main real estate asset types (retail, industrial, office, and multi-residential).

The objective of the MSCI Property Fund Index is to measure the performance of unlisted pooled structures, including the effects of cash holdings, leverage, and fund operating costs, fees, as well as returns to the underlying real estate assets. To achieve this objective, the index is constructed top-down from the financial records of real estate investment funds. To effectively represent the performance of the market, MSCI uses fund-level Net Asset Values (NAVs) and distributions to the fund's investors. The asset values used by the fund to calculate NAVs are primarily from property valuations (same appraisal process as with the Property Index). It is impossible to create a representative index using only transactions, as unlisted property funds' investments are illiquid, and transactions are infrequent.

The performance calculations of the PFI apply to individual funds and the index. The chosen measure of performance of the PFI for this research is the Net Total Return with the Modified Dietz formula. This measure considers the impact of leverage, cash, fund costs, and fund fees on the index's total return.

Starting from the gross total return for a single month *t*, the Modified Dietz formula defines total return (gross of fees) as:

$$Gross\ Fund\ Return_t = \left(\frac{Appreciation_t + NIY_t}{AvWtdEq_t}\right) \times 100$$

Where:

$$AvWtdEq_t = NAV_{t-1} + \sum_{i=1}^{n} \left[Days_{i,t} \times NCI_{i,t}\right]$$

and:

- $Appreciation_t$ is defined as the net appreciation/depreciation of all assets and interests, both realized and unrealized, not caused by capital expenditure. It includes properties, mark-to-market debt, and any other investments or liabilities in the month *t*;
- NIY_t is the net investment income in month *t*;
- $AvWtdEq_t$ is weighted average equity at month *t*;
- NAV_{t-1} is the net asset value of the underlying real estate assets in month *t-1*;
- $Days_{i,t}$ is the number of days as a portion of the number of days in the month *t*;
- $NCI_{i,t}$ is the net capital invested in month *t*.

We are interested in the index's net total return (net of fees) for this study. The issue with this measure is that MSCI only started publishing it in 2015, whereas the gross total return with the modified Dietz formula has been available since 2005. To solve this issue and obtain net total returns from 2005 to 2021, we computed the average monthly fund fee of 6.3 bps. We inserted it in the Net Total Return with the Modified Dietz formula from 2005 through 2014 (the result is a simple deduction of fee added to the gross total return formula). MSCI calculates the net total return for a single month as:

$$Net\ Fund\ Return_t = \left(\frac{Appreciation_t + NIY_t - Fees_t}{AvWtdEq_t}\right) \times 100$$

Where $Fees_t$ represents the management fees and any incentive fee in month *t* for the period of 2015 to 2021, or the average monthly fee for the 2005 to 2014 period.

Blackrock iShares S&P/TSX Capped REIT Index ETF (XRE)

This study's third real estate investment vehicle is publicly traded equity REITs. To gain exposure to the Canadian REIT market, the XRE ETF is selected. XRE aims to replicate the performance of the S&P/TSX Capped REIT Index, net of expenses since October 2002. With 20 different REITs in the fund, as of the end of September 2021, the ETF is exposed to different types of REITs, such as retail, residential, office, industrial and diversified. To obtain the quarterly returns necessary to use as inputs in the optimization model, monthly total (NAV) returns since October 2002 were downloaded directly from Blackrock's website. These monthly returns account for dividend or distributions reinvestments and are net of fund fees and expenses. Quarterly return concerning a single quarter is calculated as follows:

$$Quarterly\ Returns_t = ([(1 + r_1) \times (1 + r_2) \times (1 + r_3)] - 1) \times 100$$

Where r_t represents the monthly total (NAV) return in month *t*.

Blackrock iShares S&P/TSX 60 Index ETF (XIU)

To gain exposure to the equity market of large and established Canadian companies, we chose the iShares S&P/TSX 60 Index ETF (XIU). The XIU ETF started trading in 1990, making it the world's first ETF. In Canada, XIU is one of the largest and most liquid ETFs as it seeks to replicate the performance of the S&P/TSX 60 Index, net of expenses. As with the previous ETF covered, XRE, monthly total (NAV) returns were downloaded directly from Blackrock's website from 1999 to September 30th, 2021, and the same quarterly returns calculation method is applied.

Blackrock iShares Core Canadian Universe Bond Index ETF (XBB)

The chosen representative of the fixed income asset class is the XBB ETF. Launched at the end of the year 2000, XBB seeks to provide monthly income to investors by replicating the performance of the FTSE Canada Universe Bond Index™, net of expenses. With more than 1 400 different securities in the fund as of September 2021, the ETF provides diversification and broad exposure to the Canadian investment-grade bond market. As with XIU and XRE, monthly total (NAV) returns were downloaded directly from Blackrock's website from November 2000 to September 30th, 2021. The same methodology for computing quarterly returns as with XRE and XIU is employed.

Treasury bills – 3 months

We calculate the risk-free rate, which we use to compute the Sharpe ratios, using the annualized 3 months Canadian Treasury bill yield. The data was downloaded directly from the Bank of Canada's website from January 2005 to September 30th, 2021.

3.2. Statistical Analysis of Unadjusted Data

Since the underlying assets of the PFI and the REIT's individual securities in XRE are Canadian commercial real estate properties, this analysis starts by looking at the Property Index to determine how direct real estate assets have performed over the last two decades. The total return of the Property Index is decomposed into two components, income return, and capital growth. Figure 1 shows the total return and the split between income return and capital growth, calculated by MSCI, since the beginning of the index, December 31st, 1999.

The index has been on a clear and steady uptrend for over 20 years, with two exceptions, the Great Financial Crisis of 2007-2008 and the COVID-19 pandemic in early 2020, which caused mild drawdowns in the index trajectory. An important takeaway from these return series is that direct real estate properties are primarily income-driven investments. The income return component has had very steady growth, which significantly contributed to lowering the overall volatility of the total return index. However, appropriate decisions on market timing and property selection by individual real estate portfolio managers could lead to higher capital growth returns. This last strategy, however, is not explored in this study.

It's worth noting that there is a reinvestment risk associated with the income return component. Since the income cannot be used to buy more shares like with an ETF, respecting the reinvestment hypothesis can become delicate. On an individual property level, income can be reinvested as capital improvements, maintenance, or redevelopment, but assuming that the income is consistently reinvested throughout the life of individual assets is unrealistic. At the aggregate level, though, or for a large portfolio, reinvesting the income into acquiring, developing, or improving real estate properties is plausible.

<u>Figure 1 - Cumulative Return of Direct Canadian Real Estate from 1999 to 2021</u>

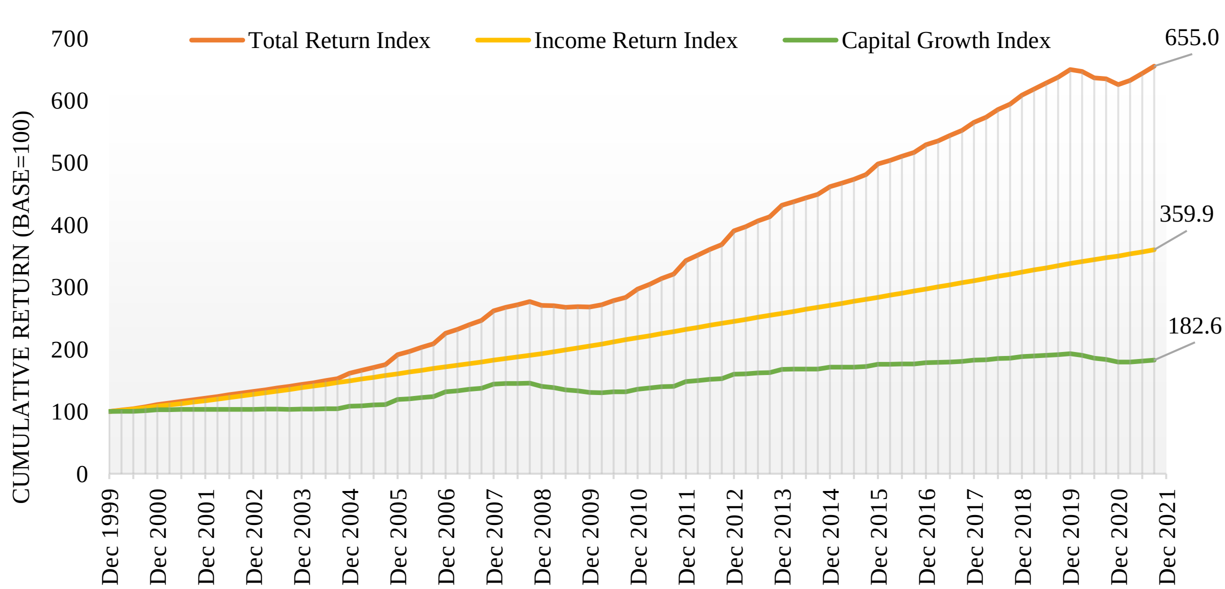

Figure 1 plots the cumulative total return, income return, and capital growth of the Property Index from December 31st, 1999, to September 30th, 2021. The calculation assumes reinvestment of the income return and uses quarterly returns. These series are indexed to a starting value of 100.

Figure 2 displays the value of a hypothetical CAD 100 investment made at the end of 2004 in the three Canadian real estate investment vehicles and Canadian stocks and bonds. For all five series, we use the total returns to calculate the growth of the investments over the 67 quarters.

We can see from figure 2 that the REIT ETF, XRE, outperformed the other four investments, closely followed by the Property Index. The XIU ETF, representing the Canadian stock market, surprisingly underperformed the Property Index and outperformed the PFI by a very small margin. Over the analyzed period, all investments significantly generated more wealth than the investment-grade Canadian bond ETF, XBB. Note that the difference in leverage is not accounted for in the return series above. The Property Index is unlevered, whereas the PFI has an average loan-to-value ratio of around 20% on the underlying properties. Also, the XRE and XIU ETFs are unlevered at the fund level. However, the underlying securities of the ETFs are impacted by the debt used at the company level. Since this study aims to measure the integration impact of either of these investment vehicles, as they are, in a pension fund portfolio, we do not adjust the returns to account for the difference in leverage. The reasoning behind this decision is that when adding XRE, XIU, or PERE funds (PFI) to a portfolio, these securities come with their implicit level of leverage. In contrast, when directly acquiring real estate properties, the use of debt is a discretionary choice by the investment manager.

<u>Figure 2 - Cumulative Total Returns from 2005 to 2021</u>

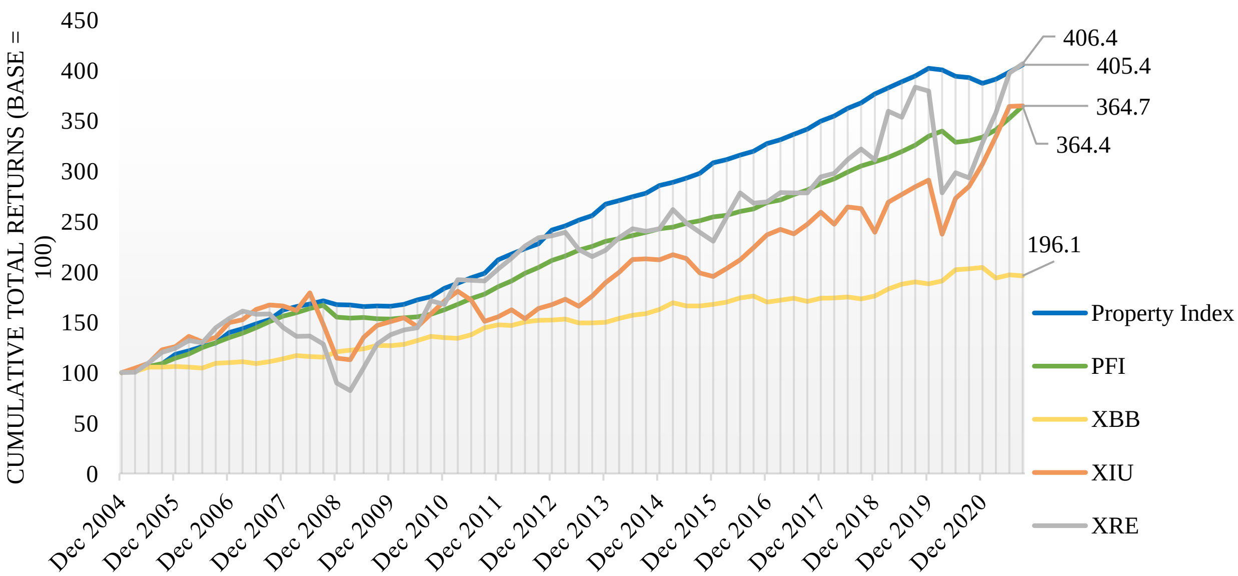

Figure 2 plots the hypothetical growth of a CAD 100 investment in each of our five series from December 31st, 2004, to September 30th, 2021. The growth of the Property Index and the PFI investments are calculated using unadjusted total returns data.

For the purpose of strategic asset allocation, the investment manager should also consider other long-term investment characteristics, like risk, along with the return. Table 1 below showcases descriptive statistics for the five datasets from 2005 to September 30th, 2021, totaling 67 quarterly returns each.

