

March 12, 2014

Me Anne-Marie Beaudoin
Corporate Secretary
Autorité des marchés financiers
800, Square Victoria, 22e étage
C.P. 246, Tour de la Bourse
Montréal, Québec H4Z 1G3
Email: consultation-en-cours@lautorite.qc.ca

The Secretary
Ontario Securities Commission
20 Queen Street West
22nd Floor, Box 55
Toronto, Ontario M5H 3S8
Email: comments@osc.gov.on.ca

Re: Request for Comments, CSA Discussion Paper and Request for Comment 81-324 Proposed CSA Mutual Fund Risk Classification Methodology

As part of a global organization whose mission is to help investors make better decisions, Morningstar Canada welcomes the opportunity to comment on the implementation of a standard risk-classification methodology for use by fund managers in the Fund Facts document.

We offer the following for the purposes of presenting our concerns and suggestions; we would be delighted to expand further on any or all of the comments/responses contained herein.

Our comments fall under three headings:

1. Single risk measures are risky
2. Standard deviation (SD) is suboptimal
3. Consider holdings-based measures

Single Risk Measures are Risky

Morningstar recognizes that without an objective and transparent standard it is difficult for investors to make meaningful comparisons or for risk classifications to be verified independently.

There are, however, unintended dangers in enforcing a single, standardized risk measure to evaluate risk across funds.

One shortcoming of any standard risk indicator is that it can suggest (to an investor) that it is the *only* measure of risk that needs consideration. In fact, investment risk is multi-faceted. There are numerous risk factors that demand consideration, including systematic risk factors (equity, interest rate, etc.) as well as fund-specific risk factors. Some of these factors can be beneficially revealed through the use of holdings-based measures, but the fundamental point here is that no single measure could possibly describe all the meaningful risks that exposure to a particular fund or ETF entails.

Another issue – and this especially applies to risk measures such as SD – arises when an investor overlooks/avoids an appropriate fund selection for inclusion in their portfolio due to its independent risk rating being higher than the investor's deemed risk tolerance. The single-measure approach is most relevant when applied to a portfolio that is dominated by, or is entirely represented by, the fund being assessed. But, if an investor's portfolio is made up of a

basket of investments (including funds), there can be a risk-*reducing* effect of including a fund that, independently, is viewed as out-of-range from a risk perspective.

For example, a conservative investor's portfolio that is missing a key sector or asset class, essential for prudent diversification (and thus risk-reduction), may demand the inclusion of a small amount of a concentrated sector mutual fund or ETF. A single-measure risk score for such a vehicle may be higher than recommended for the investor and they are consequently dissuaded from incorporating it. The irony and potential downside is that the risk of the conservative portfolio may actually be higher than it otherwise would be had the investor included the diversifying investment. Diversification as a risk-reduction activity is a sensible approach, practiced by many, and supported by decades of investment research.

To illustrate how sector diversification can reduce risk, we formed a portfolio of five sector equity ETFs designed to mimic the monthly returns of the iShares S&P/TSX 60 Index ETF¹ over the 10-year period January 2004—December 2013². Exhibit 1 shows the ETFs and their weights in the portfolio, 10-year annualized SD, and the 10-year annualized SD of the portfolio.³ Note that the SD of the portfolio is lower than that of all of the constituent ETFs, even though the SDs of some of the ETFs are 60 to 130% higher than the SD of the portfolio. This illustrates the power of diversification as first explained by Nobel laureate Harry Markowitz in 1952 and has been one of the main principles of investment practice ever since.

Exhibit 1: Diversification Effect on Standard Deviation (SD)

ETF	Category	Portfolio Weight	10-Year Annualized SD (%)
iShares S&P/TSX Capped Financials Index	Financial Services Equity	37%	15.4
iShares S&P/TSX Capped Info Tech Index	Miscellaneous - Sector Equity	10%	25.9
iShares S&P/TSX Capped Energy Index	Natural Resources Equity	32%	22.7
iShares S&P/TSX Global Gold Index	Precious Metals Equity	10%	33.2
iShares S&P/TSX Capped REIT Index	Real Estate Equity	11%	15.5
Portfolio of ETFs		100%	14.5

SD is Suboptimal

While adopting a standardized measure for fund/ETF risk presents its own inherent challenges, the choice of which measure(s) to employ also present challenges and demands careful consideration.

