
The Morningstar Rating™ for Funds

Morningstar Methodology

November 2016

Contents

- 1 Introduction
- 2 Morningstar Categories
- 4 Theory
- 8 Calculations
- 12 The Morningstar Rating: Three-, Five-, and 10-Year
- 13 Morningstar Return and Morningstar Risk Rating
- 15 The Overall Morningstar Rating
- 16 Rating Suspensions
- 17 Conclusion
- 18 Appendix 1: Risk-Free Rates Applied
- 19 Appendix 2: Methodology Changes

Introduction

This document describes the rationale for, and the formulas and procedures used in, calculating the Morningstar Rating™ for funds (commonly called the “star rating”). This methodology applies to funds receiving a star rating from Morningstar, except in Japan where these are the Ibbotson Stars.

The Morningstar Rating has the following key characteristics:

- ▶ The peer group for each fund’s rating is its Morningstar Category™.
- ▶ Ratings are based on funds’ risk-adjusted returns.

Morningstar Category

The original Morningstar Rating was introduced in 1985 and was often used to help investors and advisors choose one or a few funds from among the many available within broadly defined asset classes. Over time, though, increasing emphasis had been placed on the importance of funds as portfolio components rather than “stand-alone” investments. In this context, it was important that funds within a particular rating group be valid substitutes for one another in the construction of a diversified portfolio. For this reason, Morningstar now assigns ratings based on comparisons of all funds within a specific Morningstar Category, rather than all funds in a broad asset class.

Risk-Adjusted Return

The star rating is based on risk-adjusted performance. However, different aspects of portfolio theory suggest various interpretations of the phrase “risk-adjusted.” As the term is most commonly used, to “risk adjust” the returns of two funds means to equalize their risk levels before comparing them. The Sharpe ratio is consistent with this interpretation of “risk-adjusted.”

But the Sharpe ratio does not always produce intuitive results. If two funds have equal positive average excess returns, the one that has experienced lower return volatility receives a higher Sharpe ratio score. However, if the average excess returns are equal and negative, the fund with higher volatility receives the higher score because it experienced fewer losses per unit of risk. While this result is consistent with portfolio theory, many retail investors find it counterintuitive. Unless advised appropriately, they may be reluctant to accept a fund rating based on the Sharpe ratio, or similar measures, in periods when the majority of the funds have negative excess returns.

Standard deviation is another common measure of risk, but it is not always a good measure of fund volatility or consistent with investor preferences. First, any risk-adjusted return measure that is based on standard deviation assumes that the riskiness of a fund’s excess returns is well

captured by standard deviation, as would be the case if excess return were normally or lognormally distributed, which is not always the case. Also, standard deviation measures variation both above and below the mean equally. But investors are generally risk-averse and dislike downside variation more than upside variation. Morningstar gives more weight to downside variation when calculating Morningstar Risk-Adjusted Return and does not make any assumptions about the distribution of excess returns.

The other commonly accepted meaning of “risk-adjusted” is based on assumed investor preferences. Under this approach, higher return is “good” and higher risk is “bad” under all circumstances, without regard to how these two outcomes are combined. Hence, when grading funds, return should be rewarded and risk penalized in all cases. The Morningstar Risk-Adjusted Return measure described in this document has this property.

This document discusses the Morningstar Category as the basis for the rating, and it describes the methodology for calculating risk-adjusted return and the Morningstar Rating. Morningstar calculates ratings at the end of each month.

Morningstar Categories

Category Peer Groups

Morningstar uses the Morningstar Category as the primary peer group for a number of calculations, including percentile ranks, fund-versus-category-average comparisons, and the Morningstar Rating. The Morningstar Rating compares funds’ risk-adjusted historical returns. Its usefulness depends, in part, on which funds are compared with others.