From the table, we can see that XRE has the highest quarterly return rate averaging 2.51%. It also seems that the REIT ETF is the riskiest investment vehicle, with a quarterly standard deviation of 8.79%. The stock ETF comes in second in terms of risk and return, with an average quarterly return of 2.23% and a quarterly volatility measure of 7.35%. The higher risk for the REIT ETF versus XIU could potentially be explained by the fact that XIU has 60 securities in the fund versus only 20 for XRE. Consistent with previous studies, private real estate delivers relatively high quarterly returns coupled with significantly lower volatility than its public counterpart and than Canadian stocks. The Property Index and the PFI had average quarterly returns of 2.13% and 1.96% over the 17 years analyzed, with a standard deviation of only 1.96% and 1.80%, respectively. Let's keep in mind that we calculated these results using private real estate data unadjusted from the appraisal smoothing bias and the liquidity risk, which are known to underestimate volatility measures, as seen in the literature. Private real estate data adjusted for these two issues are presented further in this paper, and the methods used to adjust the data are also detailed in the following section. Nonetheless, with raw data, private real estate offers almost double the return of the Canadian investment-grade bond universe ETF with approximately the same level of risk despite both being income-driven investments vehicles.

Table 1 - Quarterly Descriptive Statistics

	Property Index	PFI	XBB	XIU	XRE
Observations	67	67	67	67	67
Mean	2.13%	1.96%	1.03%	2.23%	2.51%
Std-dev.	1.96%	1.80%	1.99%	7.35%	8.79%
Min	-2.26%	-6.92%	-5.09%	-22.59%	-30.22%
0.25	1.37%	1.35%	-0.13%	-1.15%	-1.64%
0.50	1.84%	2.11%	1.07%	2.93%	2.38%
0.75	2.73%	3.19%	2.32%	6.47%	6.79%
Max	9.01%	5.23%	5.89%	19.94%	27.68%

Table 1 presents summary statistics of the quarterly returns for the five selected asset classes from December 31st, 2004, to September 30th, 2021. The Property Index and the PFI are unadjusted for the appraisal smoothing and the liquidity risk.

Perhaps of more interest to institutional investors than the stand-alone performance of real estate is its role within an optimal strategic asset allocation. Indeed, the role of an asset class in risk reduction via diversification is highly regarded by investment professionals. Table 2 shows the correlation coefficients between the return series of the five Canadian asset types analyzed.

We can see that the two highest correlations are between the Property Index and the PFI at 0.7209 and between the XRE and XIU at 0.7459. In contrast, the correlation between REITs and private real estate is very low at 0.0383 with the Property Index and 0.2150 with the PFI. These results indicate that a diversified Canadian REIT portfolio seems to be affected by the same systematic (or idiosyncratic) risk as a diversified Canadian stock portfolio, while private real estate returns do not seem to be. This phenomenon can substantially impact the portfolio allocation process through the MVO framework.

Table 2 - Correlation Between Asset Classes

	Property Index	PFI	XBB	XIU	XRE
Property Index	1.0000				
PFI	0.7209	1.0000			
XBB	-0.1254	-0.3198	1.0000		
XIU	0.0845	0.1807	-0.1097	1.0000	
XRE	0.0383	0.2150	0.0816	0.7459	1.0000

Table 2 presents the correlation matrix of the quarterly returns for the five selected asset classes from December 31st, 2004, to September 30th, 2021.

Although, the appraisal valuation process can understate the very low correlation with private real estate. Another interesting result, despite the income nature of real estate as an investment, it is negatively correlated with investment-grade Canadian bonds, especially the PFI and XBB. The negative correlation between these two assets could yield excellent diversification benefits in low-risk/low-return portfolios.

4. Methodology

4.1. Adjusting Private Real Estate Data

The literature previously covered in this paper revealed two important biases impacting direct real estate returns: the effect of appraisal smoothing and the liquidity risk. Researchers know these two to underestimate the volatility of real estate returns. They can, in turn, overestimate the optimal allocation to private real estate in a mean-variance optimization (MVO) framework. In this first sub-section, we treat the raw private real estate data of the Property Index and the PFI to correct these biases by drawing adjustments techniques from the existing literature on the subject.

4.1.1. Unsmoothing Private Real Estate Returns

Since the Property Index and the PFI are primarily based on the appraisal valuation process to determine the capital growth over a period of time, they are subject to the same smoothing problem that generally occurs with these types of indices. When looking at the graphical return series and the standard deviations of these two datasets in the previous "Data" section of this paper, it seems evident that the appraisal smoothing effect is present in our private real estate indices.

A methodology must be employed to unsmooth the returns by undoing the effects of appraisal smoothing in both our private real estate indices to reveal the "true" or the effective capital returns. Research on the subject outlines that the appraised capital return of a property is expressed as a weighted average between the previous appraisal and the proper return based on current market conditions. Clayton, Geltner, and Hamilton (2001) modeled this phenomenon in the following fashion:

$$R_{App,t} = \alpha \times R_t + (1 - \alpha) \times R_{App,t-1}$$

Where:

- R_t is the effective capital growth return in quarter *t*;
- $R_{App,t}$ is the capital growth return based on an appraisal in quarter *t*;
- $R_{App,t-1}$ is the capital growth return based on an appraisal in quarter *t-1*;
- α is the confidence factor that appraisers have in the current market information.

Since Clayton, Geltner, and Hamilton (2001) are the ones who successfully estimated a reasonable level of confidence α of 0.815, that can be used to unsmooth the capital growth return without having to pose any additional assumption, it is their model that is selected in this paper to treat this matter. It is, therefore, a matter of simple algebra to rearrange the equation above to solve for the true capital growth return:

$$R_t = \frac{R_{App,t} - (1 - 0.815) \times R_{App,t-1}}{0.815}$$

Appendix B details the step-by-step method applied to the capital growth return and the income return to obtain total returns, in which the capital growth component is unsmoothed. Note that the smoothing only applies to the capital growth and not the income component of the total return.

4.1.2. Adjusting Private Real Estate Data for the Liquidity Risk

In finance, the most widely used measure of risk is the variance (or the standard deviation) of returns. However, the research covered in the literature review has demonstrated times and times again that standard variance calculation methodology fails to accurately quantify the risks of investing in private real estate by not reliably incorporating the liquidity risk associated with transacting real estate assets. In order to appropriately optimize asset allocation in an MVO framework, it is essential to adjust the private real estate data's variance to factor in the liquidity risk.

This paper borrows the variance adjustment model for liquidity and price risks from Lin and Vandell (2007). The formula they provide is:

$$(\sigma^{ex\ ante})^2 = \sigma^2 \times \left[1 + \left(\frac{Var(t)}{T_H + E[t]} \times \mu^2\right)\right]$$

Where:

- $(\sigma^{ex\ ante})^2$ is the variance adjusted for the liquidity risk;
- σ^2 is the variance of returns representing the price risk;
- $E[t]$ and $Var(t)$ are the two first moments of the time required for sale distribution;
- T_H represents the holding period as the number of years;
- μ represents the expected return.

To apply this last formula to our returns, information about the probability distribution function of the marketing period (the time required for sale) and the holding period is needed. Since neither are available to us, we use the results derived by Lin and Vandell (2007) to adjust the volatility metrics of the Property Index and the PFI. Appendix D presents their results.

Two assumptions are needed to determine the proper volatility adjustments for our private real estate datasets. First, the expected marketing period is inspired by Cheng, Lin, and Liu (2013), who defined the standard selling time in normal market conditions as being from 6 to 9 months. We chose an expected marketing period of 8 months for both datasets to stay on the more conservative side. As for the holding period, we believe the selected time horizon should be different for the Property Index than for the PFI. The reason behind this is that PERE funds have a predetermined term length. As Kuzmicki and Simunac (2008) explained in their paper, PERE funds are limited life vehicles with a life cycle typically ranging from three to eight years, meaning that the assets must be sold before the end of the fund's life.

Meanwhile, directly acquiring real estate assets does not have any time limit, and pension funds typically plan to hold these assets for an extended period. This limited life constraint of PERE funds is why a holding period of five years is assumed for the PFI, whereas we believe a holding period of ten years for the Property Index is appropriate. Under these assumptions, the variance of the unsmoothed Property Index is adjusted upwards by a factor of 30%. At the same time, the unsmoothed PFI is revised upwards by a factor of 57%.

4.2. Portfolio Simulation and Optimization Process

As mentioned earlier, this study aims to analyze the impact of integrating different Canadian real estate investment vehicles and a mix of these investment vehicles to a portfolio of Canadian stocks and investment-grade bonds (the control portfolio). Then, to measure how the risk-return performance is affected by the different asset combinations and data adjustments.

We developed a model based on the Modern Portfolio Theory (MPT) of diversification to do so. Introduced by Harry Markowitz in 1952, the MPT suggests that investors base their asset allocation decisions on the risk-return characteristics and co-movement of the assets returns. The basic concept is that by combining different types of assets, such as stocks, bonds, and real estate, into a portfolio, one can create diversification benefits and attain lower risk through imperfect correlation and co-movement between the asset classes.

The following methodology applies to every asset class combination tested, and we use whether unadjusted or adjusted private real estate data in the model. A mean-variance optimization model is used to derive our results and quantify the impacts of diversification. The first step of the mean-variance optimization process involves identifying and estimating the necessary inputs for the model. For each asset class, the expected return, the standard deviation, and the correlations between each pair of assets. But, before carrying out the optimization process, the feasible set of portfolios must be defined. This set is meant to represent the set of investment possibilities. To illustrate these possibilities, we develop a portfolio simulator designed to generate an arbitrarily large number of weighted combinations of the asset classes used as inputs, forming a set of simulated portfolios. We generate 100 000 random portfolios for each asset class combination tested for this study. The number 100 000 is assumed to be large enough to adequately represent the desired effects of the asset combination tested, even if the possible number of weighted allocations is technically infinite.

The portfolio simulator works by randomly attributing 100 000 sets of N weights, which always sums up to 1, for the N asset classes figuring in the test. Using these weights, we calculate the annualized expected return of the portfolios in the following fashion:

$$E[R_p] = \sum_{i=1}^{N} w_i \times E[R_i]$$

Where:

- $E[R_p]$ represents the expected return of a simulated portfolio;
- w_i are the simulated weights of the i = 1, …, N asset class;
- $E[R_i]$ represents the expected return approximated by taking the annualized historical mean of the i = 1, …, N asset class.

The following input is the annualized volatility of the portfolios, which is calculated with the formula:

$$\sigma_p = \sqrt{\sum_{i=1}^{N} w_i^2 \times \sigma_i^2 + 2 \times \sum_{i=1}^{n}\sum_{j=1}^{n} w_i \times w_j \times Cov_{i,j}}$$

Where:

- σ_p is the standard deviation of a simulated portfolio;
- σ_i^2 is the i = 1, …, N asset class variance;
- And the variance-covariance matrix for an N asset class portfolio is given by:

$$Cov_{i \dots N} = \begin{bmatrix} \sigma_1^2 & \cdots & \sigma_{N,1} \\ \vdots & \ddots & \vdots \\ \sigma_{1,N} & \cdots & \sigma_N^2 \end{bmatrix}$$

Then, the Sharpe ratios, which will allow us to rank the portfolios in terms of their risk-return relationship, are calculated for each simulated portfolio by dividing the expected return, minus the risk-free rate, by its standard deviation:

$$SR_p = \frac{(E[R_p] - R_f)}{\sigma_p}$$

Where:

- SR_p is the Sharpe ratio of a simulated portfolio;
- R_f is the risk-free rate of 1.3765%, calculated by averaging the annualized yield of the 3-month T-bill over the analyzed period.

Now that the portfolio simulations are done and all the necessary inputs have been calculated, the optimization process can begin. The next step is to determine the efficient frontier. The efficient frontier comprises the set of portfolios that minimize the risk for a given level of expected return with respect to certain constraints. In this model, no short selling or buying on margin is allowed, meaning that the weights are constrained to the range [0,1], mainly because these strategies rarely fit in a pension fund long-term investment plan. This also means that none of the asset classes have a strategic allocation limit inside the [0,1] range. Additionally, the sum of the weights needs to be equal to one for every portfolio simulated and optimized. Mathematically, we represent this model with the following optimization function and constraints:

$$Minimize\ \sigma_p^2 = \sum_{i-1}^{N} \sum_{j=1}^{N} w_i \times w_j \times Cov_{i,j}$$

Subject to:

$$\textstyle\sum_{i=1}^{N} w_i \times E[R_i] - E[R_p^*] = 0;$$

$$\textstyle\sum_{i=1}^{N} w_i - 1 = 0;$$

$$0 \leq w_i \leq 1 \; for \; i = 1, \ldots, N.$$

Where $E\left[R_p^*\right]$ is the range of possible returns from the feasible set of portfolios analyzed.

The minimization process quantifies the optimal portfolio allocation for each level of possible return by attributing the correct proportion of weight to each asset class. Then, the minimum-variance portfolio is identified and marks the beginning of the efficient frontier, which extends to the highest return level possibly attainable in the specific feasible set of portfolios tested. These are the frontiers that we compare between the different asset combination tests. On the efficient frontiers modelized lies the optimal portfolio. The optimal portfolio has the highest Sharpe ratio, representing the asset allocation that delivers the best risk to expected return tradeoff for a specific asset class combination test.