¹ An ETF that many investors use to represent the Canadian equity market. Its ticker is XIU.

² The correlation between monthly returns of iShares S&P/TSX 60 and those of the portfolio over this period is 97%.

³ We calculated standard deviation using the formula provided in "Annex A.5" of *CSA Discussion Paper and Request for Comment 81-324 - Proposed CSA Mutual Fund Risk Classification Methodology*.

Although SD is a widely used measure of investment risk, it suffers from some serious flaws that can cloud its description of investment risk. We believe there are better measures that can be just as easily implemented that will produce a better picture of risk for investors. The ubiquity of the SD measure should not be a reason to prefer it over other, more informative measures.

We will address the case for the inclusion of a holdings-based measure in the next section, but for the purposes of using historical performance as a proxy for measuring investment risk (the role SD plays), Morningstar strongly encourages the CSA to consider Conditional Value at Risk (CVaR) as a substitute.

SD measures the dispersion of returns around the average return of an investment. It does not differentiate between above-average and below-average returns. It also assumes that the nature of investment returns is normally distributed. This is often not the case, and as anyone who participated in the markets in 2008-2009 will attest, the need to describe so-called “non-normal” events is highly desirable.

Morningstar believes that CVaR is a better measure than SD. The reason: instead of measuring the dispersion of returns around the average, CVaR focuses entirely on possible losses. For most investors, risk includes both the *likelihood* and *magnitude* of potential losses. Since CVaR provides a summary of how severe losses can be, it is a much more useful and meaningful risk measure than the more commonplace SD.

CVaR is enjoying a widespread and growing popularity in the institutional investment world and will undoubtedly become more popular in the retail investment space. This is an opportunity to embrace a better, if imperfect proxy for investment risk. But because it is not as well known (in the mutual fund world), some misconceptions surrounding CVaR have arisen, which we briefly address here. As with all aspects of this response document, we welcome all inquiries to elaborate.

Misconception #1: CVaR is difficult to calculate

The notion that CVaR is more difficult to calculate than SD measure is at best curious. They use the same input data and both are invariably calculated by a computer program. CVaR is calculated in spreadsheets as well as proprietary performance measurement systems. Users of SD measures, when asked, are often at a loss as to *how* to calculate the measure, but this does not preclude them from effectively utilizing the measure. The same can be said for CVaR.

Misconception #2: CVaR and SD are proxies for each other

If this were true, there would be no need to calculate CVaRs. The reality is that CVaR can reveal what SD often misses: the magnitude and likelihood of an extreme negative return. Both measures are dependent on the notion that the historical return pattern is relevant to assessing risk, and to the extent that this is desirable, CVaR does a better job.

Appendix A contains a real-life example of two mutual funds⁴ that exhibit very similar SD measures, but have very different CVaRs. This is simply more informative to the risk-averse investor and can be very helpful in selecting investments.

Given that SD is suboptimal for historical risk measurement (especially in light of better measures), and given that current technology makes CVaR as straightforward to calculate as SD, and that CVaR and SD are not proxies for each other, Morningstar strongly recommends using CVaR for the historical risk measurement component of a risk-classification system.

⁴ While the funds and data are real, we have deliberately concealed their names to avoid unnecessary diversions.

Consider Holdings-based Measures

The greatest limitation of historical returns-based risk measures such as SD and CVaR is that they are entirely backwards-facing. To use these measures, one must agree that their historical performance pattern is likely to be imitated going forward – and this is clearly not guaranteed. Another problem is that these same returns are generated by the underlying securities (“holdings”) in which they are invested; should the nature of those holdings change, one can reasonably expect the performance to change as well. For example, a mutual fund that takes on a new, stylistically different manager is likely to see a change in its holdings characteristics and thereby potentially deviate from its historical record, making a historical returns-based risk measure far less relevant. A risk measure that includes an assessment of the risks associated with current holdings is clearly advantageous.

The Canadian Investment Funds Standards Committee (CIFSC)⁵ classifies all Canadian funds into various categories *based on each fund’s portfolio holdings*. Fund categories generally reflect exposures to risk factors specific to their group. For example, one of the fixed income categories is High Yield Fixed Income. These funds are disproportionately exposed to the risk of widening credit quality spreads. Another category is Global Fixed Income: these funds are exposed to the risk of the Canadian dollar strengthening relative to other currencies.

