The Canadian Pension Fund Model: A Quantitative Portrait

Alexander D. Beath, Sebastien Betermier, Chris Flynn, and Quentin Spehner

Alexander D. Beath
is a senior research analyst at CEM Benchmarking in Toronto, ON, Canada.
alex@cembenchmarking.com

Sebastien Betermier
is an associate professor of finance at the Desautels Faculty of Management at McGill University in Montreal, QC, Canada.
sebastien.betermier@mcgill.ca

Chris Flynn
is a principal at CEM Benchmarking in Toronto, ON, Canada.
chris@cembenchmarking.com

Quentin Spehner
is an analyst at CEM Benchmarking in Toronto, ON, Canada.
quentin@cembenchmarking.com

KEY FINDINGS

- From 2004 to 2018, Canada's pension funds outperformed their peers in terms of investment performance and insurance against liability risks.
- The Canadian model is cost efficient, not low cost. Canadian funds reduce costs by managing assets in-house and then redeploy resources by growing their internal capabilities and allocating more capital to strategic assets.
- The Canadian model works best for funds whose pension liabilities are indexed to inflation.

ABSTRACT

This article presents a quantitative portrait of the Canadian pension fund model. The authors show that, between 2004 and 2018, Canadian pension funds outperformed their international peers in terms of both asset performance and liability hedging. A central factor driving this success is the implementation of a three-pillar business model that consists of (1) managing assets in-house to reduce costs, (2) redeploying resources to internal investment teams for each asset class, and (3) channeling capital toward growth assets that increase portfolio efficiency and hedge liability risks. This model works best for funds whose pension liabilities are indexed to inflation.

TOPICS

Pension funds, portfolio theory, risk management, performance measurement*

*All articles are now categorized by topics and subtopics. **View at PM-Research.com**.

Over the past decades, a growing body of literature has brought to light the distinct features that characterize Canada's large public pension funds (see *The Economist* 2012, Ambachtsheer 2016, Bédard-Pagé et al. 2016). In a collection of insightful case studies, the World Bank (2017) described the Canadian model as the combination of independent governance, professional in-house management, scale, and extensive geographic and asset-class diversification. These features all contributed to the strong performance of Canadian funds over the past decades and in turn allowed the pension plans to remain well funded in spite of decreasing interest rates and increasing life expectancy.

The present study builds on this literature and presents a quantitative portrait of the Canadian pension fund model. We aim to answer several questions. How successful has the Canadian model been over the past two decades in terms of both asset performance and liability hedging? Which features of this model have contributed

the most to its success? Is the Canadian model restricted to a few flagship funds, or has it spread across the country? Can institutional funds borrow features from this model even if they are not in a position to adopt it entirely?

To conduct this analysis, we use data from CEM Benchmarking, a global benchmarking company, and analyze performance metrics, asset allocation strategies, and cost structures for 250 pension, endowment, and sovereign wealth funds across 11 countries. We split the sample into large and small funds depending on whether their assets under management exceeded $10 billion in 2018 and quantitatively analyze the distinct features of Canadian funds.

We first examine the large funds and show that, between 2004 and 2018, Canadian funds outperformed their peers on all fronts. Not only did they generate greater returns for each unit of volatility risk, they also did a superior job of hedging their pension liability risks. The ability to deliver both high return performance and insurance against liability risks is notable because hedging is typically perceived as a cost. Our results are in line with those of Ambachtsheer (2017a, 2017b), who found that Canadian funds outperformed non-Canadian funds on a risk-adjusted basis from 2006 to 2015 (see Ambachtsteer (2021) for a good review).

We then examine the factors that drive the success of large Canadian funds. We identify three distinct features that characterize their asset allocation and cost structure. The first feature is well known: Canadian funds manage on average 52% of their assets in-house, whereas non-Canadian funds only manage 23% of their assets internally. For the very large funds that manage more than $50 billion, the gap is even more pronounced: Canadian funds manage on average 80% of their assets in-house, whereas non-Canadian funds only manage 34% internally. We estimate that, by managing a high proportion of their assets in-house, Canadian funds reduce costs by approximately one-third.

The second distinctive feature of large Canadian funds is the redeployment of resources to internal investment teams for each asset class. Canadian funds use the cost savings that result from the transition to in-house management to grow their internal capabilities. We show that they outspend their peers on internally managed portfolios (18 bps on average versus 7 bps for their peers). Examples of expenses include risk management units and IT infrastructure, where Canadian funds outspend their peers by a factor of 5. In addition, even though they manage fewer assets externally, Canadian funds outspend their peers on external managers (121 bps versus 86 bps). These patterns hold true within each asset class and style.

The third distinctive feature of large Canadian funds is the allocation of capital toward assets that increase portfolio efficiency and hedge against liability risks. These assets include commodity producer stocks, real estate, and infrastructure. We find that, with the cost savings that result from the internal management of assets, Canadian funds are able to allocate 18% of their assets under management (AUM) to real assets, which tend to be more expensive to manage than stocks and bonds. In comparison, non-Canadian funds allocate just 9% of their AUM to real assets.

Altogether, these findings reveal a three-pillar business model that goes beyond a simple combination of professional in-house management, scale, and asset diversification. The Canadian model is an integrated and cost-efficient model that enables large Canadian funds to spend more on each asset class and allocate more capital to strategic assets while spending less than their peers. On average, Canadian funds spend 57 bps of their AUM each year to run their fund, whereas non-Canadian funds spend 62 bps.

We quantify the impact of these features on fund performance by running a statistical analysis in which we regress several performance metrics on a set of predictor variables. The performance metrics include the Sharpe ratio of the funds' assets,

which reflects the risk–return efficiency of their investments, and the Sharpe ratio of the funds' net portfolio, which is long their assets and short their liabilities. We also consider the correlation between the funds' assets and liabilities to evaluate the amount of liability hedging.

We find that the greater proportion of assets managed in-house significantly improves the funds' Sharpe ratios. Moreover, even though multiple asset classes positively contribute to the funds' asset performance, two asset classes play a leading role in improving both the asset–liability Sharpe ratio and asset–liability correlation: fixed income and real assets, which make up a large share of Canadian fund portfolios. Both asset classes have cash flow characteristics that are similar to the pension funds' liability profiles and therefore allow the funds to earn a risk premium while aligning the risk of their liability portfolio to their assets.

Another factor that contributes to the strong asset–liability performance of large Canadian funds is the fact that a high proportion of their pension liabilities is indexed to inflation. The explanation for this effect is that indexed liabilities tend to correlate more with growth assets than with nominal liabilities. Consequently, by having indexed liabilities, Canadian funds are able to invest in growth assets that improve both return performance and liability hedging.

We also run a series of back tests to study how US pension funds would have performed between 2004 and 2018 by adopting features of the Canadian model. These counterfactual exercises allow us to fully control for environmental factors such as the performance of Canadian real estate and fixed income. The US environment is fitting for this analysis because (1) it includes the majority of the funds in the sample and (2) it provides the opportunity to work with both the US public and corporate pension fund models.