For this study, a total of fourteen different asset class combination optimizations are tested (excluding the robustness tests combinations). The three main combinations include the control frontier, constituted of XBB and XIU only, followed by the mix of the five asset classes with unadjusted private real estate data, and finally, the combination of the five asset classes with adjusted private real estate data. These are all part of the primary results presented in the next section. We present five more efficient frontiers representing five other asset combinations as additional or secondary results. These are mainly the additions of individually selected real estate asset vehicles to the control frontier with and without adjusting the private real estate data. The other six combinations are not included in this paper since they do not provide any additional insights or valuable information on the matter. They are essentially hybrid combinations with two real estate investment vehicles in addition to XBB and XIU. As we can imagine, their results are middle frontiers between the individual real estate efficient frontiers and the full five asset classes combinations results.

4.3. Sensitivity Analysis

This first robustness test focuses on the tri real estate investment vehicle combination with the stock ETF and the bond ETF with adjusted private real estate data since this asset combination best summarizes the results of this research. We carry out a sensitivity analysis on the inputs used to generate the main efficient frontier of the simulations and the MVO model to render the results even more robust. The other efficient frontiers modelized are not tested in the sensitivity analysis because of the high number of possible combinations times the number of scenarios and the low additional value each of these tests would provide.

To conduct this analysis, two additional scenarios concerning the mean and the standard deviation of returns of the two private real estate investment vehicles are modified and contrasted with the results previously calculated (the standard scenario). These two new scenarios are the pessimistic and the catastrophic scenarios. We reduce the mean return of the Property Index and the PFI by as much as -10% and -25%, respectively. Meanwhile, their standard deviations are further increased by +10% and +25%. Note that the percentage increase in volatility is compounded (and not added) with the previous adjustments made to unsmooth the returns and account for the liquidity risk associated with private real estate assets and PERE funds. The goal of this analysis is to highlight how sensitive the MVO framework model is to changes in the input's values and also to observe how the optimal portfolios compositions would change if the Canadian economic environment of private commercial real estate deteriorated compared to the solid and steady growth this asset class has seen over the past two decades, providing insights in downside scenarios.

We conduct a second robustness test in which an additional scenario is created in an attempt to explain the abnormally strong performance of private real estate from another angle. We believe that the appraisal smoothing is likely to be at the core of the correlation and risk understatement of private real estate and that compensating for it should significantly reduce the diversification benefits that real estate creates when it is included in the efficient portfolios. To test this hypothesis, we modify the correlations by assuming that the correlation between private and public real estate investment vehicles is robust and much higher than what we calculated with our return series. We also assume that the correlation between private real estate and the stock ETF is similar to that between the REITs ETF and the stock ETF. The "unsmoothing" parameter is also enhanced to inject more volatility in the Property Index and the PFI's serial returns. This analysis aims to highlight how the results from the MVO model are affected by these changes and to observe if private real estate would still enhance portfolio performance under these new assumptions.

5. Primary Results and Interpretations

5.1. Adjusted Private Real Estate Data Results

The first adjustment made to private real estate data has been to unsmooth the capital return component to yield unsmoothed total return series for the Property Index and the PFI. The "unsmoothing" method effectively accentuated the volatility of returns which was underestimated by the appraisal valuation method. Table 3 contrasts the descriptive statistics of the two private real estate datasets before and after the adjustment. While the quarterly mean of returns almost stayed the same, the quarterly standard deviation of these two datasets went from 1.96% and 1.80% to 2.27% and 2.00% for the Property Index and the PFI, respectively. Then, another adjustment impacting the volatility of returns took place to consider the liquidity risk associated with transacting real estate assets. As explained in the methodology section, the Property Index standard deviation is increased by a factor of 1.3 to 2.95% and the PFI by a factor of 1.57, putting its quarterly standard deviation at 3.15%.

Table 3 - Unsmoothed Quarterly Descriptive Statistics

	Property Index	Unsmoothed Property Index	PFI	Unsmoothed PFI
Observations	67	67	67	67
Mean	2.13%	2.11%	1.96%	1.97%
Std-dev.	1.96%	2.27%	1.80%	2.00%
Min	-2.26%	-3.12%	-6.92%	-8.82%
0.25	1.37%	1.26%	1.35%	1.26%
0.50	1.84%	1.65%	2.11%	2.02%
0.75	2.73%	2.60%	3.19%	3.00%
Max	9.01%	10.34%	5.23%	5.53%

Table 3 presents summary statistics of the quarterly returns from December 31st, 2004, to September 30th, 2021, for the five selected asset classes, the Unsmoothed Property Index and the Unsmoothed PFI. The Unsmoothed Property Index and the Unsmoothed PFI have been adjusted for the appraisal smoothing bias.

Table 4 presents the correlation matrix before and after the adjustments. The changes are marginal; the correlation between the two private real estate indices and XBB decreased by a few points, possibly further increasing the diversification benefits of combining these asset classes. In contrast, the correlation between XIU and XRE increased ever so slightly. The most significant change is between the Property Index and the PFI, where the correlation went down from 0.7209 to 0.6660. As expected, the correlations between the three ETFs are unchanged since the adjustments were only applicable to the private real estate indices.

Table 4 - Correlation Between Asset Classes and Adjusted Private Real Estate Data

	Property Index	PFI	Adj. Property Index	Adj. PFI	XBB	XIU	XRE
Property Index	1.0000						
PFI	0.7209	1.0000					
Adj. Property Index	0.9828	0.6743	1.0000				
Adj. PFI	0.6938	0.9856	0.6660	1.0000			
XBB	-0.1254	-0.3198	-0.1411	-0.3540	1.0000		
XIU	0.0845	0.1807	0.0948	0.2148	-0.1097	1.0000	
XRE	0.0383	0.2150	0.0412	0.2432	0.0816	0.7459	1.0000

Table 4 presents the correlation matrix of the quarterly returns from December 31st, 2004, to September 30th, 2021, for the five selected asset classes, the adjusted Property Index, and the adjusted PFI. The adjusted Property Index and the adjusted PFI have been adjusted for the appraisal smoothing and the liquidity risk.

5.2. Portfolio Simulation and Optimization Results

This paper examines real estate's role in pension funds' mixed-asset portfolios using both private and public real estate indices. It reviews various real estate-related risk/return opportunities and assesses the impact on portfolio performance of integrating these real estate investment vehicles. Before debuting the analysis of the results, it is important to advise the reader that the weighted allocations explored in this study are not meant to be reproduced in a portfolio management setting. The reasoning here is that as more real estate investment vehicles are added to the model, the number of real estate dimensions to the optimization exercise increases. In doing so, the likelihood of more significant real estate allocations almost necessarily increases because, given our model constraints, the allocation to any asset class can never be less than zero. To be prudent, an investment manager should increase the number and type of stocks and bonds vehicles in the model and include other forms of alternative investment as well.

As a base case, figure 3 and table 5 show the results when only Canadian investment-grade bonds (XBB) and Canadian equities (XIU) are included in the portfolio simulation and optimization model.

<u>Figure 3 - Efficient Frontier of the Control Portfolios</u>

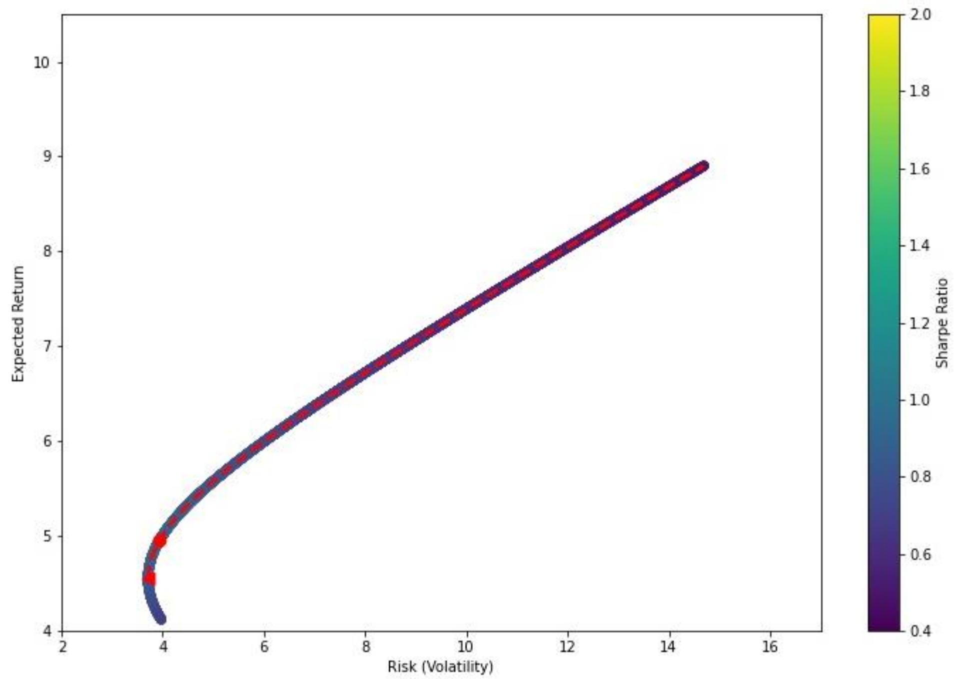

Figure 3 plots the optimized control efficient frontier constituted of XBB and XIU. The constraint confines each asset's weight (w_i) to [0, 1]. The red dotted line represents the efficient frontier, whereas the red dot and the red triangle represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio. This plot also illustrates the 100 000 simulated portfolios for the selected asset classes. The portfolios are ranked by Sharpe ratio: the darker colors (purple and blue) represent lower Sharpe ratios, and lighter colors (green and yellow) represent higher Sharpe ratios.

The optimal allocations not including real estate are, as one might expect, large allocations to investment-grade bonds for the very conservative portfolios and large allocations to equities for the more aggressive portfolios.

The optimal portfolio has a 0.91 Sharpe ratio with an 82% bond allocation and a 18% stock allocation. The heavier weight on XBB is partly explained by the asset class's low volatility, which facilitates the maximization of the Sharpe ratio. As the target portfolio return increases, the stock allocation gets more significant, and the Sharpe ratio falls. The risk compensation seems to be better for the investment-grade bond-heavy portfolios, but it comes at the cost of losing the higher expected return portfolios.

Table 5 – Optimal Asset Allocations (Control Frontier)

Target Portfolio Return	Portfolio Volatility	Sharpe-Ratio	Weights in Optimal Portfolio	
			XBB	XIU
4.55%	3.71%^{MV}	0.86	0.91	0.09
4.96%	3.94%	0.91^{Opt}	0.82	0.18
5.00%	3.99%	0.91	0.82	0.18
6.00%	6.01%	0.77	0.61	0.39
7.00%	8.82%	0.64	0.40	0.60
8.00%	11.86%	0.56	0.19	0.81
8.90%	14.69%	0.51	0.00	1.00

Table 5 provides the portfolio volatility, the Sharpe ratio, and the weights of XBB and XIU in optimal portfolios for varying target portfolio returns of the control frontier. The "MV" superscript identifies the row of the minimum variance portfolio in the control model, and the "Opt" superscript identifies the portfolio's row with the highest Sharpe ratio, the optimal portfolio.

The optimal allocations without real estate can be compared to those when we add real estate to the mix. Figure 4 and Table 6 include the three real estate investment vehicles: the Property Index, the PFI, and XRE, in addition to XBB and XIU. In this model, the private real estate returns have not been adjusted for the appraisal smoothing and the liquidity risk; we call it the unadjusted model. It is apparent from the figure that the inclusion of real estate in the asset mix drastically improved portfolio performance. The high Sharpe ratio portfolios, denoted by the yellow color, are frequent in the 6.0% to 8.5% target return range, a risk/return performance level that was far from being attainable before the addition of real estate investment vehicles.

Investors can obtain higher returns with lower risk now that real estate is included in the mix. Also, more elevated expected returns are attainable because of the contribution of XRE. The REIT ETF has a higher expected return, on average, than the stock ETF XIU, which allows portfolios to reach beyond the 9.0% target portfolio return level. However, these higher expected returns also come with a sharp increase in volatility.

Figure 4 - Efficient Frontier Including Real Estate (Unadjusted)

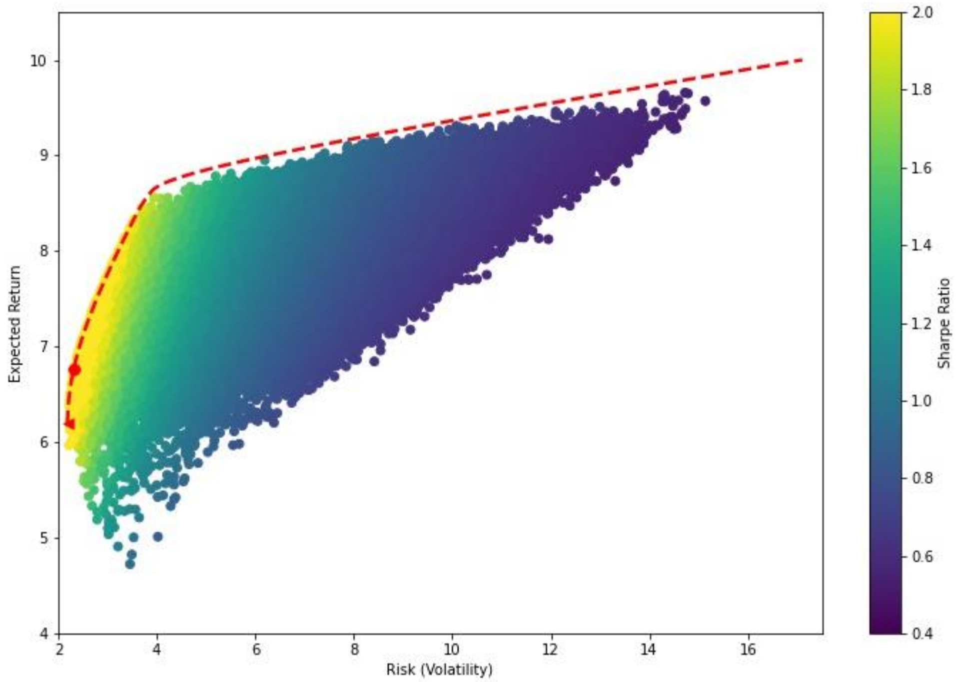

Figure 4 plots the optimized, efficient frontier constituted of XBB, XIU, XRE along with the unadjusted Property Index and the unadjusted PFI. The constraint confines each asset's weight (w_i) to [0, 1]. The red dotted line represents the efficient frontier, whereas the red dot and the red triangle represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio. This plot also illustrates the 100 000 simulated portfolios for the selected asset classes. The portfolios are ranked by Sharpe ratio: the darker colors (purple and blue) represent lower Sharpe ratios, and lighter colors (green and yellow) represent higher Sharpe ratios.