There are a multitude of holdings-based, quantitative risk measures that can be calculated (as many are currently for the CIFSC) that would assist investors in assessing the overall risk of a fund or ETF. These measures could be combined with the returns-based measure (we recommend CVaR over SD) to produce a single-measure risk metric. Morningstar is willing and able to expand on these possibilities should it interest the CSA.

⁵ In Canada, the broad risk assessment for mutual funds is achieved first through CIFSC’s standardization of the categories in which they are grouped. Canada is the global leader in this respect – in no other jurisdiction do all fund companies, media outlets and third-party data/analyses providers adhere to a single category scheme.

Issues for Comment on the Notice and Request for Comment

1. As a threshold question, should the CSA proceed with (i) mandating the Proposed *Methodology* or (ii) adopting the Proposed Methodology only as guidance for fund managers to identify the mutual fund's risk level on the prescribed scale in the Fund Facts? Are there other means of achieving the same objective than by mandating the Proposed Methodology, or by adopting it only as guidance? We request feedback from investment fund managers and dealers on what a reasonable transition period would be for this.

MORNINGSTAR COMMENT: Morningstar believes there are risks associated with utilizing a single measure to evaluate investment risks of a fund or ETF (see opening comments). However, should the CSA proceed with mandating a methodology for a standard risk assessment, we strongly recommend that it be based on a blend of measures that includes Conditional Value at Risk (CVaR) and a holdings-based approach. We believe that the use of the SD measure as the sole measure of risk does not serve the best interests of the investor.

2. We seek feedback on whether the Proposed Methodology could be used in similar documents to Fund Facts for other types of publicly-offered investment funds, particularly ETFs. For ETFs, what, if any, adjustments would we need to make to the Proposed Methodology? For instance should SD be calculated with returns based on market price or net asset value per unit?

MORNINGSTAR COMMENT: A number of risk measures work for other types of publicly offered investment funds, such as ETFs. Risk measures should be based on market price, rather than net asset value, as there is the potential for a mismatch between the two and market price best depicts an investor's experience in an ETF.

3. We seek feedback on whether you agree or disagree with our perspective of the benefits of having a standard methodology, as well as whether you agree or disagree with our perspective on the cost of implementing the Proposed Methodology.

MORNINGSTAR COMMENT: Morningstar acknowledges that a standard methodology benefits investors by enabling meaningful fund comparisons and allowing for risk classifications to be verified independently. However, we must reiterate our concerns about a sole, standard risk indicator in that it can suggest to an investor that it is the *only* measure of risk that need be considered.

4. We do not currently propose to allow fund managers discretion to override the quantitative calculation for risk classification purposes. Do you agree with this approach? Should we allow discretion for fund managers to move their risk classification higher only?

MORNINGSTAR COMMENT: Morningstar agrees with the CSA's decision to give managers no discretion to override the quantitative calculation, should it be mandated. The CSA's quantitative approach imposes a uniform methodology across the fund universe. Doing otherwise would undermine investors' ability to make fair comparisons. It would also undercut the CSA's desire to standardize its risk-classification scheme.

5. Keeping the criteria outlined in the introduction above in mind, would you recommend other risk indicators? If yes, please explain and supplement your recommendations with data/analysis wherever possible.

MORNINGSTAR COMMENT: We would urge the CSA to consider approaches that blend a returns-based measure such as CVaR with a holdings-based measure.

A risk measure is helpful to investors only if it can be used to distinguish investments that are inherently more risky from those that are inherently less risky. But risk measures that rely solely on historical returns are useful only if the risk manifested in past returns is indicative of risk that will be manifested in the future. Since this is not

necessarily the case, we recommend that fund risk also be measured using methods based on portfolio holdings, thus reflecting the inherent risks. A holdings-based measure can be combined with a returns-based measure to form an overall risk measure that takes into account both the riskiness that manifested in the past with risks currently inherent in the portfolio.

Holdings-based analysis is not new. In fact, the Canadian Investment Funds Standards Committee (CIFSC) classifies all Canadian funds into various categories based on each fund's portfolio holdings.