We find that, for US public funds, the adoption of the Canadian model would have led to a 15% absolute increase in the 15-year Sharpe ratio of the asset portfolio, a 13% increase in the 15-year Sharpe ratio of the asset–liability portfolio, and a 20% increase in the correlation between assets and liabilities. These estimates do not include any additional performance resulting from the Canadian funds' decision to spend more in each asset class. For US corporate funds, which already hedge a high proportion of their liability risks, the adoption of the Canadian model would have also led to increases in all performance metrics, but mostly in the Sharpe ratios through the reduction of costs associated with in-house management.

We then turn our attention to smaller Canadian funds and observe similar patterns. Small Canadian funds outperformed their international peers on all fronts between 2004 and 2018. Their business model also shares similarities: Small Canadian funds invest more of their assets in-house than their peers (13% versus 3% on average), allocate more capital to real assets (10% versus 7%), and redeploy more resources toward active strategies (82% versus 72%). These findings indicate that the success of the Canadian model among flagship funds has trickled down to a wide range of smaller funds.

However, the differences in allocation and costs between small and large Canadian funds reveal that there is no uniform Canadian model. Because of scale constraints, small funds only adopt a light version of the three-pillar model described. We find that small Canadian funds make a number of adjustments elsewhere: They reduce costs altogether, do more internal active management in public markets that are more accessible than private markets, invest less in hedge funds, and concentrate the bulk of in-house management in one asset class—fixed income.

Our article contributes to the research literature on the Canadian pension fund model. Prior work has emphasized other aspects of this model, such as its independent governance, risk-sharing design, and legal framework, scale, diversification, and ability to make use of short-term leverage through repo and derivative markets

(Ambachtsheer 2016; Bédard-Pagé et al. 2016; World Bank 2017; Lipshitz and Walter 2020). Here, we go a step further in analyzing the business model implemented by Canadian pension funds and show that it leads to a more cost-efficient management of pension assets and greater net performance.

Our article departs from prior research regarding the choice of performance metrics for pension funds. Prior research has focused on standard metrics of performance: risk-adjusted excess return metrics (see Coggin, Fabozzi, and Rahman 1993; Blake, Lehman, and Timmerman 1999). Although these metrics effectively capture the ability of a fund to outperform a particular benchmark, they do not necessarily reflect the fund's ability to meet its pension liabilities (Ambachtsheer 2017b). Building on the asset-liability literature (Sharpe and Tint 1990; Ezra 1991; Ang, Chen, and Sundaresan 2013), we use metrics that specifically reflect the risk–return performance of a pension fund's asset–liability management.

Our article also brings a new perspective to the question of whether pension plans should index their pensions to inflation. The challenge with nominal liabilities is that, even though their risks can be hedged by investing in long-term nominal bonds and duration overlays, these investments must be large and thus leave little room for growth assets unless substantial leverage is employed. These investments also yield low returns and therefore require large contributions from the fund sponsors to keep the plans afloat. In contrast, we show that indexed liabilities make it possible for pension funds to hedge their liability risks by owning a diversified basket of growth assets such as commodities, infrastructure, and real estate. The hedging properties of these assets are consistent with the findings of Hoevenaars et al. (2008), Amenc, Martellini, and Ziemann (2009), Lipshitz and Walter (2019), and Barras and Betermier (2020).

The rest of the article is organized as follows. We first describe the data and variable construction and then compare the return and hedging performance of Canadian and non-Canadian funds. We then characterize the defining features of the Canadian pension fund model and evaluate the impact of these features on fund performance. Finally, we discuss lessons learned from the Canadian model for other investors and regulators and conclude. The online appendix provides additional details about the construction of variables and presents additional empirical results.

DATA AND VARIABLE CONSTRUCTION

Sample

We use data from CEM Benchmarking, a Toronto-based global benchmarking company that has collected detailed annual cost and performance data from more than 1,000 pension, endowment, and sovereign wealth funds in 18 countries. There are 326 funds in 2018, the last year of our sample. In total, these funds managed slightly more than $10 trillion of assets in 2018.

We run the analysis over two time periods: 5 years (2014–2018) and 15 years (2004–2018). For each time period, we exclude funds that do not appear in the data for one or more years. The five-year sample includes 250 funds from 11 countries. These funds represent 77% of all funds in 2018 and 93% of their aggregate market value. The 15-year sample includes 105 funds, which represent 32% of all funds in 2018 and 46% of their aggregate market value. In the online appendix, we run additional analysis on a 20-year sample (1999–2018) and show that our main results are robust despite the smaller number of funds (72).

Consistent with Ambachtsheer (2017a), we categorize a fund as *large* if it manages more than $10 billion of assets as of 2018 and *small* otherwise. In some tests, we

EXHIBIT 1

Sample Summary Statistics

	Number of Funds		Average AUM	% of Public Funds	% of Pension Plans	% of Indexed Liabilities	Duration of Liabilities
	2014–2018	2004–2018					
By Size							
Large							
Canada	17	11	77,194	71	53	85	14.78
RoW	109	52	69,972	48	81	58	13.77
Small							
Canada	36	24	3,158	33	83	65	14.06
RoW	88	18	3,865	48	97	65	14.64
By Geography							
US	127	60	28,488	39	94	55	13.66
Canada	53	35	26,905	45	74	67	14.22
UK	37	0	11,803	92	100	85	15.62
Netherlands	18	5	68,861	0	94	48	14.87
Other	15	5	169,373	67	0		

NOTES: This exhibit reports summary statistics of our sample by size and geography. RoW corresponds to rest-of-world. We report the number of funds present in the 5-year sample (2014–2018) and 15-year sample (2004–2018); the average level of AUM (in millions of 2018 US dollars); the proportions of public funds and pension plans in the five-year sample; and, for the pension plans, the average proportion of liabilities that are indexed to inflation and the average duration of liabilities (in years).

also focus on the subset of *very large* funds that manage more than $50 billion as of 2018. The category rest-of-world (RoW) captures all non-Canadian funds.

Exhibit 1 provides summary statistics of the sample. The five-year sample includes 17 Canadian and 109 RoW large funds. Large Canadian funds are slightly larger than their international peers ($77 billion versus $70 billion on average in 2018). They are mostly public funds and less likely to manage the pension plan liabilities. These statistics reflect the fact that several large Canadian funds are asset managers for a large number of smaller pension plans. Among funds that also manage liabilities, large Canadian funds have a high duration and a high proportion of pension liabilities that are indexed to inflation.

The sample also includes 36 Canadian and 88 RoW small funds. Small Canadian funds tend to be slightly smaller than their peers ($3.1 billion versus $3.8 billion) and are less likely to be public and manage the pension plan liabilities.

Overall, Canadian funds make up 20% of the five-year sample. Most funds are from the United States (127), followed by Canada (53), the United Kingdom (37), the Netherlands (18), and other countries (15).