A point of interest is that every efficient portfolio includes some allocation to at least one of the real estate investment vehicles. However, the optimal allocation to real estate is quite significant, perhaps too substantial. The PFI and the investment-grade bond ETF dominate the asset allocation for the more conservative portfolios. This is due to the low volatility of these two asset classes and the negative correlation between them, generating massive diversification benefits. In the middle portion of the efficient frontier, the PFI and XBB are gradually substituted by a greater allocation in the Property Index since it delivers higher expected returns while still maintaining pretty low volatility and low correlation with XBB. Then, the 8.5% target return marks a sudden turning point above which most asset classes cannot deliver the expected returns, so the asset allocation quickly shifts to the REIT ETF. XRE has a less performant risk/return trade-off characteristic than most asset classes, meaning that a more significant share in XRE causes a drastic change in the slope of the efficient frontier. From this point, an increase in the target portfolio return comes with an even more considerable increase in portfolio volatility. As expected, the REITs dominate the more aggressive portfolios but, surprisingly so, without any weight allocated to the XIU ETF. XIU seems to be stochastically dominated in every portion of the efficient frontier.

Table 6 – Optimal Asset Allocations Including Real Estate (Unadjusted)

Target Portfolio	Portfolio	Sharpe-	Weights in Optimal Portfolio				
Return	Volatility	Ratio	Property Index	PFI	XBB	XIU	XRE
6.20%	2.19%^{MV}	2.20	0.04	0.49	0.46	0.01	0.00
6.77%	2.32%	2.33^{Opt}	0.21	0.44	0.33	0.02	0.00
7.00%	2.43%	2.31	0.28	0.41	0.29	0.02	0.00
8.00%	3.22%	2.01	0.57	0.31	0.08	0.02	0.02
9.00%	6.27%	1.22	0.68	0.00	0.00	0.00	0.32
10.00%	17.10%	0.50	0.03	0.00	0.00	0.00	0.97

Table 6 provides the portfolio volatility, the Sharpe ratio, and the weights of the unadjusted Property Index, the unadjusted PFI, XBB, XIU, and XRE in optimal portfolios for varying target portfolio returns of the control frontier. The "MV" superscript identifies the row of the minimum variance portfolio in the unadjusted model, and the "Opt" superscript identifies the portfolio's row with the highest Sharpe ratio, the optimal portfolio.

The optimal portfolio of this asset mix has a Sharpe ratio of 2.33, more than twice as much as the 0.91 Sharpe ratio of the optimal portfolio for the control efficient frontier. This higher performance can be attributed to a 21% allocation in the Property Index, a 44% allocation to the PFI, a 33% allocation to XBB, and a small 2% allocation to XIU (0% on XRE). A stunning total of 65% of the optimal portfolio is allocated to real estate in this model using unadjusted real estate data.

From the results above, it seems that the stochastic dominance and the diversification benefits of real estate are overstated since the private real estate data hasn't been adjusted yet for the appraisal smoothing and the liquidity risk. Figure 5 and Table 7 below present the same asset class mix as above; this time, using adjusted private real estate data, we call it the adjusted model. Let's look at how the adjustments impact the results.

Figure 5 - Efficient Frontier Including Real Estate (Adjusted)

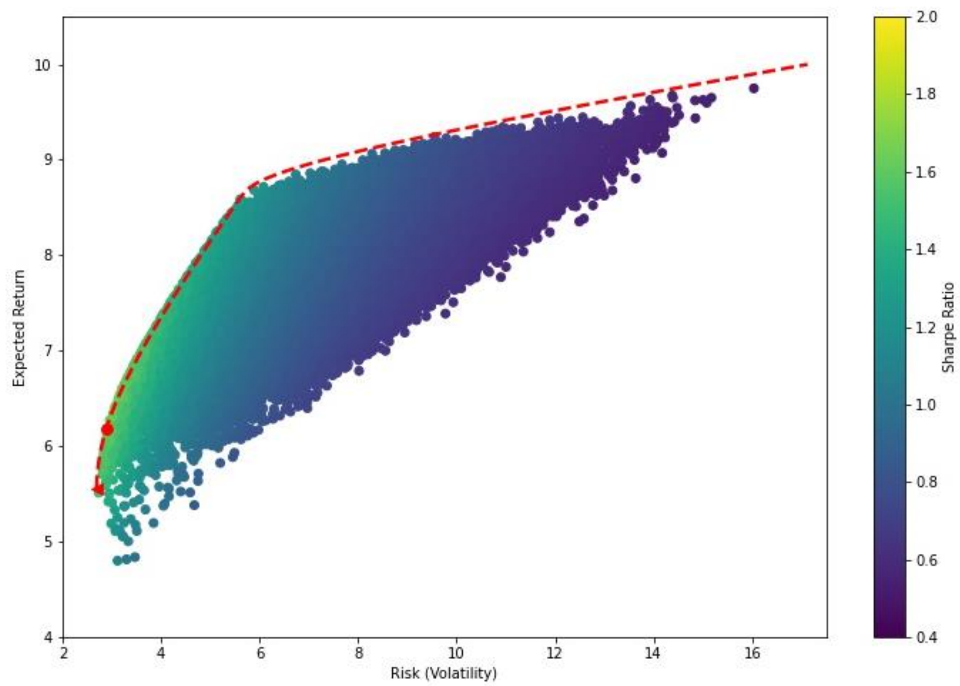

Figure 5 plots the optimized, efficient frontier constituted of XBB, XIU, XRE along with the adjusted Property Index and the adjusted PFI. The constraint confines each asset's weight (w_i) to [0, 1]. The red dotted line represents the efficient frontier, whereas the red dot and the red triangle represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio. This plot also illustrates the 100 000 simulated portfolios for the selected asset classes. The portfolios are ranked by Sharpe ratio: the darker colors (purple and blue) represent lower Sharpe ratios, and lighter colors (green and yellow) represent higher Sharpe ratios.

The figure shows that the higher Sharpe ratio portfolios, or the "yellow" portfolios, are now barely attainable by a tiny proportion of portfolios. Even for those, the Sharpe ratio does not exceed 1.66. A is a considerable reduction compared to the previous model. Since the adjustments increased the volatility of returns associated with the Property Index and the PFI, any portfolio heavily allocated to these asset classes is translated to the right in the visual representation above as their risk/return performance diminishes. This translation explains the "flattening" effect between the adjusted and the unadjusted models in the 6.0% to 8.5% target portfolio return range. Above the 8.5% mark, the simulated portfolios of the two models exhibit similar performance.

Even after the adjustments, the performance is still remarkably better with real estate than without any. This model is considered more representative of the impact on portfolio performance of integrating the three real estate investment vehicles to Canadian stocks and bonds portfolios.

Table 7 – Optimal Asset Allocations Including Real Estate (Adjusted)

Target Portfolio Return	Portfolio Volatility	Sharpe-Ratio	Weights in Optimal Portfolio				
			Property Index	PFI	XBB	XIU	XRE
5.56%	2.69%^{MV}	1.55	0.07	0.27	0.63	0.03	0.00
6.18%	2.89%	1.66^{Opt}	0.22	0.21	0.51	0.06	0.00
7.00%	3.60%	1.56	0.42	0.18	0.32	0.06	0.02
8.00%	4.80%	1.38	0.66	0.11	0.12	0.04	0.07
9.00%	7.31%	1.04	0.66	0.00	0.00	0.00	0.34
10.00%	17.12%	0.50	0.03	0.00	0.00	0.00	0.97

Table 7 provides the portfolio volatility, the Sharpe ratio, and the weights of the adjusted Property Index, the adjusted PFI, XBB, XIU, and XRE in optimal portfolios for varying target portfolio returns of the control frontier. The "MV" superscript identifies the row of the minimum variance portfolio in the adjusted model, and the "Opt" superscript identifies the portfolio's row with the highest Sharpe ratio, the optimal portfolio.

Table 7 details the performance measures and weighted allocations for various portfolio return targets. As with the previous model, every efficient portfolio includes some allocation to at least one of the real estate investment vehicles. However, in this model, the allocation to these real estate investment vehicles is lower in every portion of the efficient frontier except for the most aggressive portfolios, where the allocation to REITs gets close to or equal to 100%. For the more conservative portfolios, the PFI still has a noteworthy presence. Nevertheless, it is noticeably smaller than in the previous model. For example, the minimum variance portfolio has a 27% allocation to the PFI, down from 49% in the unadjusted model, while the investment-grade bond ETF saw its weight increase from 46% to 63%. Even with this shift in allocation, the conservative portfolios are still primarily dominated by the combination of the PFI and XBB. As the target portfolio return increases in the 7.0% to 9.0% range, the weight attributed to the Property Index gets materially heavier even after "unsmoothing" its returns and after the upward adjustment of its standard deviation. Meanwhile, the allocation to XIU stayed insignificantly small. Another point of interest concerning the middle part of the efficient frontier is that in the adjusted model, the weighting attributed to XBB remains superior to the PFI, while the PFI significantly outweighed XBB in the unadjusted model. Unsurprisingly, as the expected return transitions upwards to the more aggressive return targets, the allocation shifts to REITs, and the performance observed corresponds to the version of the unadjusted model.

The optimal portfolio of this efficient frontier has a Sharpe ratio of 1.66, which is still remarkably higher than the optimal portfolio of the control efficient frontier. This Sharpe ratio seems more realistic and more reasonably attainable than the 2.33 in the unadjusted model, which leads us to believe that the adjustments made to the real estate return data were justified and even necessary. This optimal portfolio is allocated in the following fashion: 22% on the Property Index, 21% on the PFI, 51% on XBB, 6% on XIU, and 0% on XRE. The total allocation to real estate is 43%, which is significantly lower than the previous model but still much higher than the scientific consensus and also than what is observed in practice.

5.3. Results Synthesis and Discussion

The optimization results indicate that Canadian pension funds would be better served by including some form of real estate in their asset allocation than not, based on their targeted risk/return objectives. Consistent with previous studies on real estate allocation in optimized portfolios, the most attractive feature of private real estate is its lack of correlation with the other asset classes, followed by its low standard deviation of returns. Even after correcting for the risk understatement of private real estate caused by the smoothness of returns and the liquidity risk, real estate remains a desirable feature of a well-diversified portfolio, although in lower allocation than before the adjustments. The relation between the unadjusted efficient frontier and the adjusted efficient frontier is similar to Bond, Hwang, and Richards (2006). They adjusted the private real estate returns for the liquidity risk and found that the unadjusted efficient frontier lies above the adjusted efficient frontier, as we also represent in figure 6. The gap between the two models is wide in the lower volatility range of 2.0% to 6.0%. In the adjusted model, the allocation to private real estate investment vehicles is significantly smaller. Yet, as we move towards more volatile portfolios, the allocation remains broadly similar because of the unanimous shift towards REITs.

We deduce insights concerning the role of each real estate investment vehicle in a mixed-asset portfolio by looking at the way the allocation to real estate varies throughout the different portfolio target returns. For more risk-averse investors, the addition of unlevered private real estate properties or core style PERE funds to a stocks and investment-grade bonds portfolio can significantly improve the risk/return performance, as measured by the materially higher Sharpe ratios of the portfolios in the optimization results analyzed. On the other hand, for the investors who prefer high-risk portfolios, public REITs seem to be the right vehicle serving this risk-aversion preference. Consistent with Mutahi and, Othieno (2015), REITs provide more return but are much more volatile than private real estate and, in this case, even slightly more than stocks. Levering up private real estate properties and investing in value-add or opportunistic style PERE funds could potentially offer an alternative to REITs. However, these investment strategies are not explored in this paper.