Fixed income funds provide a good example of how holdings-based analysis can provide important information about the risk of a fund that cannot be discerned from historical return-based measures (like SD and CVaR) alone. Please see Appendix B for a useful analysis of various CIFSC fixed income categories that illustrate this point.

6. We believe that SD can be applied to a range of fund types (asset class exposures, fund structures, manager strategies, etc.). Keeping the criteria outlined in the introduction above in mind, would you recommend a different Volatility Risk measure for any specific fund products? Please supplement your recommendations with data/analysis wherever possible.

MORNINGSTAR COMMENT: We agree that SD, along with other risk measures such as CVaR, can be applied to a range of fund types.

7. We understand that it is industry practice (for investment fund managers and third party data providers) to use monthly returns to calculate SD. Keeping the criteria outlined in the introduction above in mind, would you suggest that an alternative frequency be used? Please specifically state how a different frequency would improve fund risk disclosure and be of benefit to investors. Please supplement your recommendations with data/analysis wherever possible.

MORNINGSTAR COMMENT: Morningstar believes that monthly returns are appropriate to calculate SD and other historically based risk measures such as CVaR.

8. Keeping the criteria outlined in the introduction above in mind, should we consider a different time period than the proposed 10-year period as the basis for risk rating disclosure? Please explain your reasoning and supplement your recommendations with data/analysis wherever possible.

MORNINGSTAR COMMENT: Morningstar believes that the proposed time period will cover most of the oldest share classes and is therefore appropriate.

9. Keeping the criteria outlined in the introduction above in mind, should we consider an alternative approach to the calculation by series/class? Please supplement your recommendations with data/analysis wherever possible.

MORNINGSTAR COMMENT: Morningstar does not believe that it is necessary to apply the CSA's proposed methodology to individual series/classes. Each series/class of a fund has identical fund holdings and therefore bears equivalent levels of risk. While it is true that returns vary by series/class, differences in SD are slight to non-existent.

10. Keeping the criteria outlined in the introduction above in mind, do you agree with the criteria we have proposed for the use of a reference index for funds that do not have sufficient historical performance data? Are there any other factors we should take into account when selecting a reference index? Please supplement your recommendations with data/analysis wherever possible.

MORNINGSTAR COMMENT: Morningstar agrees with the CSA's reference-index criteria. Relying on correlations alone could provide an incomplete picture. By using a combination of returns- and holdings-based approaches, the reference index is more likely to match a fund's risk profile. Morningstar believes that the CSA should consider CIFSC category-based benchmarks as potential proxies because they are better proxies for the investor experience than market-based benchmarks.

11. Keeping the criteria outlined in the introduction above in mind,
 - i. Do you agree with the proposed number of risk bands, the risk band break-points, and nomenclature used for risk band categories?
 - ii. Do the proposed break points allow for sufficient distinction between funds with varying asset class exposures/risk factors?

If not, please propose an alternative, and indicate why your proposal would be more meaningful to investors. Please supplement your recommendations with data/analysis wherever possible.

MORNINGSTAR COMMENT: Morningstar does not have a strong opinion on the proposed number of risk bands and the risk band break-points but we fundamentally disagree with the CSA's proposal to fix the risk-band break points.

The fundamental problem is that values of the ranges were presumably selected to represent the riskiness of specific asset classes over some historical period, but there is no guarantee that the values will continue to do so in the future, as the risk levels of asset classes change over time. For this reason we would favour a system with floating risk bands.

We illustrate the problems with fixing the risk band break-points to a given set of values in Exhibit 2 below. Exhibit 2 shows the 10-year rolling annualized SD of the following four equity market indexes using the formula shown in section 5 of the Annex over the period of January 1990 to December 2013:

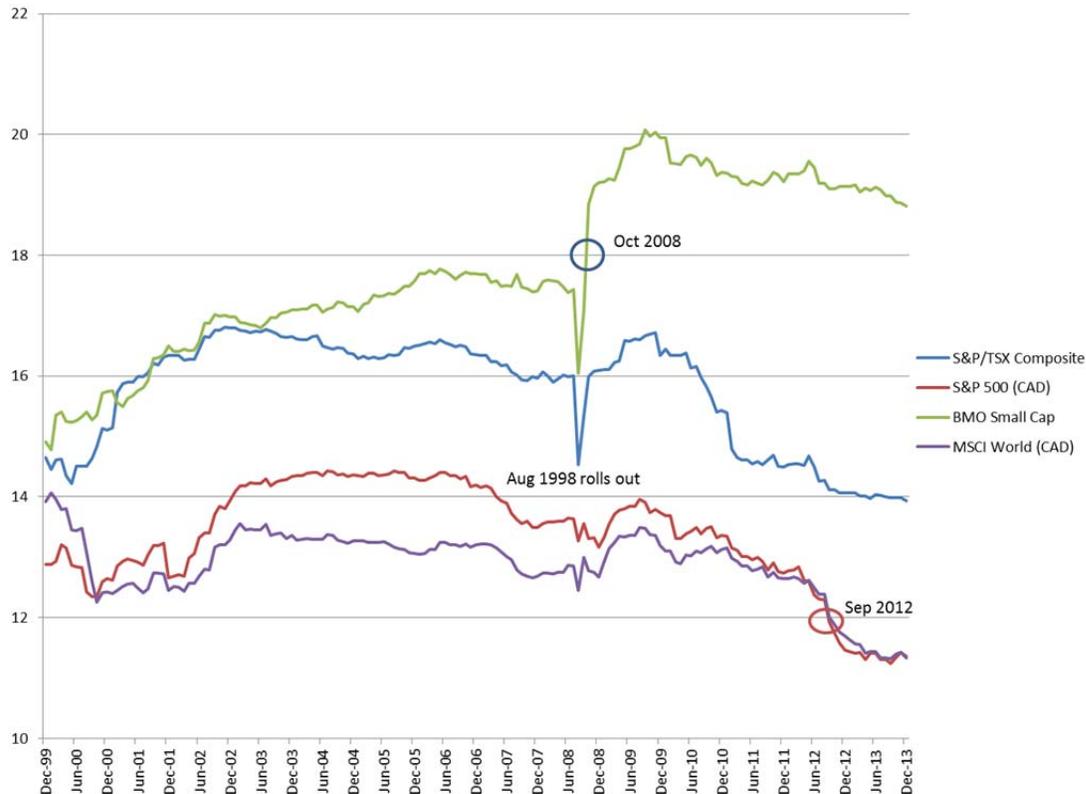
- S&P/TSX Composite
- S&P 500 (converted to CAD at Bank of Canada spot rates)
- BMO Small Cap Blended (Weighted)
- MSCI World (in CAD converted at spot prices)

What we see is that although the rolling SD of the S&P/TSX Composite remained within the "Medium to High" range of 12-18%, it did so with considerable variation within that range. The rolling SD of the BMO Small Cap Index had even more variation and crossed over from the "Medium to High" range to the "High" range in October 2008 and has remained there ever since, so that its 12-month average risk classification eventually changed as well. The rolling SDs of the S&P 500 and MSCI World indexes fell below 12% in September 2012, moving these indexes from the "Medium to High" range into the "Medium" range where they have remained ever since, so that their 12-month average risk classifications eventually changed as well.

Further complicating matters is the way in which a single month's return can cause changes to rolling SDs as illustrated by the sudden drops in the 10-year rolling SDs of the S&P/TSX Composite when the extreme negative returns of August 1998 roll out in August 2008 and then spike up during the market crash that started in September 2008.

It is important to note that the changes in risk classifications do not represent changes in the investment strategies of the indexes and hence do not represent any changes to the exposures to risk factors.

Exhibit 2: Rolling 10-Year Annualized SDs of Equity Market Indexes



12. Do you agree with the proposed process for monitoring risk ratings? Keeping the criteria outlined in the introduction above in mind, would you propose a different set of parameters or different frequency for monitoring risk rating changes? If yes, please explain your reasoning. Please supplement your recommendations with data/analysis wherever possible.

MORNINGSTAR COMMENT: Morningstar believes that the approach suggested by the CSA is reasonable.

13. Is a 10-year record-retention period too long? If yes, what period would you suggest instead and why?

MORNINGSTAR COMMENT: Morningstar has no comment to make on record-retention periods.

14. Please comment on any transition issues that you think might arise as a result of risk classification changes that are likely to occur upon the initial application of the Proposed Methodology. How would fund managers and dealers propose to minimize the impact of these issues?