Construction of Portfolios

We construct several portfolios for each fund. The *asset portfolio* includes all assets under management. We split the asset portfolio into several style categories: internal versus external management and active versus passive management. We also split the asset portfolio into six asset classes: public equity (stocks), fixed income (which includes government and corporate bonds and asset-backed securities), real assets (which include real estate, infrastructure, and land), hedge funds, private equity, and private credit (which includes private corporate loans and nonsecuritized mortgages).

If the fund is responsible for paying out pensions, we construct a *liability portfolio*. This portfolio is modeled as a mix of long-term real and nominal local bonds that matches (1) the duration of liabilities and (2) the proportion of liabilities indexed to

inflation that are reported by the fund.[1] We do not have information on the funded status of the fund, so we estimate the value of the liability portfolio assuming that the fund is 100% funded.

Finally, we construct an *asset–liability portfolio* that is long the fund's assets and short its liabilities. This portfolio is highly informative because its expected return indicates how well the asset portfolio is able to meet the fund's liability needs.

Construction of Performance Metrics

We introduce six metrics to analyze the performance of Canadian funds.

Asset-based metrics. Our first set of metrics is based on the fund's asset performance. We focus primarily on metrics that capture total fund performance rather than just the active management component. To this end, we use the Sharpe ratio of the asset portfolio as a measure of the risk–return trade-off of the fund's investments. The Sharpe ratio is calculated as the ratio of the arithmetic average of the fund's realized return in excess of the one-year interest rate on local nominal government bonds to the portfolio's volatility.

The volatility corresponds to the long-term volatility of the fund's asset portfolio return, as opposed to the portfolio's realized volatility. We calculate it in three steps. First, we split each asset class into several categories and estimate a detailed covariance matrix of these categories based on long time series of monthly benchmark returns. We use approximately 30–40 benchmarks depending on the fund. Second, we calculate for each year the annualized variance of the fund's asset portfolio as a function of (1) the policy weights reported by the fund in that year and (2) the long-term covariance matrix of benchmark returns. Third, we compute the average variance across all years in the sample and then obtain the portfolio volatility as the square root of this variance.[2] The details of the calculation are given in the online appendix.

The use of benchmark returns to calculate the covariance matrix has multiple advantages. It allows us to work with longer time series and higher-frequency data than the reported annual fund returns. Data on benchmark returns are usually available at the monthly level since 2000. The use of benchmark returns also addresses concerns about return smoothing in private markets because the benchmarks are public market proxies. Finally, it allows us to incorporate derivatives positions such as currency swaps. For example, if a Canadian fund hedges 50% of its US large-cap equity portfolio, then we treat the overall position as the sum of two investments: 50% in a hedged US large-cap equity benchmark and 50% in the equivalent nonhedged benchmark.

Building on work by Levy (2017), we also consider the fund's geometric average return as another performance metric. The rationale is that a fund that is able to cumulate high returns over long periods of time requires smaller contributions from the plan sponsor than if it were to fully invest in long-term government bonds. In other words, the fund's plan is less expensive to run.

Finally, we quantify the value added from the fund's active management by comparing the performance of the asset portfolio to that of the fund's policy portfolio, which is provided by the fund. In the online appendix, we empirically verify that the policy portfolio generally provides an accurate representation of the risk profile of the fund's asset portfolio.

Liability-based metrics. Our second set of performance metrics is based on the fund's ability to pay its liabilities. Building on the asset–liability literature (Sharpe and Tint 1990; Ezra 1991; Ang, Chen, and Sundaresan 2013), we calculate the Sharpe

[1] In a minority of cases, the information about indexation and duration of liabilities is approximated by CEM Benchmarking based on the country's demographics.

[2] This calculation makes use of the additive properties of the variance statistic, which are commonly used in models of time-varying volatility (Engle, Focardi, and Fabozzi 2013).

ratio of the asset–liability portfolio to evaluate its risk–return trade-off. The numerator of this ratio is the difference between the arithmetic average return of the asset portfolio and the arithmetic average return of the liability portfolio. The denominator is the volatility of this long–short portfolio that we calculate using the three-step procedure described previously. In this volatility calculation, we include duration overlays as positions in long-term bonds that perfectly correlate with the nominal component of the fund's liability portfolio.

We also consider the geometric average return of the fund's asset–liability portfolio. We compute it as the difference between the geometric average return of the fund's asset portfolio and the geometric average return of the fund's liability portfolio. A high net return indicates that the firm's assets are growing fast relative to its liabilities. Finally, we include the correlation between the asset and liability portfolios as a measure of liability hedging. A high correlation implies that the fund is hedging a high proportion of its liability risks. The correlation can be directly inferred from the volatility calculation of the fund's asset–liability portfolio. All performance metrics are based on returns that are expressed in the fund's local currency. Average returns and Sharpe ratios are net of costs.

PERFORMANCE COMPARISON

We now compare the return and hedging performance of large and small Canadian and RoW funds. Exhibit 2 displays the six performance metrics over the 5-year and 15-year samples.

Large Canadian funds outperformed their peers on all fronts over both time periods. In terms of asset performance, Canadian funds generated a higher Sharpe ratio (93% versus 60% for the peers over the five-year sample). The difference is significant at the 1% level based on a two-tailed unpaired *t*-test. Canadian funds also generated a higher average return (7.9% versus 6.1%) and a higher value added (0.6% per year versus 0.2%). The same findings hold for small Canadian funds.

The liability-based performance metrics indicate that Canadian funds also did a better job hedging their liability risks than did their peers. For example, large funds generated an average Sharpe ratio of 55% on their asset–liability portfolio over the five-year sample, which exceeds the 25% threshold that Ambachtsheer (2017b) considered as the aspirational "goldilocks" standard for pension funds. In contrast, RoW funds generated an average asset–liability Sharpe ratio of 12% during this period. Additionally, the average return on the Canadian funds' asset–liability portfolio was four times higher than that of their peers (4% versus 1.2%), and the correlation between their asset and liability portfolio was almost twice as high (48% versus 28%).

In the online appendix, we report the performance metrics for funds of other pension models for which we have a sufficient number of observations. There are five distinct pension models that we broadly characterize as the Canadian model, the Dutch model, the UK public funds model, the US corporate funds model, and the US public funds model. We find clear differences across models. In stark contrast with Canadian funds, US public and corporate funds scored low on all performance metrics. Their assets returned 2.5% less per year than Canadian funds (5.5%–5.6% versus 8%), had half their Sharpe ratio (48%–55% versus 93%), and were just high enough to cover the return of their liabilities. Additionally, US public funds had the lowest asset–liability correlation (15%). Equally interesting is the contrast in performance between UK public funds and Dutch funds, which highlights a common trade-off between asset performance and liability hedging. UK public funds generated higher returns than Dutch funds (8.5% versus 6.3%) but exhibited a lower correlation between assets and liabilities (29% versus 38%). In light of this trade-off, the fact that Canadian funds perform