It can be assumed that the returns of major asset classes (domestic equities, foreign equities, domestic bonds, and so on) will, over lengthy periods of time, be commensurate with their risk. However, asset class relative returns may not reflect relative risk over ordinary investor time horizons. For instance, in a declining interest-rate environment, investment-grade bond returns can exceed equity returns despite the higher long-term risk of equities; such a situation might continue for months or even years. Under these circumstances many bond funds outperform equity funds for reasons unrelated to the skills of the fund managers.

A general principle that applies to the calculation of fund star ratings follows from this fact; that is, the relative star ratings of two funds should be affected more by manager skill than by market circumstances or events that lie beyond the fund managers’ control.

Another general principle is that peer groups should reflect the investment opportunities for investors. So, categories are defined and funds are rated within each of the major markets around the world. Morningstar supports different category schemes for different markets based on the investment needs and perspectives of local investors. For example, Morningstar rates high-yield

bond funds domiciled in Europe against other European high-yield bond funds. For more information about available categories, please contact your local Morningstar office.

Style Profiles

A style profile may be considered a summary of a fund's risk-factor exposures. Fund categories define groups of funds whose members are similar enough in their risk-factor exposures that return comparisons between them are useful.

The risk factors on which fund categories are based can relate to value-growth orientation; capitalization; industry sector, geographic region, and country weights; duration and credit quality; historical return volatility; beta; and many other investment style factors. The specific factors used are considered to be a) important in explaining fund-return differences and b) actively controlled by the fund managers.

Because the funds in a given category are similar in their risk-factor exposures, the observed return differences among them relate primarily to security selection ("stock-picking") or to variation in the timing and amount of exposure to the risk factors that collectively define the category ("asset weighting"). Each of these, over time, may be presumed to have been a skill-related effect.

Note that if all members of a fund category were uniform and consistent in their risk factor exposures, and the risk factors were comprehensive, there would be no need to risk-adjust returns when creating category-based star ratings. However, even within a tightly defined category, the risk exposures of individual funds vary over time. Also, no style profile or category definition is comprehensive enough to capture all risk factors that affect the returns of the funds within a category.

In extreme cases where the funds in a category vary widely in their risk factor exposures (that is, it is a "convenience category"), a star rating would have little value and is not assigned. For example, in the United States, ratings are not assigned to funds in the bear-market category because these funds short very different parts of the market. In Europe, ratings are not assigned to funds in the guaranteed category.

Defining Fund Categories

The following considerations apply when Morningstar defines fund categories:

- ▶ Funds are grouped by the types of investment exposures that dominate their portfolios.
- ▶ In general, a single return benchmark should form a valid basis for evaluating the returns for all funds in a single category (that is, for performance attribution).
- ▶ In general, funds in the same category can be considered reasonable substitutes for the purposes of portfolio construction.
- ▶ Category membership is based on a fund's long-term or "normal" style profile, based on three years of portfolio statistics. Supplemental analysis includes returns-based style analysis, review of strategy disclosure from fund literature, and qualitative review by analysts.

Theory

Expected Utility Theory

Morningstar Risk-Adjusted Return is motivated by expected utility theory, according to which an investor ranks alternative portfolios using the mathematical expectation of a function (called the utility function) of the ending value of each portfolio. This is a helpful framework to model decision-making under uncertainty.

Let W be the ending wealth within a portfolio being considered and $u(\cdot)$ be the investor's utility function. The expected utility of the portfolio is $E[u(W)]$.