MORNINGSTAR COMMENT: Morningstar has no comment on transition issues.

In conclusion, Morningstar recognizes that the current lack of objective and transparent standards makes it difficult for investors to make meaningful comparisons or for the classifications to be verified independently. The challenge, though, with a standard risk indicator is that it can suggest to an investor that it is the only measure of risk that need be considered.

Should the CSA proceed with mandating a standard risk-classification methodology for use by mutual fund managers in the Fund Facts document, Morningstar calls on it to rethink its proposed adoption of SD. We strongly urge the CSA to consider CVaR, which focuses on the likelihood and magnitude of the extreme losses that investors regard as the most meaningful form of risk, as well as holdings-based measures, and to consider the idea of blending returns-based and holdings-based approaches.

Sincerely,

Morningstar Research Inc.

A handwritten signature in blue ink, appearing to read 'Scott C. Mackenzie', written over a horizontal line.

Scott C. Mackenzie
President and CEO

Appendix A

CVaR and SD are not proxies for each other. We dispel this misconception by presenting two funds that have similar SDs but very different CVaRs. (See Exhibit A1.)

The case for CVaR begins with an explanation of Value at Risk (VaR). Value at Risk describes the left tail of a fund's distribution of returns and how much capital can be lost over a given period. A 5% VaR means that there is a 5% chance of losing a certain dollar amount or more in 12 months. Applying this idea to returns, the 5% VaR is the negative of the fifth percentile of the return distribution. For example, suppose that the fifth percentile of the monthly return distribution of a fund is minus 6%. If so, its one-month 5% VaR is 6%. This means there is a 5% chance of losing \$600 or more on a \$10,000 investment over one month.

Conditional Value at Risk (CVaR) goes further and accounts for possible losses beyond VaR; it is the expected or average loss of capital that would occur should VaR be breached. Therefore, CVaR is always greater than VaR. CVaR is a measure of tail risk, focusing entirely on possible losses. For example, a one-month 5% CVaR of 10% means that the average one-month loss in the 5% of worst cases is 10%.

Exhibit A1 demonstrates that CVaR is a more revealing measure of risk than SD and why CVaR and SD are not proxies for each other. Exhibit A1 shows summary statistics on two actual equity funds that are sold in Canada that have similar 10-year annualized SD but different levels of one-month 5% CVaRs, which we refer to as Fund A and Fund B. All of the statistics are based on the 120 monthly total returns of the funds over the period from January 2004 to December 2013. Exhibit A1 shows that, while Fund A has a similar SD to Fund B, it has a much higher one-month 5% CVaR.⁶

⁶ We calculated VaR and CVaR in Microsoft Excel using Visual Basic macros that implement the methodology for calculating these statistics presented in Appendix 26A of *Frontiers of Modern Asset Allocation* by Paul D. Kaplan (John Wiley & Sons, 2012).