Figure 6 - Efficient Frontiers Synthesis

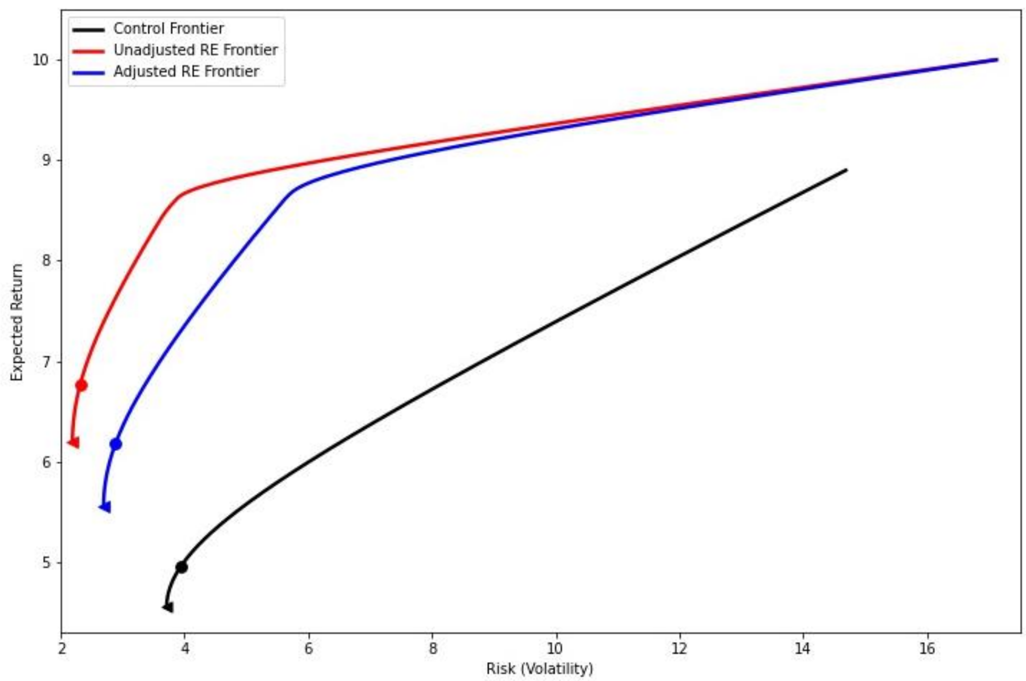

Figure 6 syntheses the three efficient frontiers previously illustrated. The black line represents the control efficient frontier constituted of XBB and XIU. The blue line represents the efficient frontier constituted of XBB, XIU, XRE with the adjusted Property Index and the adjusted PFI. The red line represents the efficient frontier constituted of XBB, XIU, XRE with the unadjusted Property Index and the unadjusted PFI. The dot and triangle on each frontier represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio, respectively.

Even though a multitude of studies also find substantial allocations to real estate in their optimization results and that a 40% optimal allocation or so to real estate resulting from an MVO framework is not exceptional or even rare in the literature on the matter, the consensus for the optimal allocation seems to be around the 15% to 25% range. Our 43% allocation to real estate in the optimal portfolio of the adjusted model exceeds this range. It is far above the observed average of 12% for the Canadian defined benefit pension plans of 2019 (Appendix A). However, our results are close to the total of 38% that Canadian defined benefit pension plans allocated to alternative asset classes that same year. Perhaps adding other alternative asset classes to our models could reduce the overall allocation to real estate by sharing the total weight.

Since our asset classes exhibit similar risk/return and correlation measures to what is found in previous studies, the difference in asset allocation results can be attributed, in our opinion, to the dimensionality problem and to the outperformance of the real estate investment vehicles during the period compared to stocks and bonds. In our mixed-asset models, the number of real estate dimensions to the optimization exercise is three out of five in total, which means that the over-allocation of real estate could be partly explained by this dynamic. In most studies, the number of real estate dimensions to the optimization exercise is one out of three to two out of four. In practice, pension funds also often include other alternative asset classes like private equity, infrastructure, and commodities which act as competitors to real estate in the strategic allocation process. To more accurately estimate the optimal allocation to real estate for a practical application, prudence would require increasing the number of asset classes and increasing the types of common-stock and bonds indices (e.g., various "styles" or from different geographical regions) in the model. As for the outperformance of real estate, a sensitivity analysis is conducted in the next section to test the robustness of our results against a change in the risk/return inputs of the models.

6. Robustness Check and Additional Results

6.1. Robustness Check #1

In the following, we analyze how robust the adjusted efficient frontier results are with respect to a change in the input's value of the Property Index and the PFI. Since the period analyzed is relatively short (17 years of data) and the quarterly observations are limited, a sensitivity analysis of the private real estate returns helps reduce the probability that the strong results of private real estate might be a product of a specifically favorable economic environment. A misestimation of the actual returns caused by the appraisal valuation process is another strong reason for the analysis.

From the literature, we know about the inefficiencies of private real estate markets, which often lead to an overestimation of the risk/return performance. We also know that real estate liquidity and transaction volume tend to dry up in more difficult economic environments. This begs the question, were the adjustments severe enough? Furthermore, would private real estate investments still be beneficial to pension funds' portfolio performance if the economic conditions of the private real estate sector had been less favorable over the last two decades?

To test the robustness of our results with stricter private real estate inputs, we add two new scenarios: the pessimistic and the catastrophic scenarios. In which the mean return of the Property Index and the PFI are negatively impacted by a -10% and a -25% drop, respectively. Meanwhile, their standard deviations are further increased by +10% and +25%. These statistical changes allow us to illustrate and analyze how poorer returns and riskier performances fare up in the MVO model. Figure 7 and Table 8 highlight how the two new additional scenarios compare to the base case (the efficient control frontier), the efficient public frontier (XRE with XBB and XIU), and to the standard scenario (the adjusted efficient frontier including real estate). We present the pessimistic and the catastrophic scenarios illustrated individually in appendix E.

Here, the increase in volatility is compounded with the previous adjustments made to the private real estate indices. Compared to raw data, the total inflation of their standard deviations is 66% for the Property Index and 91% for the PFI in the pessimistic scenario, coupled with a -10% reduction of their mean return. In the catastrophic scenario, the inflation of their standard deviations is 88% for the Property Index and 118% for the PFI. On top of the -25% reduction of their mean return.

With inferior inputs, it is not a surprise to see in the figure below the efficient frontiers of the pessimistic and the catastrophic scenarios under the standard scenario. However, despite these penalties, the efficient frontiers of the two new scenarios are still notably higher on the chart than the public and control frontier. It implies that even in a significantly less favorable private real estate economic environment, the inclusion of directly held real estate properties or core style PERE funds can improve pension funds' portfolios' risk/return performance. An interesting thing to point out is that, even if the new scenarios lowered the efficient frontier, the optimal portfolios, marked by the orange and the greenish dot, almost stayed stationary regarding their portfolio volatility. This could indicate that in deteriorating conditions, the resilience of private real estate could still act as a good portfolio stabilizer.

Figure 7 - Efficient Frontier Sensitivity Analysis 1

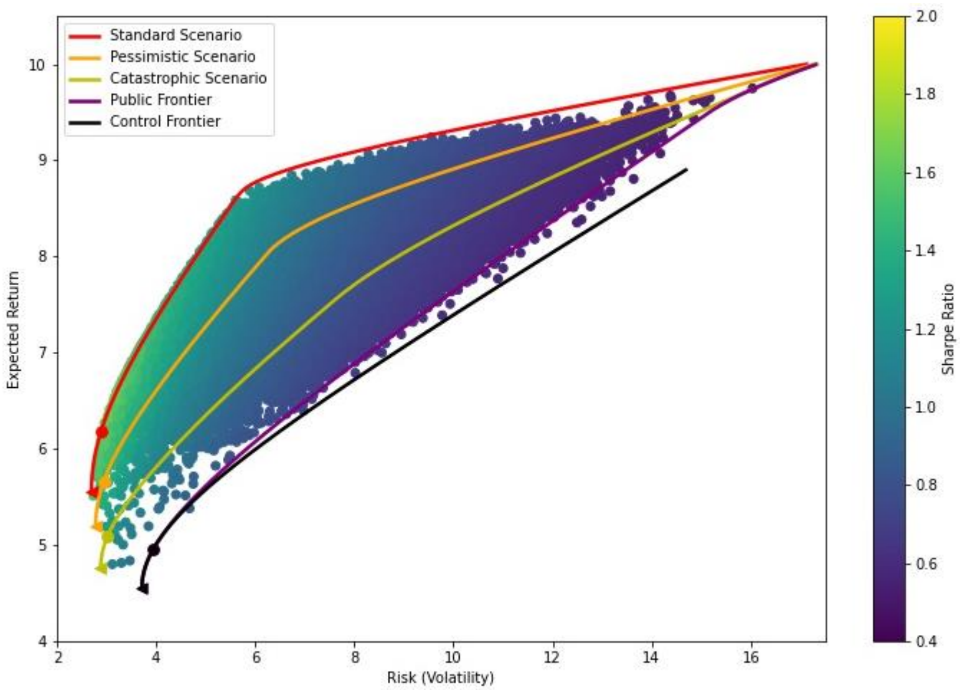

Figure 7 plots the efficient frontier, including adjusted private real estate under three different scenarios against two benchmarks. The lines represent the efficient frontiers, whereas the dot and the triangle on the lines represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio, respectively, in the analyzed scenario. The constraint confines each asset's weight (wi) to [0, 1]. The red line (the standard scenario) is the efficient frontier constituted of XBB, XIU, XRE with the adjusted Property Index and the adjusted PFI, as seen earlier in Figures 5 & 6. The orange line (the pessimistic scenario) represents the same asset combination as the standard scenario, but the Property Index and the PFI both had the mean return reduced by 10%, and their volatility increased by 10%. The yellow line (the catastrophic scenario) represents the same asset combination as the standard scenario, in which the Property Index and the PFI had the mean return reduced by 25%, and their volatility increased by 25%. The two benchmarks in this figure are the Public Frontier (constituted of XBB, XIU, and XRE) and the Control Frontier (comprised of XBB and XIU only). This plot also illustrates the 100 000 simulated portfolios for the selected asset classes. The portfolios are ranked by Sharpe ratio: the darker colors (purple and blue) represent lower Sharpe ratios, and lighter colors (green and yellow) represent higher Sharpe ratios.

Table 8 details the performance measures and the weighted allocations for various portfolio return targets of the pessimistic and catastrophic scenarios. For the pessimistic scenario displayed in Panel A, the optimal portfolio of the efficient frontier has a Sharpe ratio of 1.46, down from 1.66 in the standard scenario. This portfolio allocates 18% to the Property Index, 21% to the PFI, 55% to XBB, 6% to XIU, and 0% to XRE. The total allocation to real estate is 39% in this case. For the catastrophic scenario displayed in Panel B, the optimal portfolio of the efficient frontier has a

Sharpe ratio of 1.23, which is still higher than the 0.91 Sharpe ratio of the optimal portfolio from the public frontier and the control frontier but significantly lower than the 1.66 found in the adjusted model.

When we compare the optimal portfolio of the standard scenario to the catastrophic scenario, we observe that the drop in the Sharpe ratio is almost entirely explained by the lower expected return since the standard deviation only increased by 13 basis points. The optimal portfolio weights in the catastrophic scenario are dispersed as follows: 13% on the Property Index, 17% on the PFI, 62% on XBB, 8% on XIU, and 0% on XRE. The total allocation to real estate is still pretty high at 30% despite the significantly poorer performance of the private real estate investment vehicles taken individually, demonstrating the strength of low correlation in the MVO framework.

Table 8 - Optimal Asset Allocation Sensitivity Analysis 1

Panel A: Pessimistic Scenario

Target Portfolio Return	Portfolio Volatility	Sharpe-Ratio	Weights in Optimal Portfolio				
			Property Index	PFI	XBB	XIU	XRE
5.19%	2.77%^{MV}	1.37	0.06	0.24	0.67	0.03	0.00
5.66%	2.95%	1.46^{Opt}	0.18	0.21	0.55	0.06	0.00
6.00%	3.24%	1.43	0.27	0.18	0.46	0.09	0.00
7.00%	4.58%	1.23	0.53	0.09	0.24	0.07	0.07
8.00%	6.19%	1.07	0.79	0.00	0.01	0.06	0.14
9.00%	10.54%	0.72	0.43	0.00	0.00	0.00	0.57
10.00%	17.28%	0.50	0.02	0.00	0.00	0.00	0.98

Panel B: Catastrophic Scenario

Target Portfolio Return	Portfolio Volatility	Sharpe-Ratio	Weights in Optimal Portfolio				
			Property Index	PFI	XBB	XIU	XRE
4.76%	2.88%^{MV}	1.17	0.04	0.21	0.70	0.05	0.00
5.09%	3.02%	1.23^{Opt}	0.13	0.17	0.62	0.08	0.00
6.00%	4.33%	1.07	0.36	0.05	0.40	0.12	0.07
7.00%	6.36%	0.88	0.54	0.00	0.15	0.14	0.17
8.00%	8.87%	0.74	0.50	0.00	0.00	0.16	0.34
9.00%	12.73%	0.60	0.22	0.00	0.00	0.20	0.58
10.00%	17.12%	0.50	0.00	0.00	0.00	0.04	0.96

Table 8 Panel A provides the portfolio volatility, the Sharpe ratio, and the weights of the adjusted Property Index, the adjusted PFI, XBB, XIU, and XRE in optimal portfolios for varying target portfolio returns for the pessimistic scenario. Panel B provides the portfolio volatility, the Sharpe ratio, and the weights of the adjusted Property Index, the adjusted PFI, XBB, XIU, and XRE in optimal portfolios for varying target portfolio returns for the catastrophic scenario. The "MV" superscript identifies the row of the minimum variance portfolio, and the "Opt" superscript identifies the portfolio's row with the highest Sharpe ratio, the optimal portfolio, for their respective scenarios.