The form of the utility function that is used often in portfolio theory has the following characteristics:

1. More expected wealth is always better than less expected wealth.
This means that the utility function must always be positively sloped, so $u'(\cdot) > 0$.
2. The utility function must imply risk aversion, and risk is always penalized.
The investor prefers a riskless portfolio with a known end-of-period value to a risky portfolio with the same expected value. For example, a fund that produces a steady 2% return each month is more attractive than a fund that has volatile monthly returns that average out to 2% per month. This can be written as:

$$[1] \quad u(E[W]) > E[u(W)]$$
 From probability theory, it follows that this can be true only if $u(\cdot)$ is everywhere a concave function, so $u''(\cdot) < 0$.
3. No particular distribution of excess returns is assumed.
Expected utility theory does not rely on any assumptions about whether a fund's returns distribution, other than that it be well-behaved, is normally or lognormally distributed. This is in contrast to other measures of risk-adjusted return that use standard deviation or variance as the main measure of risk. While many funds' returns are approximately lognormally distributed, utility theory will also work for those that are not, such as funds that use extensive options strategies.
4. The investor's beginning-of-period wealth has no effect on the ranking of portfolios.
It is reasonable to assume that the investor's risk aversion does not change with the level of investor wealth, that is, those more-wealthy individuals are not universally more or less risk-averse than less-wealthy individuals. Individuals with the same attitudes toward risk and the same opportunity set will choose the same investments, regardless of their level of wealth.

One form of a utility function that has these characteristics and that is used often in portfolio theory is called "constant relative risk aversion." Relative risk aversion (RRA) describes the degree to which wealth affects an investor's level of risk aversion, and this is measured based on the shape of the utility function with respect to wealth:

$$[2] \quad RRA(W) = -\frac{Wu''(W)}{u'(W)}$$

By assuming that RRA is a constant value (that is, the level of wealth will not change the investor's attitude toward risk), the equations for the utility function can be written as follows:

$$[3] \quad u(W) = \begin{cases} -\frac{W^{-\gamma}}{\gamma} & \gamma > -1, \gamma \neq 0 \\ \ln(W) & \gamma = 0 \end{cases}$$

where: γ is a parameter that describes the degree of risk aversion, specifically, $RRA(.) = \gamma + 1$.

Because end-of-period wealth (W) is a function of beginning wealth and total return, these equations can be rewritten as follows, where there is a certain level of utility associated with each level of total return.

$$[4] \quad u(W_0(1+TR)) = \begin{cases} W_0^{-\gamma}u(1+TR) & \gamma > -1, \gamma \neq 0 \\ \ln(W_0) + u(1+TR) & \gamma = 0 \end{cases}$$

where:

W_0 = beginning-of-period wealth

TR = total return on the portfolio being evaluated so that $W = W_0(1+TR)$

The value of W_0 does not affect the curvature of utility as a function of TR , and so it does not affect how the investor ranks portfolios.

Degree of Risk Aversion

Gamma (γ) represents the degree of risk aversion. In theory, it can be any number of values.

When γ is less than -1 , the investor is risk-loving, rather than risk-averse. This investor might be indifferent between a steady fund that always earns 2.5% each month and a volatile fund that is expected to earn 2% on average each month. This investor likes risk.

When γ is -1 , the degree of risk aversion is zero, meaning that the investor is indifferent between a riskless choice and a risky choice as long as the arithmetic average expected return is the same. This investor is indifferent between a steady fund that always earns 2% per month and a volatile fund that is expected to earn 2% on average (for example, equal likelihood of negative 4%, 2%, or 8% each month), even though the volatile fund could lose money.

When γ is 0, the investor is indifferent between a riskless choice and a risky choice as long as the geometric average expected return is the same. This investor is indifferent between a steady fund that always earns 1.88% and a volatile fund expected to earn 2% on average, with an equal likelihood of negative 4%, 2%, or 8% each month. (The geometric average of those volatile expected returns is 1.88%.) An initial investment in either portfolio is expected to grow to the same amount after one year.

The risk premium is the amount of extra expected return demanded by the investor to compensate for the possibility of losing money in the risky portfolio versus the riskless portfolio. When γ is 0, this investor requires a risk premium of 0.12% per month, the difference between the arithmetic average return of the risky portfolio and the riskless return. In this case, the riskless return is the same as the geometric average return.