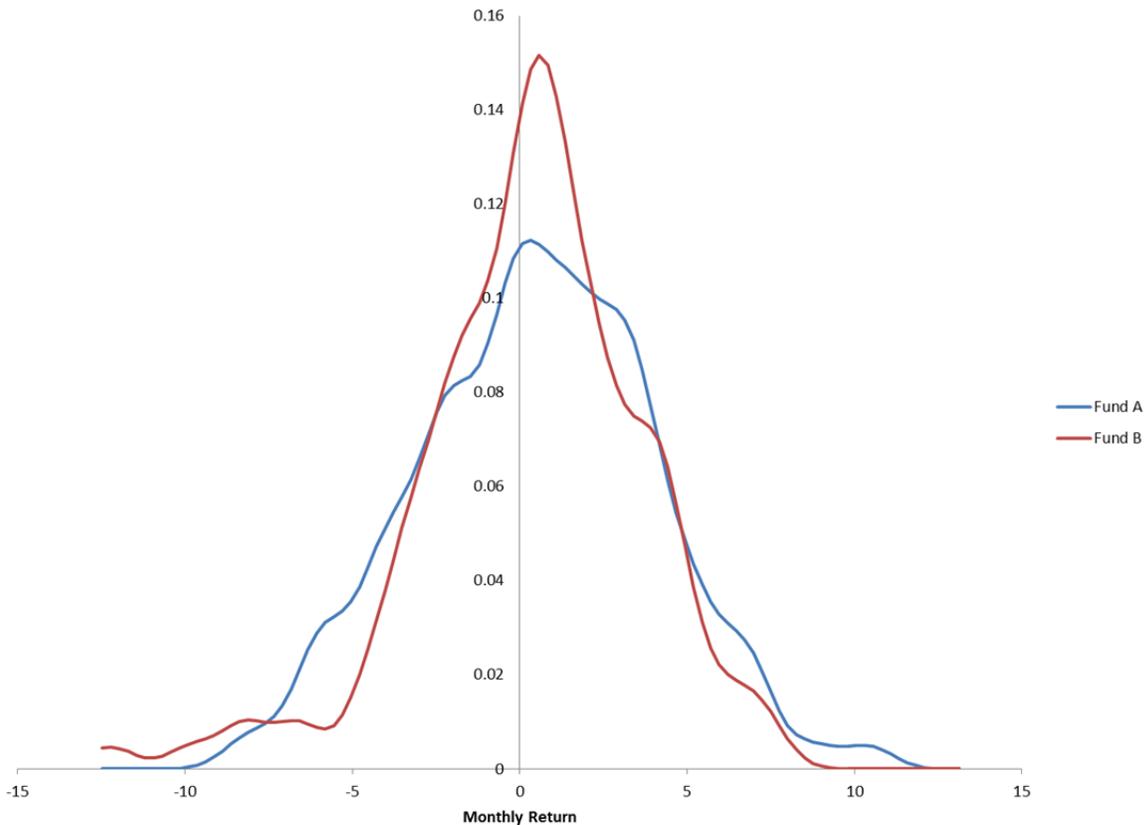
Exhibit A1: Summary Statistics on Example Funds: January 2004—December 2013

	Fund A	Fund B
Category:	Health Care Equity	Global Equity
Compound Annual Return (%):	5.43	1.02
Monthly Average Return (%):	0.50	0.15
Annualized SD (%):	12.38	12.49
Skewness:	0.00	-1.10
Excess Kurtosis:	-0.13	2.87
1-Month 5% VaR (%):	5.68	6.58
1-Month 5% CVaR (%):	6.85	10.06

In order to understand why these funds have such different CVaRs when they have such similar SDs, we need to look at the distributions of their returns. In Exhibit A2 we present the distribution curves of the two funds.⁷ We see that the left tails of their return distributions are quite different. While the left tail of Fund A is quite thin, the left tail of Fund B is quite fat. In other words, while the left tail of Fund A is like that of a normal distribution in that there are effectively no outcomes beyond three SDs to the left of the mean, the left tail of Fund B contains outcomes well beyond that point. Specifically, the worst three monthly returns on Fund A were -8.27%, -7.42%, and -6.20%, which are 2.45, 2.22, and 1.88 SDs below the mean, respectively. This pattern is typical of a thin-tailed distribution in that there are no returns worse than three SDs below the mean. In contrast, the worst three monthly returns on Fund B were -15.69%, -12.23%, and -9.81%, which are 4.12, 3.43, and 2.76 SDs below the mean, respectively. Such occurrences of returns that fall below three SDs below the mean are typical of fat-tailed distributions and explain why there can be significant differences in CVaRs (6.85% vs. 10.06% per month) when there is little difference in SDs.

⁷ We created these distribution curves using the methodology presented in *ibid.*