EXHIBIT 2

Performance Metrics by Fund Size and Sample Period

		Assets			Asset—Liabilities		
		Sharpe Ratio	Average Return	Value Added	Sharpe Ratio	Average Return	Correl.
Large Funds	**2014–2018**						
	Canada	0.93	0.079	0.006	0.55	0.040	0.48
	RoW	0.59	0.061	0.002	0.12	0.012	0.28
	2004–2018						
	Canada	0.75	0.075	0.005	0.34	0.03	0.46
	RoW	0.62	0.070	0.002	0.17	0.02	0.20
Small Funds	**2014–2018**						
	Canada	0.82	0.073	0.005	0.38	0.028	0.46
	RoW	0.62	0.059	0.000	0.10	0.011	0.34
	2004–2018						
	Canada	0.75	0.075	0.006	0.28	0.022	0.42
	RoW	0.61	0.067	0.002	0.17	0.019	0.30

NOTES: This exhibit reports average performance metrics for large and small Canadian and RoW funds over two time periods: 2014–2018 and 2004–2018. For each sample, we consider metrics based on (1) the portfolio of the fund's assets and (2) the portfolio that is long the fund's assets and short its liabilities. Asset-portfolio metrics include the annualized Sharpe ratio, the annualized geometric average return, and the annualized performance relative to the fund's policy portfolio (value added). Asset-liability portfolio metrics include the Sharpe ratio, the geometric average return, and the correlation between the fund's assets and liabilities. Details on the calculation of these statistics are provided in the main text.

strongly on both fronts is noteworthy and suggests that it is possible to meet both objectives.

FEATURES OF THE CANADIAN MODEL

To gain insights into the drivers of the performance of Canadian funds, we examine their asset allocation and cost structure.

A Three-Pillar Business Model

We break down the funds' asset allocation and cost structure as follows. For each fund in the sample, we split the asset portfolio in three separate ways: (1) by internally and externally managed accounts, (2) by passively and actively managed accounts, and (3) by asset class. Then, for each subportfolio, we calculate its average percentage allocation in the fund between 2014 and 2018 and its average cost in basis points as a proportion of its size.[3] Finally, we calculate the average allocation and cost of each portfolio across Canadian and RoW funds and report the results in Exhibit 3. For confidentiality reasons, the cost for each asset class has been standardized so that it is always equal to 100 bps for the average Canadian fund.

[3] Costs include all external investment costs (i.e., manager base fees, performance fees, and carried interest) and internal costs associated with the investment process (i.e., costs of staff making investment decisions, support costs, consulting costs, CIO, CEO, board costs, IT, human resources, investment operations, and building and custody costs). Costs exclude brokerage commissions, implicit costs such as bid–ask spreads, costs of leverage, and property management costs associated with real estate, although returns are net of all said expenses. Costs and returns both exclude costs associated with benefit administration. When calculating the cost of an asset class, we exclude cases in which the fund does not invest in the asset class for the full five years. This is to avoid fixed costs involved with launching a new asset class. We also exclude the top 1% and bottom 1% cost values to reduce the impact of outliers.

EXHIBIT 3

Fund Asset Allocation and Cost Structure

		Internal/ External		Passive/ Active		Asset Class					
	Total	Int.	Ext.	Pas.	Act.	Stocks	F.I.	Real	Hedge	P.E.	Credit
Large Funds											
Allocation (%)											
Canada	100	52	48	19	81	37	29	18	7	7	1
RoW	100	23	77	21	79	42	34	9	6	6	1
Cost (bps)											
Canada	57	18	121	6	75	100	100	100	100	100	100
RoW	62	7	86	6	81	83	161	155	100	133	253
Small Funds											
Allocation (%)											
Canada	100	13	87	18	82	48	36	10	1	2	2
RoW	100	3	97	28	72	45	37	7	6	4	1
Cost (bps)											
Canada	49	10	57	11	61	100	100	100	100	100	100
RoW	56	10	61	9	77	86	135	107	92	102	101

NOTES: This exhibit reports the breakdown of portfolio allocation and costs for large and small Canadian and RoW funds between 2014 and 2018. We split portfolios three ways: by internal versus external management, by passive versus active management, and by asset class (stocks, fixed income [F.I.], real assets [Real], hedge funds [Hedge], private equity [P.E.], and private credit [Credit]). For each portfolio, we report the average percentage allocation inside the fund's asset portfolio and the average cost of the portfolio in basis points as a proportion of the portfolio's size. The cost for each asset class has been standardized so that it is always equal to 100 bps for Canadian funds. Details on the calculation of these statistics are provided in the main text.

We first examine the model of large Canadian funds. Exhibit 3 highlights three distinct features that characterize their allocation and cost structure. The first feature is the internal management of assets. Large Canadian funds manage 52% of their assets in-house on average, whereas RoW funds only manage 23% of their assets internally. In the online appendix, we find that this gap is even more pronounced for the very large funds that manage more than $50 billion: Canadian funds manage 81% of their assets in-house on average whereas non-Canadian only manage 39% internally. These results are in line with the findings of Ambachtsheer (2017a) and the World Bank (2017).

The immediate consequence of this shift to in-house management is a substantial reduction in costs. We observe in Exhibit 3 that Canadian funds indeed spend far less than RoW funds in most asset classes. For example, RoW funds spend 61% more than Canadian funds in fixed income, 55% more in real assets, 33% more in private equity, and 153% more in private credit. We will return to this point in the next section and provide a more precise estimate of the cost reduction associated with in-house management.

The second distinct feature of large Canadian funds is the redeployment of resources to internal investment teams for each asset class. We see in Exhibit 3 that Canadian funds spend 18 bps inside their internally managed portfolio on average, whereas RoW funds only spend 7 bps. Moreover, even though Canadian funds invest less externally, when they do outsource investments to external managers, they tend to outspend their peers. We find that Canadian funds spend 121 bps inside their externally managed portfolio, whereas RoW funds only spend 86 bps.

These cost differentials could potentially be driven by differences in the proportions of assets that are actively and passively managed—active management tends

to be relatively more expensive. However, Exhibit 3 shows that this is not the case. Canadian and RoW funds actively manage a fairly similar proportion of assets. In the online appendix, we go one step further and break down every asset class into subportfolios that are managed internally, externally, passively, and actively. We find that Canadian funds spend more than their peers in 13 of the 14 subportfolios, regardless of whether they are actively or passively managed. The evidence thus indicates that Canadian funds use the cost savings associated with in-house management to spend more in each asset class and style.

Some of the very large funds in the CEM dataset report additional details about their expenses. This information allows us to gain clearer insights into how Canadian funds allocate their additional expenses.[4] We find that, among the 6 large Canadian funds and 17 RoW funds that report risk management expenses, Canadian funds spend 0.73 bps of their total portfolio on risk management teams, whereas RoW funds only spend 0.18 bps. Similarly, Canadian funds spend 3.39 bps of their portfolio on IT infrastructure, whereas RoW funds spend 0.60 bps. These cost differentials are large, even once we account for the fact that Canadian funds manage a greater proportion of their assets in-house.