6.2. Robustness Check #2

Intuitively, reducing returns and accentuating the volatility noticeably narrows the spread of the adjusted efficient frontier, including real estate with the benchmarks. Varying the inputs value of the Property Index and the PFI in Robustness Check #1 didn't effectively explain the abnormal outperformance of the two private real estate indices in the model. However, the gap between the models still needs more investigating, in our opinion.

We now turn to the smoothing problem caused by the appraisal valuation process to explain this abnormal performance. Maybe the adjustment to unsmooth the returns of the private real estate indices is not significant enough to reflect the proper variation of values and also if the appraisal-based valuation process could be at the core of the correlation understatement with publicly traded assets. We ask, would private real estate still be beneficial to pension funds portfolios if its correlation with REITs was higher at 0.9 (as sharing similar underlying assets could implicate) and if its correlation with the stock market would be similar as with the REITs at 0.7459? Also, what happens to private real estate if we couple the modified correlations with a significant change in "α" the appraiser confidence factor (the "unsmoothing" factor) from 0.815 to 0.4, further accentuating the variation of returns? See Figure 8 for the results.

This figure highlights the diversification benefits of private real estate, assuming correlation parity between private and public markets. The gap between the public and control frontier shows the benefits of adding XRE to the XBB and XIU portfolios. In contrast, the gap between the public and the orange frontier (correlation parity scenario with enhanced "unsmoothing" parameter) shows the additional benefits of integrating private real estate investment vehicles to stocks, investment-grade bonds, and REITs portfolios. Interestingly, even under the assumptions stated above, directly held unlevered real estate properties and core PERE funds still seem beneficial to portfolios from a mean-variance optimization standpoint. Although, these benefits are noticeably smaller than the other previously explored scenarios.

Figure 8 - Efficient Frontier Sensitivity Analysis 2

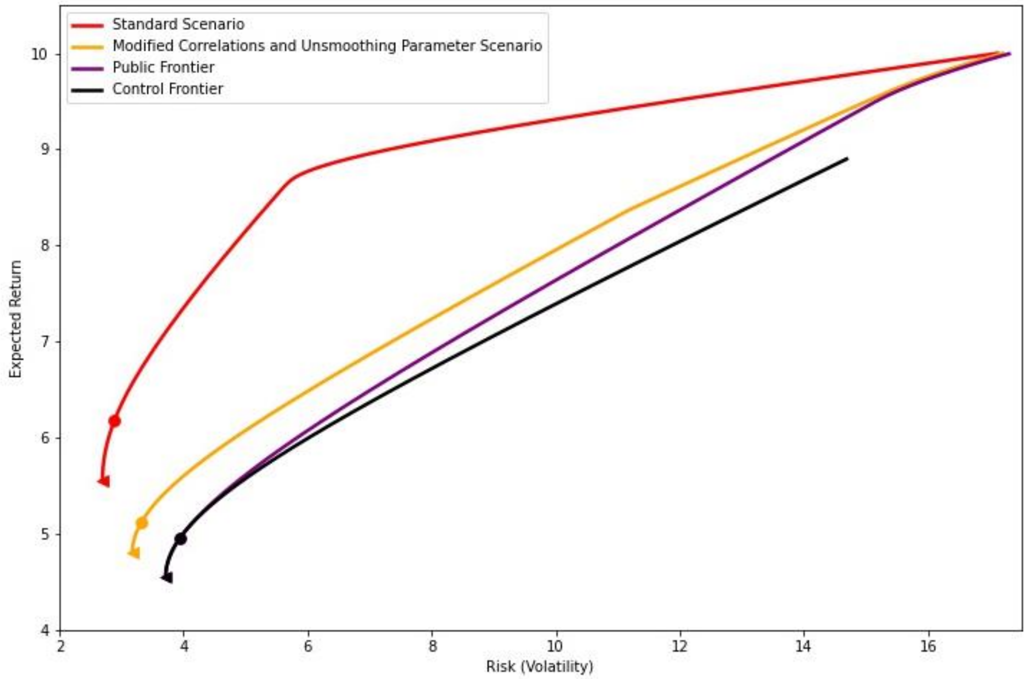

Figure 8 plots the efficient frontier, including adjusted private real estate under a scenario against two benchmarks. The lines represent the efficient frontiers, whereas the dot and the triangle represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio, respectively, in the analyzed scenario. The constraint confines each asset's weight (wi) to [0, 1]. The red line (the standard scenario) is the efficient frontier constituted of XBB, XIU, XRE with the adjusted Property Index and the adjusted PFI. The orange line represents the efficient frontier constituted of XBB, XIU, XRE with the adjusted Property Index and the adjusted PFI in which the correlation between private (Property Index and PFI) and public real estate (XRE) is 0.9 and, the correlation between private real estate and the stock market (XIU) is 0.7459. The appraiser confidence factor "α" is also changed from 0.815 to 0.4. The two benchmarks in this figure are the Public Frontier (constituted of XBB, XIU, and XRE) and the Control Frontier (comprised of XBB and XIU only).

Table 9 details the performance measures and weighted allocations for the new scenario. The details of the public frontier are reported in Table 10 with the other optimal asset allocation of individual real estate investment vehicles.

Even though every efficient portfolio includes some allocation to at least one of the real estate investment vehicles, the beneficial impacts on portfolio performance, measured by the Sharpe ratios, are now notably smaller and seem more realistic. In the adjusted efficient frontier (the standard scenario), the addition of the three real estate investment vehicles allowed the simulation of an optimal portfolio with a Sharpe ratio of 1.66 with a total allocation to real estate of 43%. Now, with the new assumptions included in the model above, the Sharpe ratio of the optimal portfolio, including our three real estate investment vehicles, has fallen to 1.13, with a total allocation to real estate of 24%.

Table 9 - Optimal Asset Allocation Sensitivity Analysis 2

Target Portfolio Return	Portfolio Volatility	Sharpe-Ratio	Weights in Optimal Portfolio				
			Property Index	PFI	XBB	XIU	XRE
4.81%	3.17%^{MV}	1.08	0.00	0.17	0.83	0.00	0.00
5.12%	3.32%	1.13^{Opt}	0.04	0.20	0.75	0.01	0.00
6.00%	4.82%	0.96	0.18	0.28	0.54	0.00	0.00
7.00%	7.35%	0.76	0.30	0.16	0.31	0.23	0.00
8.00%	10.13%	0.65	0.38	0.10	0.12	0.34	0.06
9.00%	13.30%	0.57	0.26	0.08	0.00	0.36	0.30
10.00%	17.21%	0.50	0.00	0.00	0.00	0.04	0.96

Table 9 provides the portfolio volatility, the Sharpe ratio, and the weights of the adjusted Property Index, the adjusted PFI, XBB, XIU, and XRE in optimal portfolios for varying target portfolio returns for the Modified Correlations and Unsmoothing Parameter scenario.

From the observed results, we hypothesize that the appraisal-based valuation process of private real estate partly explains the abnormal performance of private real estate in the MVO framework by significantly smoothing the variations of real estate values through time and by understating the correlation with the other asset classes. However, even after the additional assumptions to compensate for the appraisal valuation smoothing bias, our results demonstrate that real estate is still beneficial to portfolio performance in every portion of the efficient frontier.

6.3. Additional Results

The optimal allocations without real estate can be compared to the allocations when real estate investment vehicles are individually added to the control portfolio. Figure 9 presents the efficient frontiers of individual real estate investment vehicles (for the unadjusted and adjusted versions of the model) and the efficient control frontier. In contrast, Table 10 highlights the different allocations of the adjusted version of the model for the three real estate investment vehicles taken individually. The difference between the efficient frontier incorporating REITs and the many efficient frontiers comprised of private real estate is noticeable. The addition of REITs to the control portfolios generates limited diversification benefits. These benefits are primarily in the upper portion of the frontier, where the allocation to XRE gets progressively heavier.

Figure 9 - Efficient Frontiers of Individual Real Estate Investment Vehicles

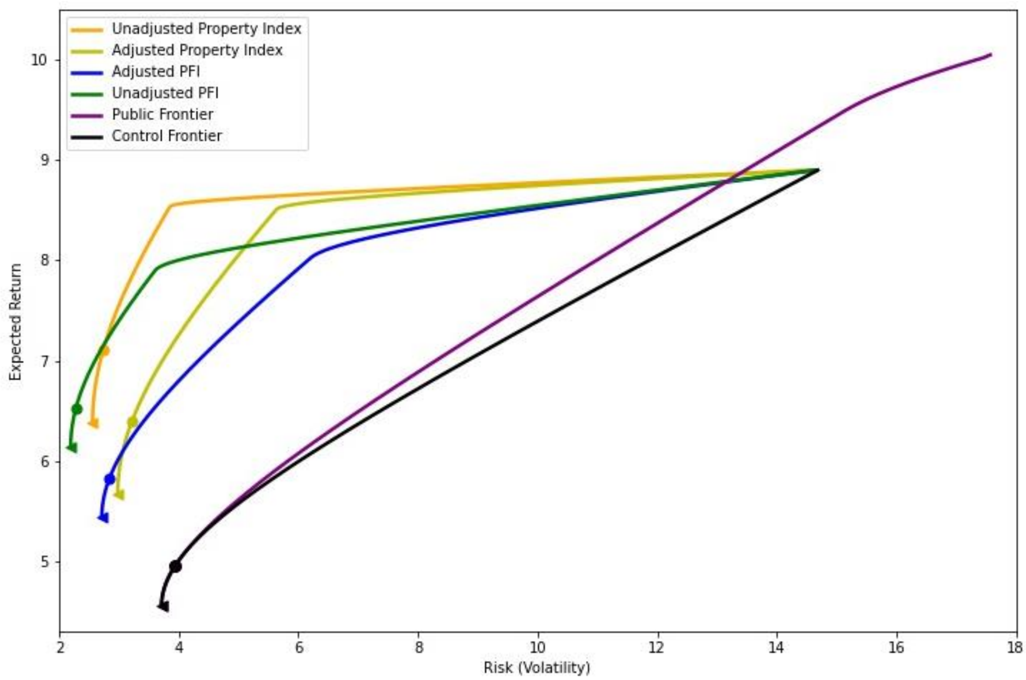

Figure 9 syntheses efficient frontiers of individual real estate investment vehicles. The black line represents the control efficient frontier constituted of XBB and XIU. The purple line represents the efficient public frontier constituted of XBB, XIU, and XRE (the Public Frontier). The blue line represents the efficient frontier constituted of XBB, XIU, XRE with the adjusted PFI. The yellow line represents the efficient frontier constituted of XBB, XIU, XRE with the adjusted Property Index. The green line represents the efficient frontier constituted of XBB, XIU, XRE with the PFI. The orange line represents the efficient frontier comprised of XBB, XIU, XRE with the unadjusted Property Index. The dot and triangle on each frontier represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio, respectively.

The addition of REITs allows the efficient frontier to reach higher in terms of returns, but at the expense of adding more volatility. On the other hand, integrating some form of private real estate into the control portfolios seems to generate massive diversification benefits. The Property Index and the PFI allowed the portfolios to generate more returns for less risk. The shape of these frontiers is consistent with what Ross, Westerfield, Jaffee, and Robert (2008) illustrated when the correlations between the asset classes are very low. The very low correlations between the Property Index and the PFI with XBB or XIU noticeably increased portfolios' performance on a Sharpe ratio basis. From a diversification standpoint, private real estate seems more beneficial to portfolios than its public counterpart, the REITs. This is not surprising when the high correlation coefficient between XIU and XRE is considered. Figure 9 also highlights the adjustment's impact on the Property Index and the PFI individually, which are consistent with what we previously analyzed in the mixed-asset models. Table 10 contrasts the three optimal portfolios.

Table 10 - Optimal Asset Allocation of Individual Real Estate Vehicles

Panel A: Adjusted Property Index with XBB and XIU

Target Portfolio Return	Portfolio Volatility	Sharpe-Ratio	Weights in Optimal Portfolio: Property Index	XBB	XIU
5.68%	2.98%^{MV}	1.44	0.30	0.65	0.05
6.40%	3.22%	1.56^{Opt}	0.45	0.48	0.07
7.00%	3.73%	1.51	0.57	0.34	0.09
8.00%	4.93%	1.34	0.76	0.12	0.12
8.90%	14.69%	0.51	0.00	0.00	1.00

Panel B: Adjusted PFI with XBB and XIU

Target Portfolio Return	Portfolio Volatility	Sharpe-Ratio	Weights in Optimal Portfolio: PFI	XBB	XIU
5.44%	2.71%^{MV}	1.50	0.32	0.66	0.02
5.82%	2.84%	1.57^{Opt}	0.40	0.56	0.04
6.00%	2.97%	1.55	0.43	0.51	0.06
7.00%	4.33%	1.30	0.64	0.26	0.10
8.00%	6.16%	1.08	0.84	0.01	0.15
8.90%	14.69%	0.51	0.00	0.00	1.00

Panel C: XRE with XBB and XIU (Public Frontier)

Target Portfolio Return	Portfolio Volatility	Sharpe-Ratio	Weights in Optimal Portfolio: XRE	XBB	XIU
4.55%	3.71%^{MV}	0.86	0.00	0.91	0.09
4.96%	3.94%	0.91^{Opt}	0.00	0.82	0.18
6.00%	5.83%	0.79	0.12	0.63	0.25
7.00%	8.30%	0.68	0.23	0.45	0.32
8.00%	10.98%	0.60	0.34	0.27	0.39
9.00%	13.76%	0.55	0.46	0.09	0.45
10.00%	17.34%	0.50	0.96	0.00	0.04

Table 10 Panel A, B, and C provide the portfolio volatility, the Sharpe ratio, and the weights of the adjusted Property Index, the PFI, and XRE, respectively individually, with XBB and XIU in optimal portfolios for varying target portfolio returns. The "MV" superscript identifies the row of the minimum variance portfolio, and the "Opt" superscript identifies the portfolio's row with the highest Sharpe ratio.