**Exhibit A2: Monthly Return Distributions on Example Funds:
January 2004—December 2013**



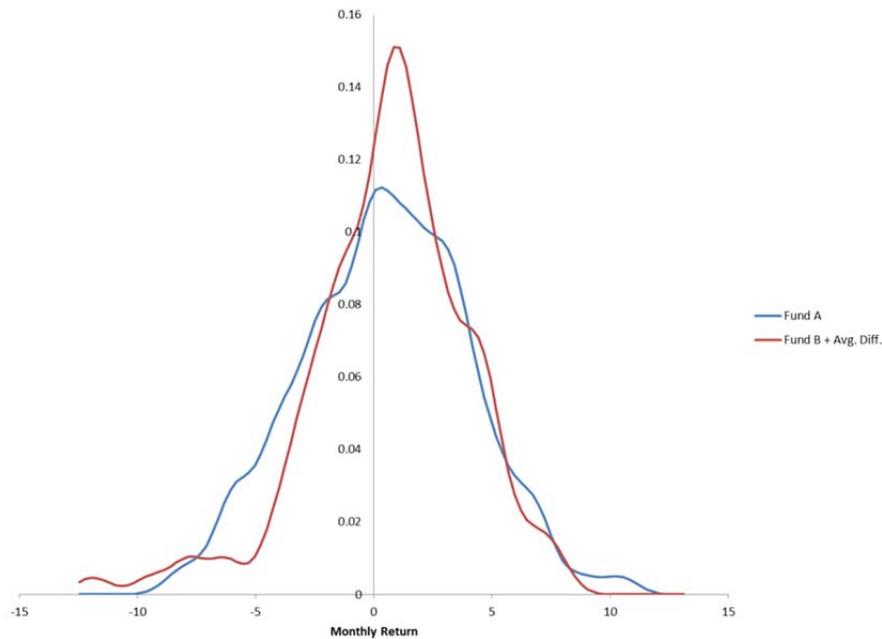
The difference in the CVaRs of Fund A and Fund B is due in part to Fund B's lower average monthly return. To take this into account, we formed a hypothetical monthly return series by adding the difference in the average monthly returns between Fund A and Fund B ($0.50\% - 0.15\% = 0.35\%$) to each of the monthly returns on Fund B. Exhibit A3 shows the summary statistics of Fund A and this hypothetical version of B. Tellingly, while the hypothetical fund has a compound annual return that is almost the same as that of Fund B, its CVaR remains substantially higher than that of Fund A.

In Exhibit A4 we present the distribution curves of Fund A and the hypothetical version of Fund B. We see that the left tails of their return distributions remain quite different. The left tail of the hypothetical version of Fund B contains returns that are substantially lower than Fund A's worst return. Recall that the worst return on Fund A was -8.27% , which is 2.45 SDs below the mean. In contrast, the worst three monthly returns of Fund B's hypothetical version were -14.34% , -11.87% , and -9.46% which are 4.12, 3.43, and 2.76 SDs below the mean. Hence, even after adjusting the returns on Fund B to bring the average return and compound annual return into line with Fund A, the left tails of the return distributions still present significant differences in CVaRs (6.85% vs. 9.70% per month).

**Exhibit A3: Summary Statistics on Example Funds with Average Difference Adjustment:
January 2004—December 2013**

	Fund A	Fund B + Average Difference
Category:	Health Care Equity	Global Equity
Compound Annual Return (%):	5.43	5.40
Monthly Average Return (%):	0.50	0.50
Annualized SD (%):	12.38	12.49
Skewness:	0.00	-1.10
Excess Kurtosis:	-0.13	2.87
1-Month 5% VaR (%):	5.68	6.23
1-Month 5% CVaR (%):	6.85	9.70

**Exhibit A4: Monthly Return Distributions on Example Funds with Average Difference Adjustment:
January 2004—December 2013**



Appendix B

A major source of risk for fixed income funds is interest rate risk. A standard measure of the interest rate sensitivity of a fixed income security is modified duration. To measure the interest rate sensitivity of a fixed income fund, Morningstar calculates the asset-weighted average modified duration of the fund's underlying securities. For domestic investment-grade fixed income funds, interest rate risk is the primary source of risk.

Exhibit B1 plots 10-year annualized SD against the average modified duration of 133 distinct Canadian fixed income funds.⁸ Of these, there are 18 funds in the Global Fixed Income category, 22 in the High Yield Fixed Income category, and 93 from various domestic investment-grade categories. For the domestic investment-grade funds, there is a strong correlation between historical SD and current average duration⁹. Hence for these funds, historical SD is a good indicator of risk. However, for the global and high yield funds, there is no relationship between SD and average current duration. Furthermore, the SDs vary over a wide range across the funds, indicating that these funds are inherently subject to varying amounts of either exchange-rate risk (global funds) or credit risk (high-yield funds) that cannot be discerned without examining the portfolio holdings.

⁸ The 10-year SD is measured over the period ended Dec. 31, 2013.

⁹ The R^2 of the regression line is 90%, so that the correlation is $\sqrt{0.9} = 95\%$.

Exhibit B1: SD vs. Current Average Duration on Fixed Income Funds