The third distinct feature of large Canadian funds is the allocation of capital toward assets that increase portfolio efficiency and hedge against liability risks. In particular, we see in Exhibit 3 that Canadian funds allocate on average twice as much of their AUM to real assets (18%) as RoW funds (9%). Real assets are generally more expensive to manage than stocks and bonds but help funds to hedge against both their interest rate and inflation risks (Amenc, Martellini, and Ziemann 2009).

In the online appendix, we verify the effectiveness of real assets in hedging liability risks by calculating the correlations between the assets and nominal and indexed liabilities for a representative fund of each pension model.[5] Real assets generally correlate strongly with the funds' indexed liabilities. In Canada, the correlation between the funds' real assets and indexed liabilities is as high as 44%. The savings gained from the transition to in-house management make it possible for Canadian funds to invest heavily in these assets.

One might argue that the high allocation to real assets is simply part of a broader shift by Canadian funds toward private asset classes that yield a liquidity risk premium. This explanation, however, would imply that Canadian funds generally invest more in all types of private assets. We find that this is not the case. In Exhibit 3, we see that Canadian funds' average allocation to private equity (7%) and private credit (1%) is roughly similar to that of RoW funds.[6] We elaborate on this point later and show that, among private assets, real assets have the most appealing risk–return profile for asset–liability management.

This asset–liability approach implemented by Canadian funds is not limited to real assets. In the online appendix, we show that their other assets generally correlate more with their liabilities than is the case for non-Canadian funds. For example, the correlation between the representative Canadian fund's stock portfolio and its indexed liabilities is 23%, whereas the equivalent correlation for the representative US fund is 13%. This result partly stems from the fact that Canadian funds invest more in

[4] Since 2010, CEM Benchmarking has compiled an additional survey on detailed expenses for the largest funds (see Macintosh and Scheibelhut 2012). We report average costs for 2017–2018.

[5] As mentioned in the previous section, correlations are based on the covariance matrix of benchmark returns. We aggregate these correlations at the asset class level using the average portfolio weights of the representative fund of each model between 2014 and 2018.

[6] In the online appendix, we run the same analysis for the subset of very large funds and find that, although Canadian funds invest more in private equity and private credit than do RoW funds, the difference is largest for real assets.

Canadian stocks, which are tilted toward commodity producer stocks that provide a natural hedge against inflation.

Altogether, we see in Exhibit 3 that Canadian funds spend on average 57 bps of their AUM to run their investments. In comparison, RoW funds spend 62 bps. Our analysis therefore shows that the three-pillar business model implemented by large Canadian funds allows these funds to spend more in each asset class and allocate more capital to strategic assets while spending less than their peers.

In the online appendix, we further analyze the allocation and cost structure of the five pension fund models introduced earlier. Canadian funds rank first in terms of assets managed in-house, allocation to real assets, and amounts spent inside the internally and externally managed portfolios. We also find that the RoW group is quite heterogeneous. Dutch funds invest mostly in bonds (47%), invest little in hedge funds (1%) and private equity (2%), and operate a low-cost model (35 bps) in large part due to the predominance of fixed income and a relatively low proportion of active management (72%). In comparison, UK public funds invest heavily in stocks (52%), hedge funds (11%), and private equity (7%). These differences help to explain why Dutch funds generate lower returns than UK public funds but better hedge their liability risks. Finally, US funds outsource the bulk of their investments to external managers (82% for public funds and 93% for corporate funds) and therefore run high-cost budgets (65–73 bps).[7]

The bottom part of Exhibit 3 reveals similar patterns in the business model of small Canadian funds. They tend to invest more in-house than their international peers (13% versus 3% for small RoW funds), more in real assets (10% versus 7%), and more actively (82% versus 72%). The evidence therefore suggests that the Canadian model has trickled down to smaller funds.

However, the evidence suggests that there is no uniform Canadian model. Small Canadian funds invest less in-house and less in real assets and private equity than do large funds. These results partly come from the fact that small funds do not have the scale needed to manage 50% of every asset class in-house. There are a number of fixed costs associated with setting up a professional investment team in-house. Without the scale required to manage a large proportion of their assets in-house, small Canadian funds have limited ability to invest in the private asset classes that are more expensive to manage.

Because of this constraint, small Canadian funds end up making adjustments elsewhere. They invest more in equities, which are cheaper to manage than private assets. They also invest less in hedge funds, which have been shown to generally underperform on a risk-adjusted basis (Heale, Beath, and Heuberger 2018). Small Canadian funds also cut down total costs to 49 bps of their asset portfolio. Finally, in the online appendix, we report the proportion of assets managed actively and internally within each asset class. Small Canadian funds invest more actively in equities and concentrate the bulk of in-house management in one asset class: fixed income.

Detailed Cost Analysis of Canadian Model Features

Because the three-pillar business model described earlier involves a complete restructuring of the cost structure, it is helpful to tease apart how each feature affects costs. To do so, we conduct a series of counterfactual exercises in which we progressively insert features of the Canadian model into the business model of a RoW fund and quantify the impact of each feature on the fund's total costs.

[7] Andonov, Bauer, and Cremers (2017) argued that the public funds' large investments in stocks and private equity are driven by incentives to report a better funding status according to US regulation (see also Novy-Marx 2013).

EXHIBIT 4

Cost Structure of Large Canadian Funds as a Proportion of Total AUM

		Asset Class					
	Total	Stocks	F. I.	Real	Hedge	P. E.	Credit
Canadian Model							
(1) RoW	60.6	10.5	5.7	10.7	13.3	20.0	0.5
(2) + in-house management	41.6	10.1	3.0	5.5	9.9	12.8	0.3
(3) + increased investment in each asset class	45.9	14.0	2.9	6.3	9.9	12.3	0.5
(4) + revised asset allocation	51.3	12.4	2.5	12.6	10.5	12.8	0.6
Canadian Model (excl. in-house mgmt)							
(5) Total	72.6	12.7	4.5	22.0	14.3	18.1	0.9

NOTES: This exhibit reports the results of a series of counterfactual exercises in which we add features of the Canadian model to a representative large-size RoW fund and estimate their impact on costs for each asset class. Costs are expressed in basis points as a proportion of the fund's size. The total cost column is the sum of the cost of each asset class and excludes any overheads. Row 1 reports the average cost by asset class for the representative RoW fund between 2014 and 2018. Row 2 assumes that the RoW fund adopts the same proportion of in-house management within each asset class as the average large-size Canadian fund. Row 3 assumes that the RoW fund also adopts the Canadian fund's proportion of active management within each asset class and spending level on each portfolio. Row 4 assumes that the RoW fund also adopts the Canadian fund's asset allocation. Row 5 assumes that the RoW fund adopts all features of the Canadian fund except for the proportion of in-house management. Details of the calculations are provided in the main text.