When γ is greater than 0, the investor demands a larger risk premium for choosing the risky portfolio. Specifically, the risk premium must be larger than the difference between the arithmetic and geometric average returns. With $\gamma=2$, the investor is indifferent between a steady fund that always earns 1.65% per month and the volatile fund above that is expected to earn 2% on average, with equal likelihood of obtaining negative 4%, 2%, or 8% each month. In this case, the risk premium is 0.35% per month.

In practice, most models assume investors are risk-averse and therefore, γ must be greater than negative 1.

Morningstar's Formulation of Utility Theory

Morningstar uses expected utility theory with a few specific conditions as the basis for Morningstar Risk-Adjusted Return. Morningstar recognizes that the investor always has a choice to buy a risk-free asset instead of holding a risky portfolio. Therefore, Morningstar measures a fund's excess returns over and above the risk-free rate (RF). In comparing risky portfolios to the risk-free asset, we assume that the investor initially has all wealth invested in the risk-free asset and beginning-of-period wealth is such that end-of-period wealth, so invested, will be USD 1.

Hence:

$$[5] \quad W_0 = \frac{1}{1 + RF}$$

The utility function can be restated in terms of total return (TR), the risk-free rate (RF), and geometric excess returns (ER) as follows:

$$[6] \quad u(W_0(1 + TR)) = \left(\frac{1 + TR}{1 + RF} \right) = u(1 + ER) = \begin{cases} -\frac{(1 + ER)^{-\gamma}}{\gamma} & \gamma > -1, \gamma \neq 0 \\ \ln(1 + ER) & \gamma = 0 \end{cases}$$

where:

$$[7] \quad ER = \text{the geometric excess return} = \frac{1+TR}{1+RF} - 1$$

Applying expected utility theory to risk-adjusted return implies that it is possible to quantify how investors feel about one distribution of returns versus another. A return distribution with high expected return and low risk is preferable to one with low expected return and high risk. But investors typically face a trade-off between risk and return. At some point, the level of risk becomes too high and the investor is willing to settle for a lower expected return to reduce risk. (Or, the level of expected return becomes too low and the investor is willing to take on more risk in order to potentially achieve higher returns.)

Morningstar uses expected utility theory to determine how much return a model investor is willing to trade off, reducing the risk of loss. Morningstar Risk-Adjusted Return measures the guaranteed riskless return that provides the same level of utility to the investor as the variable excess returns of the risky portfolio. We call this riskless return the “certainty equivalent” geometric excess return.

For example, an investor might be indifferent between a moderately risky fund generating 12% return (what we observe) and a riskless fund generating 8% return (as determined by the utility function). In that case, the investor is willing to give up 4% in return in order to remove the risk. By converting all return series to their riskless equivalents, Morningstar can compare one fund with another on a risk-adjusted basis. This equalizes the playing field for funds in the same category that have different exposures to risk factors.

Let $ER^{CE}(\gamma)$ denote the certainty equivalent geometric excess return for a given value of γ . The following formula states that the level of utility is the same between the certainty equivalent geometric excess return and the expected excess returns of the fund:

$$[8] \quad u(1+ER^{CE}(\gamma)) = E[u(1+ER)]$$

Hence:

$$[9] \quad 1 + ER^{CE} = \begin{cases} (E[(1+ER)^{-\gamma}])^{-\frac{1}{\gamma}}, & \gamma > -1, \gamma \neq 0 \\ e^{E[\ln(1+ER)]}, & \gamma = 0 \end{cases}$$

Morningstar defines Morningstar Risk-Adjusted Return, $MRAR(\gamma)$, as the annualized value of the certainty equivalent, ER^{CE} , using the time series average of $(1+ER)^{-\gamma}$ as an estimate of $E[(1+ER)^{-\gamma}]$. That is, Morningstar uses historical excess returns as the basis for expected excess returns rather than relying on analysts’ forecasts or other probabilities of future returns.