From Panel A, the optimal portfolio of the efficient frontier constituted of the adjusted Property Index with XBB, and XIU has a 1.56 Sharpe ratio, which is a bit lower than the adjusted PFI with XBB and XIU in panel B at a 1.57 Sharpe ratio. Apart from this discrepancy, the two portfolios differ because the addition of the adjusted Property Index allows for more return and more volatility than the addition of the adjusted PFI. The allocation to real estate is also heavier in Panel A at 45% versus 40% in Panel B. In both cases, most of the weight is allocated to XBB with very little weight on XIU. In Panel C, the optimal portfolio of the efficient frontier integrating XRE with XBB and XIU has the same allocation and Sharpe ratio as the optimal portfolio of the control efficient frontier since no weight is allocated to XRE in the lower portion of the frontier because of its high volatility and high correlation with XIU.

The analysis of portfolio optimization with individual real estate investment vehicles confirms our previous statement from the primary results. For more risk-averse investors, the addition of unlevered private real estate properties or core style PERE funds is preferable to REITs. REITs seem more appropriate to investors seeking higher returns and are more risk tolerant.

7. Conclusion

The increasing popularity of alternative assets, in particular real estate, amongst the large Canadian pension funds necessitates a more thorough examination of the portfolio-enhancing characteristics of these investments' vehicles and their interactions with the traditional asset classes. This article examines real estate's role in Canadian pension funds' mixed-asset portfolios using private directly held real estate properties, private equity real estate funds, and publicly traded equity REITs, making it the first to integrate PERE funds in addition to REITs and direct properties in a portfolio optimization setting. Moreover, combining the corrections for the appraisal smoothing bias and the liquidity risk is a unique feature of this study. The article also complements and updates the scarce literature on portfolio optimization with real estate in the Canadian market.

As usual, when using historical data to guide investment decisions, one must be careful in assigning too much importance to the exact numbers that come out of the analysis since there is no guarantee that the future will resemble the past. However, several general conclusions come out from the results of the analysis.

In the first test, when we add unadjusted real estate investment vehicles to the control frontier, the Sharpe ratio of the optimal portfolio increases significantly (from 0.91 to 2.33 in the case of the unadjusted model). The increase is perhaps too substantial and possibly unrealistic. Portfolio performance seems to be overstated and highlights the need to adjust private real estate data. We then compare the unadjusted model to the adjusted model in which the Property Index and the PFI's data are adjusted to account for the effect of appraisal smoothing and the liquidity risk. The resulting optimal portfolio from this latter model shows a 1.66 Sharpe ratio, demonstrating that even after correcting for these biases, real estate remains a desirable feature of a well-diversified portfolio, although in lower allocation than before the adjustments. Drawing from the results of the adjusted model, we conclude that for the investors who prefer low-risk portfolios, the addition of unlevered private real estate properties or core style PERE funds is the optimal investment vehicle to improve the performance of a stock and investment-grade bond portfolio. Conversely, for those who prefer high-risk portfolios, the addition of public REITs seems to be the right choice to serve this risk-aversion preference.

In every model tested, the allocation to real estate is high, mainly because of the lower volatility and low correlation coefficients with the other asset classes. Allocation to real estate around 40% or more seems too high for pension funds, but interestingly, it approaches a type of asset allocation frequently observed for Canadian households (homeowners). The robustness tests allowed us to analyze how the allocation to real estate would change if the smoothing effect from the appraisal valuation process underestimated the volatility and the correlation measures significantly more than our adjustments accounted for. In the strictest scenario tested (the correlation parity scenario with enhanced unsmoothing parameter), the optimal portfolio yields a 1.13 Sharpe ratio with an allocation to real estate slightly over 20%. This result demonstrates that real estate is still a portfolio performance enhancer under more demanding assumptions.

Overall, the diversification benefits of including Canadian commercial real estate investment in mixed-asset portfolios are clear. Although private real estate investment vehicles present higher liquidity risks than traditional asset classes and the valuation issues also impact the reliability of the volatility and correlation measures. The pension funds that are willing to hurdle these obstacles have the opportunity to achieve better long-run performance. The most promising avenues of research for real estate portfolio optimization would be to replace core style PERE funds with more aggressive funds, like value-add and opportunistic style funds. Otherwise, the use of long-horizon measures of risk and return seems promising as an alternative method to counter the appraisal smoothing problem. Although, this method would require data on a longer time frame than what is available to us in this study.

8. Bibliography

1. Alex S. MacNevin, 2013. “Comparative Tax Advantages of Canadian Pension Funds as Investors in Real Estate,” *Canadian tax journal*, vol. 61, no. 1, p. 41-78.
2. Andre, Kuzmicki and D. Simunac, 2008. “Private Equity Real Estate Funds: An Institutional Perspective,” *RealPac paper*, p. 1-20.
3. Bond, Shaun A., S. Hwang, and K. Richards, 2006. “Optimal Allocation to Real Estate Incorporating Illiquidity Risk,” *Journal of Asset Management*, vol 7, no. 1, p. 2-16.
4. Cheng, P., Z. Lin, and Y. Liu, 2013. “Liquidity Risk of Private Assets: Evidence from Real Estate Markets,” *The Financial Review*, vol. 48, p. 671-696.
5. Chmiel, J. and J. Rodrigues, 2019. “The Canadian REIT Almanac”, Raymond James.
6. Clayton, J., D. Geltner and S.W. Hamilton 2001. “Smoothing in Commercial Property Valuations: Evidence from Individual Appraisals,” *Real Estate Economics*, vol. 29, no. 3, p. 337- 360.
7. Clayton, J, and G. MacKinnon, 2002. “Departures from NAV In REIT Pricing: The Private Real Estate Cycle, The Value of Liquidity and Investor Sentiment.” Working Paper.
8. Ciochetti, Brian A., T.M. Craft, and J.D. Shilling, 2002. “Institutional Investors' Preferences for REIT Stocks,” *Real Estate Economics*, vol. 30, no. 4, p. 567-593.
9. Conner, P., and Y. Liang, 2004. “The complex interaction between real estate cap rates and interest rates,” *Briefings in Real Estate Finance*, vol 4, no 3, p. 185-197.
10. Felista, Mutahi and O. Ferdinand, 2015. “Real Estate Investment Trusts versus Direct Real Estate Investments: A Portfolio Optimization Approach,” *Journal of Finance and Investment Analysis*, vol. 4, no. 4, p. 57-75.
11. Firstenberg, P.M., S.A. Ross and R.C. Zisler, 1988. “Real Estate: The Whole Story,” *Journal of Portfolio Management*, vol. 14, no. 3, p. 22-34.
12. Geltner, D., 1993. “Estimating market Values from Appraised Values Without Assuming an Efficient Market,” *Journal of Real Estate Research*, vol. 8, no. 3, p. 325-346.
13. Geltner, D. and W. Goetzmann, 2000. “Two Decades of Commercial Property Returns: A Repeated-Measures Regression-Based Version of the NCREIF Index,” *Journal of Real Estate Finance and Economics*, vol. 21, p. 5–21.
14. He, Lin T, J.R. Webb, and N. Myer, 2003. “Interest Rate Sensitivities of REIT Returns,” *International Real Estate Review*, vol. 6, no. 1, p. 1-21.
15. Hudson-Wilson, Susan, F.J. Fabozzi, and J. Gordon, 2003. “Why Real Estate?”, *Journal of Portfolio Management*, Special Real Estate Issue.
16. Hudson-Wilson, Susan, F.J. Fabozzi, J. Gordon, M. Anson, and S.M. Giliberto, 2005. “Why Real Estate?”, *Journal of Portfolio Management,* Special Real Estate Issue, vol. 31, no. 5, p. 12-21.
17. Ibbotson, R. and L.B. Siegal, 1984. “Real Estate Returns: A Comparison with Other Investments,” *AREUEA Journal*, vol. 12, p. 219-242.

18. Ilkka Tomperi, 2010. “Performance of private equity real estate funds,” *Journal of European Real Estate Research*, vol. 3, no. 2, p. 96-116.
19. Joseph L. Pagliari Jr., 2017. “Another Take on Real Estate's Role in Mixed-Asset Portfolio Allocations,” *Real Estate Economics*, vol. 45, no. 1, p. 75-132.
20. Kieran, Farrelly, and S. Simon, 2019. “The risk and return of private equity real estate funds,” *Global Finance Journal*, vol. 42, p. 1-13.
21. Lin, Z., and K. Vandell, 2007. “Illiquidity and pricing biases in the real estate market,” *Real Estate Economics*, vol. 35, p. 291–330.
22. Markowitz, H. 1952. “Portfolio Selection,” *Journal of Finance*, vol. 7, p. 77–91.
23. Òscar Jordà, M. Schularick, and A.M. Taylor, 2019. “The Total Risk Premium Puzzle,” *NBER Working Paper*, no. 25653.
24. Òscar Jordà, K. Knoll, D. Kuvshinov, M. Schularick, and A.M. Taylor, 2019. “The Rate of Return on Everything, 1870–2015”, *The Quarterly Journal of Economics*, vol. 134, Issue 3, p. 1225–1298.
25. Pierre Vachichin, 2019. “The diversification effects of the inclusion of real estate assets in a Canadian investment portfolio: A 2010 to 2018 perspective”, Master of Science in Finance, Simon Fraser University.
26. Quan, D.C., and J.M. Quigley, 1991. “Price Formation and the Appraisal Function in Real Estate Markets,” Journal of Real Estate Finance and Economics, vol. 4, no. 2, p. 181-202.
27. Ross, Westerfield, Jaffee, and Robert, 2008. “Corporate Finance,” McGrawHill Ryerson, 5th Canadian edition, page 302.
28. Shiller, Robert J. 2014. “Speculative Asset Prices,” *American Economic Review*, vol. 104, no. 6, p. 1486-1517.
29. Stefek, D. and R. Suryanarayanan, 2012. “Private and Public Real Estate: What Is the Link?”, *Journal of Alternative Investments*, vol. 14, no. 3, p. 66-76.
30. Thomas R. Arnold, Ling, C. David, and A. Naranjo, 2019. “Private Equity Real Estate Funds: Returns, Risk Exposures, and Persistence,” *Journal of Portfolio Management Real Estate Special Issue*, p. 24-42.
31. Thomas R. Arnold, Ling, C. David, and A. Naranjo, 2021. “Private Equity Real Estate Fund Performance: A Comparison to REITs and Open-End Core Funds,” *Journal of Portfolio Management*, Real Estate 2021, p. 1-20.
32. Xiong, James X., and T.M. Idzorek, 2011. “The Impact of Skewness and Fat Tails on the Asset Allocation Decision.”, *Financial Analysts Journal*, vol. 67, no. 2, p. 23–35.

9. Appendices

Appendix A: RBC Canadian Defined Benefit Pension Survey 2021

Figure 10 - 2019 Canadian Defined Benefit Pension Plans Asset Allocation

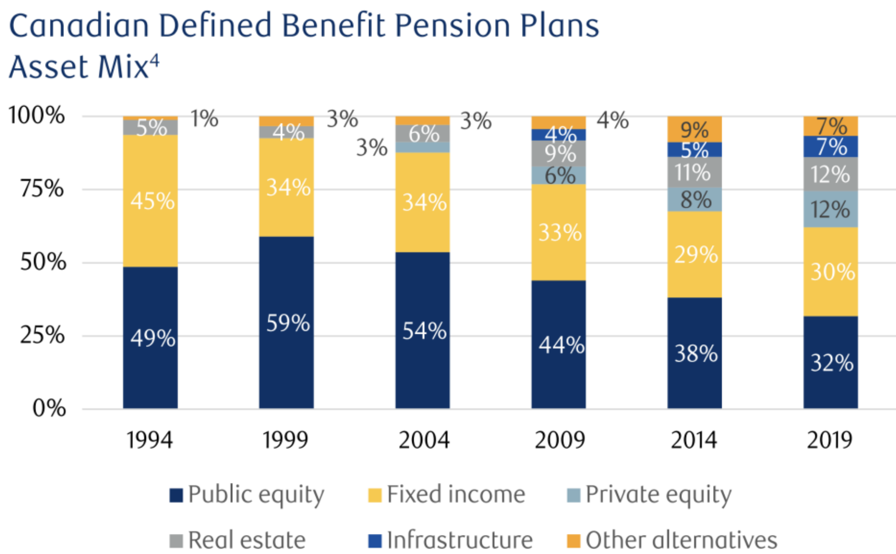

Figure 10 presented by RBC Investors & Treasury Services. "Pathways to Sustainability: Canadian Defined Benefit Pension Survey 2021". RBC Investor & Treasury Services' fourth annual survey of Canadian defined benefit pension plans, conducted in November of 2020, reflects the perspectives for 2021 of 122 pension plans from across the country.