This counterfactual analysis is based on the representative large-size Canadian and RoW funds described in Exhibit 3. For each fund, we further break down every asset class into an internally managed passive portfolio, an externally managed passive portfolio, an internally managed active portfolio, and an externally managed active portfolio. For each subportfolio, we calculate the average allocation and cost exactly as in Exhibit 3. We then progressively modify the RoW fund by adopting the Canadian fund's (1) proportion of assets managed internally inside each asset class, (2) cost of each sub-portfolio and proportion of assets managed actively inside each asset class, and (3) overall asset allocation.

Exhibit 4 reports the funds' average cost allocation aggregated at the asset class level. Unlike Exhibit 3, the costs here are expressed as proportions of the funds' asset portfolio. The first row summarizes the cost structure for the representative RoW fund. The fund spends approximately 10 bps on public equities, 6 bps on fixed-income markets, 10 bps on real asset markets, 13 bps on hedge fund strategies, 20 bps on private equity, and 1 bp on private credit. In total, these costs (which exclude overheads and total portfolio expenses) amount to 60.6 bps of the asset portfolio.[8]

In the second row, we assume that the RoW fund adopts the same proportion of internal management inside each asset class as the average Canadian fund. Costs drop down to 41.6 bps as the result of this transition. This means that, on average, increasing the proportion of assets managed in-house by one percentage point decreases total costs by 0.7 bps. In the case of large Canadian funds, going from 23% to 52% managed in-house reduces total costs by approximately one-third. The cost reduction is as high as 50% for fixed income and real assets, where the shift to in-house management is the largest.

In the third row, we assume that the RoW fund adopts the same weight of passive and active management as well as the same cost budget as the average Canadian

[8]The costs of individual categories do not exactly add up to the total costs reported in Exhibit 4 for several reasons. There are additional expenses for derivative overlays. We also do not include the cost of an asset class if the fund did not invest in that asset class for the full five years. Finally, funds incur additional overhead expenses at the total portfolio level.

fund within each subportfolio. Because of the increased spending on each subportfolio, costs go up by 4.3 bps to a total of 45.9 bps.

In the fourth row, we assume that the RoW fund also adopts the Canadian-model asset allocation that features more real assets. Costs go up again by approximately 5.4 bps to 51.3 bps but remain 9.3 bps lower than the cost of the average RoW fund.[9] Interestingly, we note that the total cost on real assets is not significantly higher than in the first row. Therefore, by managing their assets internally, Canadian funds are able to double their exposure to real assets for approximately the same cost.

In the fifth row, we repeat the analysis but exclude the transition to in-house management. In other words, we examine what the cost structure of the RoW fund would look like if it spent more on each asset class and adopted the Canadian-model asset allocation but did not change the proportion of assets managed in-house. We see that the total cost would jump from 60.6 bps (first row) to 72.6 bps. This jump is greater than it would be if the fund were to manage its assets internally (41.6 bps to 51.3 bps). This means that the transition to in-house management not only reduces costs but also makes it less costly for the fund to restructure the remaining parts of its portfolio.

DRIVERS OF PERFORMANCE

We next evaluate the extent to which the business model implemented by Canadian funds has been responsible for their past performance.

Regression Analysis of Performance

In the online appendix, we run a comprehensive statistical analysis in which we regress several performance metrics on a set of predictor variables. The performance metrics consist of the Sharpe ratios of the funds' asset and asset–liability portfolios and the correlation between the funds' assets and liabilities. We calculate these metrics over the 5-year and 15-year samples. The asset-based predictor variables include the funds' average proportion of assets actively managed during these sample periods, the average proportion of assets invested in each asset class, the average fund size, and the average proportion of assets managed internally conditional on being either a large or a small fund. We split the proportion of in-house management by size group because, as explained earlier, large funds have greater ability to scale up internal operations. The liability-based variables include the funds' average duration of liabilities and the average proportion of liabilities indexed to inflation.

Unsurprisingly, we find that all asset classes positively contribute to the funds' asset Sharpe ratio in the 15-year sample, which speaks to the importance of building a well-diversified portfolio for long-term asset performance. However, not all asset classes protect funds against liability risks. Funds that invest heavily in stocks and private equity tend to have strong return performance but low correlation between assets and liabilities. By contrast, investments in fixed income and real assets tend to increase both the asset–liability Sharpe ratio and the asset–liability correlation. These asset classes have cash flow characteristics that are similar to the pension funds' liability profiles and hence allow funds to capture a risk premium while aligning the risk of their liability portfolio to their assets.

The proportion of assets managed in-house also leads to a high Sharpe ratio. This effect is most significant among large funds that can afford to build large internal

[9] The total cost of 51.3 bps does not include the overheads and total portfolio expenses that are relatively large for Canadian funds (the total cost reported in Exhibit 4 is 57 bps).

management teams. The fact that this effect is significant after having controlled for the funds' asset allocation confirms our hypothesis that in-house management brings advantages beyond the ability to allocate capital to strategic assets. These findings are consistent with prior research on the impact of in-house management on value creation (Macintosh and Scheibelhut 2012; Beath 2015).

Finally, we find that a fund's liability structure plays an important role in driving its return and hedging performance. Specifically, funds with a higher proportion of liabilities indexed to inflation tend to have stronger asset–liability performance metrics. The explanation is that indexed liabilities tend to correlate more with growth assets than with nominal liabilities. Large Canadian funds benefit from this effect because they have a greater fraction of indexed liabilities than RoW funds.

Backtests for US Funds

One might be concerned that the regression analysis does not fully control for environmental factors. For example, it may be that the Canadian funds' high Sharpe ratio was partly driven by the booming real estate market in Canada between 2004 and 2018. To address this concern, we run a series of backtests in which we study how US pension funds would have performed by adopting features of the Canadian model. The US environment is fitting for this analysis because (1) it includes the majority of funds in the sample and (2) it gives us the opportunity to work with two distinct pension models: the US public and corporate pension fund models.

We set up the backtests as follows, using the same three performance metrics as in the previous section (i.e., the asset Sharpe ratio, asset–liability Sharpe ratio, and asset–liability correlation). We begin by constructing representative US public and corporate funds from the average allocation and liability profile of large public and corporate US funds that do both asset and liability management over the five-year sample. We abstract from the funds' active performance by calculating past returns from benchmarks for each asset category. We then progressively swap features of both funds for those of the average large-size Canadian fund and repeat the historical performance analysis. The Sharpe ratios initially abstract from fees, but we incorporate fees in the final backtest.

Exhibit 5 reports the results. The first row shows the performance metrics for the representative US public fund. In the second row, we swap the fund's liability structure for that of the representative Canadian fund to make sure we do not adapt Canadian-model recommendations to the wrong type of fund. The asset–liability Sharpe ratio and the correlation both increase because the Canadian fund has a higher proportion of indexed liabilities.

In the next set of rows, we swap the asset allocation to that of the average Canadian fund. We do so in several steps. First, we change the allocation at the asset class level (third row). The Canadian fund's allocation includes more fixed income and real assets and less private equity. We then tilt the allocation inside each asset class (fourth row). The Canadian fund's stock allocation includes more global and emerging-market equities, its bond allocation includes more long-term bonds, and its real asset allocation includes more infrastructure and fewer real estate investment trusts. We treat Canadian equity and fixed income as home assets and, hence, replace them with US equivalents.

The Canadian-model allocation leads to a significant increase across all performance metrics from the US public model allocation. Comparing rows 2 and 4, we find that the 15-year asset Sharpe ratio increases from 70% to 83%, the 15-year asset–liability Sharpe ratio increases from 16% to 19%, and the asset–liability correlation increases from 22% to 36%. We emphasize that this increase in performance results from an integrated reorganization of the portfolio that goes beyond a simple allocation

EXHIBIT 5

Backtests for Large US Pension Funds

	Asset Sharpe Ratio		Asset–Liab. Sharpe Ratio		Asset–Liab. Corr.
	5-yr	15-yr	5-yr	15-yr	
US Public Funds					
(1) US public fund	0.428	0.703	0.019	0.113	0.151
(2) + CA liability profile	0.428	0.703	0.061	0.160	0.218
(3) + CA asset allocation	0.468	0.778	0.046	0.148	0.308
(4) + CA tilts inside asset classes	0.561	0.830	0.122	0.191	0.358
(5) + additional tilts	0.524	0.838	0.125	0.254	0.420
(6) + CA cost structure	0.536	0.850	0.158	0.287	0.420
US Corporate Funds					
(1) US corporate fund	0.583	0.838	0.024	0.106	0.345
(2) + CA liability profile	0.583	0.838	0.155	0.220	0.450
(3) + CA asset allocation	0.555	0.816	0.133	0.205	0.409
(4) + CA tilts inside asset classes	0.561	0.830	0.122	0.191	0.358
(5) + additional tilts	0.524	0.838	0.125	0.254	0.420
(6) + CA cost structure	0.548	0.862	0.192	0.321	0.420

NOTES: This Exhibit reports the results of backtests in which we study how representative large-size US public and corporate funds would have performed between 2014 and 2018 and 2004 and 2018 by adopting features of the Canadian model. We consider three performance metrics: the Sharpe ratio of the funds' asset portfolio, the Sharpe ratio of the portfolio that is long the fund's assets and short their liabilities, and the correlation between the funds' assets and liabilities. Unlike Exhibit 2, performance metrics are based on benchmark returns within each asset class as opposed to fund realized returns. Rows 2–6 of each panel assume different specifications of the US fund. Row 2 swaps the US fund's liability profile for that of a representative large-size Canadian fund. Row 3 also swaps the overall asset allocation for the Canadian fund's. Rows 4 and 5 also swap the portfolio inside each asset class for the Canadian fund's. The difference between the rows is that, in Row 4, we substitute investments in Canadian stocks and fixed income for US equivalents, whereas in Row 5, we substitute Canadian stocks and fixed income for emerging equity and US long-term bonds, which highly correlate with US liability risks. Row 6 also swaps the overall cost structure for that of the Canadian fund's.

toward real assets. For example, we see in the third row that the change in overall asset allocation is not sufficient—it increases the asset–liability correlation from 22% to 31% but decreases the asset–liability Sharpe ratio. However, when the change in asset allocation is combined with the restructuration of the real asset portfolio, all performance metrics go up.

In the fifth row of Exhibit 5, we take the analysis one step further and swap the capital that Canadian funds invest in Canadian equities and fixed income for assets that have a strong correlation with US pension liability risks. Specifically, we allocate the Canadian equity share to emerging markets, which have a high correlation with US indexed liabilities, and the Canadian bond share to US long-term bonds, which have a high correlation with nominal liabilities. Once again, we observe a jump across all performance metrics. The 15-year asset Sharpe ratio increases from 83% to 84%, the 15-year asset–liability Sharpe ratio increases from 19% to 25%, and the asset–liability correlation increases from 36% to 42%.

In the sixth and final row of Exhibit 5, we incorporate the difference in costs on each asset class by US public and Canadian funds, which amounts to approximately 10 bps of the total portfolio. The resulting Sharpe ratios increase once again.

Altogether, the combined impact of these changes on fund performance is large. Comparing rows 1 and 6, the 15-year asset Sharpe ratio increases from 70% to 85%, the 15-year asset–liability Sharpe ratio increases from 11% to 28%, and the asset–liability correlation increases from 15% to 42%. These estimates do not include any

additional performance resulting from the Canadian funds' decision to spend more on each asset class.

Exhibit 5 also reports the results for the representative US corporate fund. Once again, we see that the adoption of the Canadian model yields higher performance metrics on all fronts: The 15-year asset Sharpe ratio increases from 84% to 86%, the 15-year asset–liability Sharpe ratio increases from 11% to 32%, and the asset–liability correlation increases from 35% to 42%.

However, the US public and corporate models differ in a couple of important ways. First, a corporate fund with Canadian liability profile and a US asset model already has a high asset–liability correlation (45%). The reason is that the US corporate model prioritizes fixed-income investments (40%) that provide a strong hedge against nominal liabilities. Consequently, the increase in correlation that results from the adoption of the Canadian model is smaller than for the US public fund.

Second, the main performance increase for the corporate model comes from the last row of Exhibit 5, in which we incorporate the effect of costs. Because US corporate funds outsource the majority of their investments to external managers, they are relatively expensive to run. We confirm this point in the online appendix: US corporate funds spend 12% more than Canadian funds inside their stock portfolio, 99% more in fixed income, 82% more in real assets, 3% more in hedge funds, 38% more in private equity, and 196% more in private credit. As a result of this cost differential, US corporate funds hedge the same amount of pension liability risks as Canadian funds but pay a higher cost to do so.

LESSONS LEARNED FROM THE CANADIAN MODEL

Our analysis of the Canadian pension fund model raises the question of whether this model can be replicated elsewhere. Doing so may not always be straightforward. As the prior literature has made clear, moving the investment team in-house requires independent corporate governance and competitive compensation schemes to attract and retain talent (Drucker 1976; Ambachtsheer 2007; World Bank 2017). It also requires a regulation system that provides flexibility regarding the ability to manage balance sheet shortfalls. For example, a system like the Dutch model that enforces strict solvency requirements on a marked-to-market basis makes it challenging to invest in growth assets that do not perfectly correlate with the funds' liabilities (Frijns 2010).

Nevertheless, our findings yield a number of insights for pension fund investors and regulators around the world. First, it is possible to adopt some features of the Canadian model even if the fund is not in a position to use the model entirely. We observe, for example, that small Canadian funds, which do not have the scale to fully replicate the model of large funds, have been able to increase the proportion of in-house management, allocation to real estate and infrastructure, and active management in select asset classes. These funds have performed well in terms of both portfolio efficiency and liability hedging over the past 15 years.

Second, the cost reduction resulting from in-house management opens the door to a broad range of additional resources. These resources include greater ability to (1) recruit sophisticated investment teams for each asset class and (2) shift the portfolio toward strategic assets. We find that the incremental cost of restructuring the portfolio also decreases once it is managed in-house. This may be because of synergies associated with internal management. For example, expenses spent on risk management and IT infrastructure benefit all investment teams at once.