Appendix B: Unsmoothing Appraisal Based Private Real Estate Returns

The MSCI/REALPAC Canada Quarterly Property Index and the MSCI/REALPAC Canada Quarterly Property Fund Index publish income, capital, and total returns to Canadian commercial real estate properties and funds. As smoothing applies to the capital value of properties rather than the total return, the following equation is used to unsmooth the capital returns from the series:

$$R_t = \frac{R_{App,t} - (1 - 0.815) \times R_{App,t-1}}{0.815}$$

To account for this effect, the following procedure is employed:

1. Using the quarterly capital growth return provided by MSCI, an index is constructed, starting with an arbitrary value of 100.
2. We calculate a variable representing the dollar value of income based on the value index for each quarter. The representation of income in dollar form is calculated by multiplying the property value index from above by that quarter’s income return percentage.
3. An unsmoothed index of property values is constructed using the equation above on the value index created in step 1 to generate a time series representing unsmoothed property values.
4. We divide the dollar representation of income by the unsmoothed value index for each quarter to recreate an accurate series of the pro-rata income returns.
5. To complete the procedure, the unsmoothed capital return is added to the income return to yield an unsmoothed total return series for each quarter.

Appendix C: Types and Styles of PERE Funds

Table 11 - Types and Styles of PERE Funds

FIGURE 1 – Types of Private Equity Real Estate Funds[6]		
Type of Fund	**Target Net IRRs**	**Typical Characteristics**
Core	7–9%	• Well-diversified, low risk/return strategy • Traditional asset classes (i.e. office, retail, industrial, multi-family) in established locations • Well occupied and well maintained assets with stable cash flows • Generally little or no debt is employed (i.e. up to 30% levered) • Income stream represents significant part of expected total return
Core-Plus	9–12%	• Moderate risk/return strategy • Core-type assets, some of which may require some form of value-add enhancement • Assets located in either primary or secondary locations • Moderate amount of leverage employed (i.e. up to 55% levered)
Value-Add	12–16%	• Moderate-to-high risk/return strategy • Opportunity to add value through operating, re-leasing, and/or re-development • Leverage employed is higher (i.e. up to 70% levered) • Value appreciation comprises significant part of expected total return
Opportunistic	>16%	• High risk/return strategy • Re-positioning of poorly managed, obsolete, and/or vacant assets • Acquisitions of entire companies with portfolios of assets and operating platforms in-place • New-build development or conversion projects • International focus pursing opportunities in established and emerging markets • Leverage employed is higher (i.e. over 70% levered)
Mezzanine Lending / Distressed Debt	9–12%	• Originate loans at terms which are more aggressive than traditional lenders • May contain profit-sharing in addition to higher interest rates and fees charged • Acquire distressed loans from lenders at discount prices (i.e. below par value) • Participate in higher-yielding tranches of mortgages (i.e. non-rated CMBS; first-loss positions) • Generally not averse to owning the assets in the event such loans default
Fund-of-Funds	12–16%	• Funds which invest in a number of third-party managed PERE funds • Aggregate capital of smaller investors and invest such capital on their behalf • Provide smaller investors access to larger, more exclusive PERE funds • Offer a diversified investment strategy through a single channel

Table 11 describes each PERE fund style. From Andre, Kuzmicki, and D. Simunac, 2008. "Private Equity Real Estate Funds: An Institutional Perspective," RealPac paper, p. 6.

Appendix D: Liquidity Bias Factors

Table 12 - Liquidity Bias Factors

Table 3 ■ Liquidity bias between *ex ante* and *ex post* return estimates: The U.S. commercial property market.

Expected Marketing Period (Months)	The Ratio of *Ex Ante* Variance to *Ex Post* Variance				
	Holding Period				
	0	1 Year	5 Years	10 Years	20 Years
0	1.00	1.00	1.00	1.00	1.00
2	2.21	1.17	1.04	1.02	1.01
4	3.42	1.61	1.15	1.08	1.04
6	4.64	2.21	1.33	1.17	1.09
8	5.85	2.94	1.57	1.30	1.16
10	7.06	3.75	1.87	1.47	1.24
12	8.27	4.64	2.21	1.66	1.35
14	9.49	5.57	2.61	1.89	1.47
16	10.70	6.54	3.04	2.14	1.61
18	11.91	7.55	3.52	2.42	1.76
20	13.12	8.58	4.03	2.73	1.93
24	15.55	10.70	5.16	3.42	2.32

Source: NCREIF property performance index, 1980Q1 to 2004Q4.
Note: The table above shows by how much the variance of the returns, after taking the uncertainty of marketing period into account, is greater than that given by traditional estimation methods using the unadjusted sales price as an unbiased estimate of market value. For example, if an investor holds a property for ten years and the expected marketing period is eight months, the *ex ante* risk faced by the investor is 30% higher than that given by traditional *ex post* estimation.

Table 12 provides the appropriate adjustments of the variance for the liquidity risk. From Lin, Z., and K. Vandell, 2007. "Illiquidity and pricing biases in the real estate market," Real Estate Economics, vol. 35, p. 310.

Appendix E: Robustness Tests Individual Efficient Frontiers

Figure 11 - Efficient Frontier of the Pessimistic Scenario

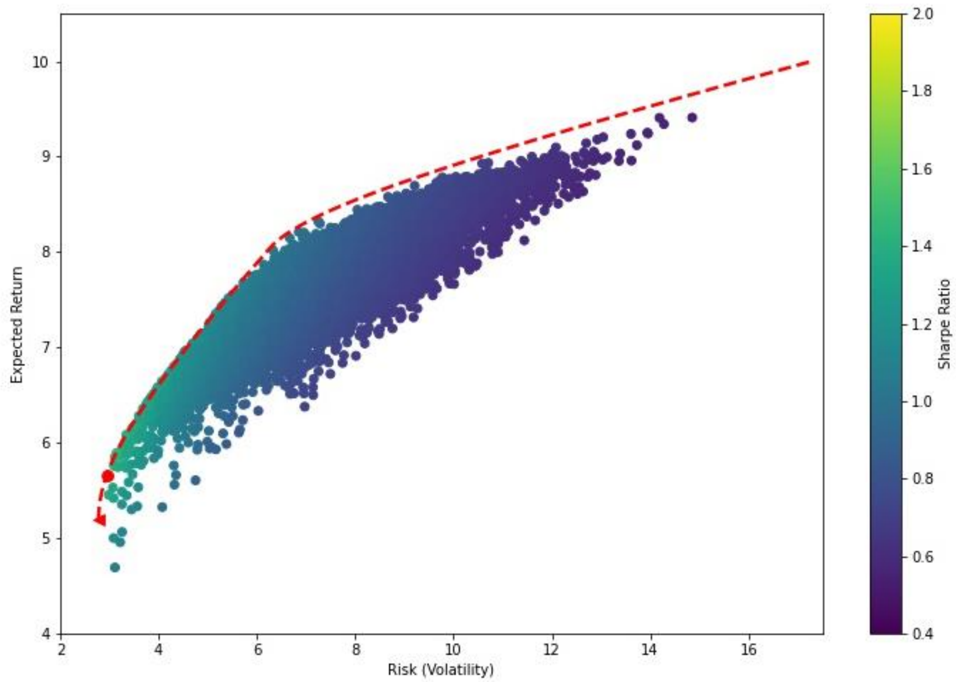

Figure 11 plots the optimized, efficient frontier constituted of XBB, XIU, XRE with the adjusted Property Index and the adjusted PFI. The Property Index and the PFI had their mean return reduced by 10%, and their volatility increased by 10%. The constraint confines each asset's weight (w_i) to [0, 1]. The red dotted line represents the efficient frontier, whereas the red dot and the red triangle represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio. This plot also illustrates the 100 000 simulated portfolios for the selected asset classes. The portfolios are ranked by Sharpe ratio: the darker colors (purple and blue) represent lower Sharpe ratios, and lighter colors (green and yellow) represent higher Sharpe ratios.

Figure 12 - Efficient Frontier of the Catastrophic Scenario

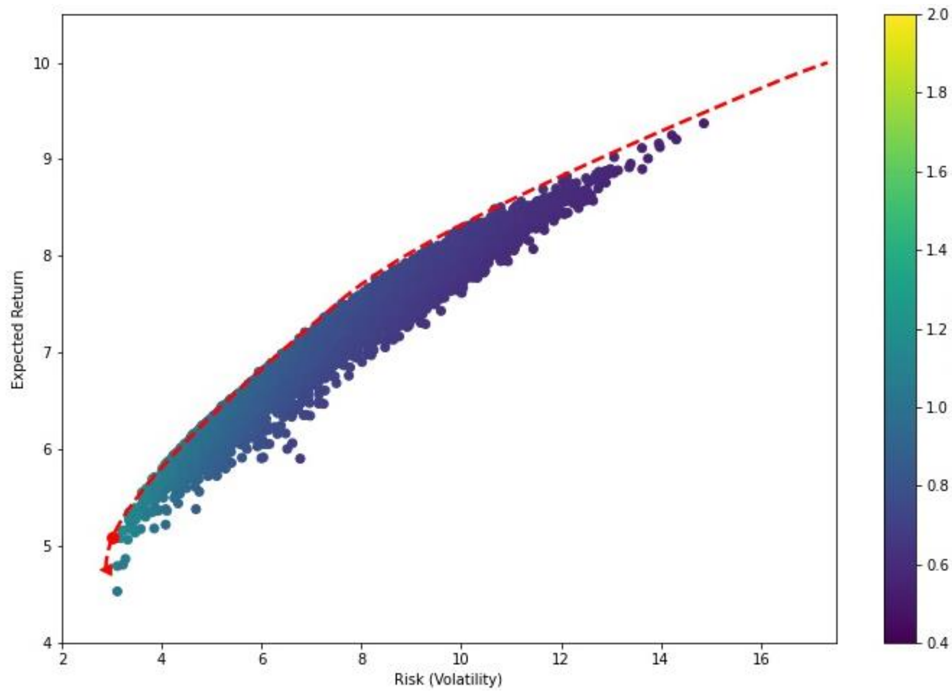

Figure 12 plots the optimized, efficient frontier constituted of XBB, XIU, XRE with the adjusted Property Index and the adjusted PFI. The Property Index and the PFI had their mean return reduced by 25%, and their volatility increased by 25%. The constraint confines each asset's weight (w_i) to [0, 1]. The red dotted line represents the efficient frontier, whereas the red dot and the red triangle represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio. This plot also illustrates the 100 000 simulated portfolios for the selected asset classes. The portfolios are ranked by Sharpe ratio: the darker colors (purple and blue) represent lower Sharpe ratios, and lighter colors (green and yellow) represent higher Sharpe ratios.

Figure 13 - Efficient Frontier with Modified Correlations and Unsmoothing Parameter

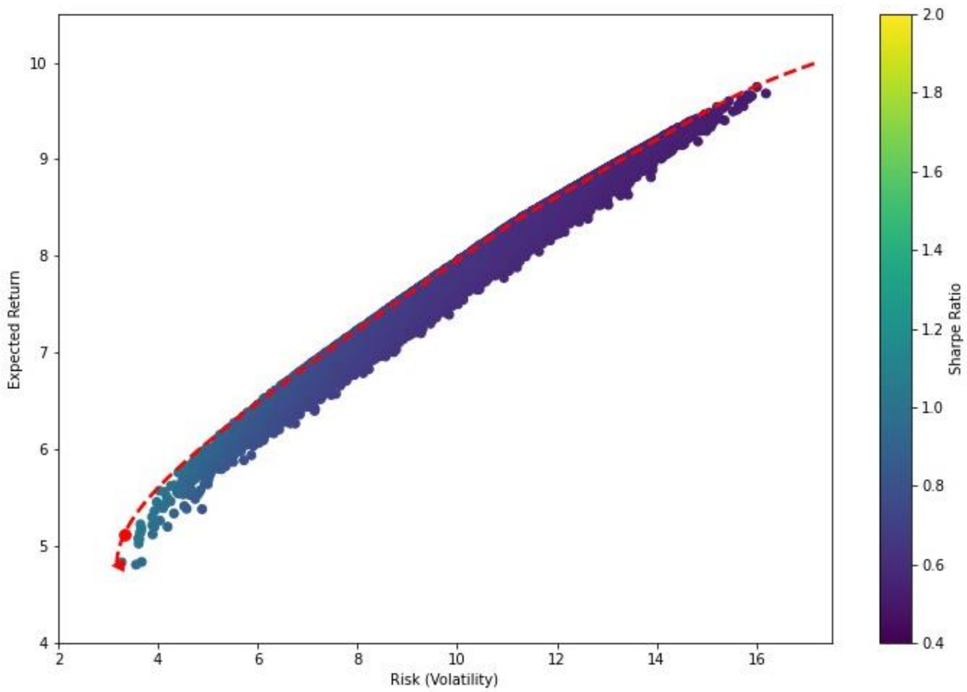

Figure 13 plots the optimized, efficient frontier constituted of XBB, XIU, XRE with the adjusted Property Index and the adjusted PFI. The correlation between the Property Index and the PFI with XIU and XRE has been fixed at 0.7459 and 0.9, respectively. The "unsmoothing" factor "α" has also been increased from 0.185 to 0.6. The constraint confines each asset's weight (w_i) to [0, 1]. The red dotted line represents the efficient frontier, whereas the red dot and the red triangle represent the optimal portfolio (maximum Sharpe ratio) and the minimum variance portfolio. This plot also illustrates the 100 000 simulated portfolios for the selected asset classes. The portfolios are ranked by Sharpe ratio: the darker colors (purple and blue) represent lower Sharpe ratios, and lighter colors (green and yellow) represent higher Sharpe ratios.