Third, the gains associated with the three-pillar model implemented by Canadian funds go beyond value creation relative to a benchmark portfolio. The gains are a

more efficient asset–liability portfolio and greater protection against pension liability risks. In the backtests of the previous section, we abstract from value creation by estimating these fund performance metrics using benchmark returns and show that the gains associated with the adoption of the Canadian model are economically large.

Fourth, there seem to be important benefits associated with having pension liabilities indexed to inflation. Pension funds can partly hedge against these liability risks by investing in a diversified basket of growth assets that also increase return performance. This asset–liability strategy is appealing because it does not require the funds to invest exclusively in low-yield bonds and duration overlays to hedge against their liability risks. This strategy may therefore contribute to the long-term sustainability of pension plans.

CONCLUSION

This article presents a quantitative portrait of the Canadian pension fund model. We show that a central factor driving the strong performance of Canadian funds has been the implementation of a cost-efficient business model that consists of (1) managing assets in-house to reduce costs, (2) redeploying resources to internal investment teams for each asset class, and (3) allocating capital to strategic assets that increase portfolio efficiency and hedge against liability risks. We also show that a complementary factor driving the Canadian funds' success is the indexation of their pension plan liabilities. This feature has made it easier for large Canadian pension funds to hedge against their liability risks by owning a diversified mix of growth assets. Additionally, we find a large number of smaller Canadian funds have adopted a light version of this business model that is more feasible for their size.

The years ahead will put the Canadian model to the test, as the severe impact of COVID-19 on commercial real estate, equities, and corporate bonds will undeniably hurt the funds' assets in the short run. How resilient is the Canadian model to a global pandemic? Will two-pronged strategies that increase asset performance and hedge against liability risks change in the post-pandemic world? We leave these questions for future research.

ACKNOWLEDGMENTS

This research is funded by the National Pension Hub of the Global Risk Institute. We thank Keith Ambachtsheer, Yuedan Liu, Tom Scheibelhut, and participants of the NPH Insight Series 2020 for insightful comments.

REFERENCES

Ambachtsheer, K. P. *Pension Revolution: A Solution to the Pensions Crisis*. Hoboken, NJ: Wiley, 2007.

——. *The Future of Pension Fund Management*. Hoboken, NJ: Wiley, 2016.

——. "The 'Canada Model' for Pension Fund Management: Past, Present, and Future." *The Ambachtsheer Letter*, August 1, 2017a.

——. "Managing Balance Sheet Risk: Do Pension Funds Know What They Are Doing?" *The Ambachtsheer Letter*, June 2017b.

——. 2021. "The 'Canadian Pension Model': Past, Present, and Future." *The Journal of Portfolio Management*, forthcoming.

Amenc, N., L. Martellini, and V. Ziemann. 2009. "Inflation-Hedging Properties of Real Assets and Implications for Asset-Liability Management Decisions." *The Journal of Portfolio Management* 35 (4): 94–110.

Andonov, A., R. M. Bauer, and M. J. Cremers. 2017. "Pension Fund Asset Allocation and Liability Discount Rates." *Review of Financial Studies* 30 (8): 2555–2595.

Ang, A., B. Chen, and S. Sundaresan. 2013. "Liability-Driven Investment with Downside Risk." *The Journal of Portfolio Management* 40 (1): 71–87.

Barras, L., and S. Betermier. "Why Do Homeowners Invest the Bulk of Their Wealth in Their Home?" Working paper, 2020, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2776350.

Beath, A. "Value Added by Large Institutional Investors between 1992–2013." CEM Benchmarking, Inc., 2015, https://www.cembenchmarking.com/ri/insight/5.

Bédart-Pagé, G., A. Demers, E. Tuer, and M. Tremblay. "Large Canadian Public Pension Funds: A Financial System Perspective." *Financial System Review*, pp. 33–38, Bank of Canada, 2016.

Blake, D., B. Lehman, and A. Timmerman. 1999. "Asset Allocation Dynamics and Pension Fund Performance." *The Journal of Business* 72 (4): 429–461.

Coggin, T., F. Fabozzi, and S. Rahman. 1993. "The Investment Performance of U.S. Equity Pension Fund Managers: An Empirical Investigation." *Journal of Finance* 48: 1039–1055.

Drucker, P. F. *The Unseen Revolution: How Pension Fund Socialism Came to America*, 1st ed. New York: Harper & Row, 1976.

The Economist. 2012. "Canada's Pension Funds: Maple Revolutionaries." Finance & Economics Section, Mars 3 Ed.

Engle, R., S. Focardi, and F. Fabozzi. "ARC/GARCH Models in Applied Financial Econometrics." In *Encyclopedia of Financial Models*, edited by F. Fabozzi. Hoboken, NJ: Wiley, 2013.

Ezra, D. 1991. "Asset Allocation by Surplus Optimization." *Financial Analysts Journal* 47: 51–57.

Frijns, J. 2010. "Dutch Pension Funds: Aging Giants Suffering from Weak and Inconsistent Risk Management." *Rotman International Journal of Pension Management* 3 (2): 16–21.

Heale, M., A. Beath, and E. Heuberger. "Hedging Fund Reality Check," CEM Benchmarking, Inc., 2018, https://www.cembenchmarking.com/ri/insight/15.

Hoevenaars, R., R. Molenaar, P. Schotman, and R. Steenkamp. 2008. "Strategic Asset Allocation with Liabilities: Beyond Stocks and Bonds." *Journal of Economic Dynamics and Control* 32: 2939–2970.

Leibowitz, M., S. Koegelman, and L. Bader. 1992. "Asset Performance and Surplus Control: A Dual-Shortfall Approach." *The Journal of Portfolio Management* 18 (2): 28–37.

Levy, M. 2017. "Measuring Portfolio Performance: Sharpe, Alpha, or the Geometric Mean?" *The Journal of Investment Management* 15 (3): 1–17.

Lipshitz, C., and I. Walter. "Bridging the Gaps: Public Pension Funds and Infrastructure Finance," New York University, Stern School of Business, 2019.

——. 2020. "Public Pension Reform and the 49th Parallel: Lessons from Canada for the U.S." *Financial Markets, Institutions & Instruments* 29 (4): 121–162.

Macintosh, J., and T. Scheibelhut. 2012. "How Large Pension Funds Organize Themselves: Findings from a Unique 19-Fund Survey." *Rotman International Journal of Pension Management* 5 (1): 34–42.

Novy-Marx. 2013. "Logical Implications of the GASB's Methodology for Valuing Pension Liabilities." *Financial Analysts Journal* 69 (1): 26–32.

Sharpe, W., and L. Tint. 1990. "Liabilities—A New Approach." *The Journal of Portfolio Management* 16 (2): 5–10.

World Bank Group. "The Evolution of the Canadian Pension Model: Practical Lessons for Building World-Class Pension Organizations." International Bank of Reconstruction and Development/World Bank, Washington DC, 2017.