With $\gamma \neq 0$, Morningstar Risk-Adjusted Return is defined as follows:

$$[10] \quad MRAR(\gamma) = \left[\frac{1}{T} \sum_{t=1}^T (1 + ER_t)^{-\gamma} \right]^{\frac{12}{\gamma}} - 1$$

where:

ER_t = the geometric excess return in month $t = -1 \frac{1+TR_t}{1+RF_t}$

TR_t = total return for the fund in month t

RF_t = return for the risk-free asset in month t

T = the number of months in the time period

When $\gamma=0$, MRAR is the annualized geometric mean of the excess returns:

$$[11] \quad MRAR(0) = \left[\prod_{t=1}^T (1 + ER_t) \right]^{\frac{12}{T}} - 1$$

A rating system based solely on performance would rank funds on their geometric mean return, or equivalently, MRAR(0) or Morningstar Return. A rating system that provides a heavier penalty for risk requires that $\gamma > 0$.

Morningstar's analysts have concluded that $\gamma=2$ results in fund rankings that are consistent with the risk tolerances of typical retail investors. Hence, Morningstar uses a γ equal to two in the calculation of its star ratings.

Because MRAR is expressed as an annualized return, it can be decomposed into a return component, Morningstar Return or MRAR(0), and a risk component, Morningstar Risk. Morningstar Risk is calculated as $MRAR(0) - MRAR(2)$, or Morningstar Return – Morningstar Risk-Adjusted Return.

Calculations

Overview

There are three steps to calculate Morningstar Risk-Adjusted Return. The calculations are done on a monthly basis first and then the results are annualized.

1. Total Return: Calculate monthly total returns for the fund. Do an additional adjustment for tax-advantaged dividends where appropriate.
2. Morningstar Return: Calculate or collect monthly total returns for the appropriate risk-free rate. Adjust returns for the risk-free rate to get Morningstar Return.
3. Morningstar Risk-Adjusted Return: Adjust Morningstar Return for risk to get MRAR. Morningstar Risk is then calculated as the difference between Morningstar Return and Morningstar Risk-Adjusted Return.

The annualized returns are the same or lower after each adjustment, as shown below.

Exhibit 1 3-Year Annualized Returns Are The Same or Lower After Each Adjustment

	Total Return %	Adjust for Risk Free Rate ↓	Morningstar Return %	Adjust for Risk ↓	Morningstar Risk-adj Return %	Morningstar Risk
Focus Fund A	6.60		6.53		5.55	0.98
Growth Fund Inv	10.21		10.14		8.72	2.11

Morningstar calculates percentile ranks in category for all of these data points. By studying these percentile ranks, one can determine which factor had the most impact on the fund's rating.

Exhibit 2 3-Year Percentile Ranks Decomposition

	Total Return % Rank	Morningstar Return % Rank	Morningstar Risk-adj Return % Rank	Morningstar Risk
Focus Fund A	83	6.53	5.55	0.98
Growth Fund Inv	51	10.14	8.72	2.11

Total Return

Morningstar calculates a fund's total return for a given month, t, as follows:

$$[16] \quad TR_t = \left\{ \frac{P_e}{P_b} \prod_{i=1}^n \left(1 + \frac{D_i}{P_i} \right) \right\} - 1$$

where

TR_t = total return for the fund for month t

P_e = end of month NAV per share

P_b = beginning of month NAV per share

D_i = per share distribution at time i

P_i = reinvestment NAV per share at time i

n = number of distributions during the month

Distributions include dividends, distributed capital gains, and return of capital. This calculation assumes that the investor incurs no transaction fees and reinvests all distributions paid during the month.

The cumulative total return is:

$$[17] \text{TR}_c = \prod_{t=1}^T (1 + \text{TR}_t) - 1$$

where

TR_c = cumulative total return for the fund

TR_t = total return for the fund for month t

T = number of months in the period (for example, three, five, or 10 years)

Tax Adjustment

In reality, an investor's total return will be reduced by any taxes that must be paid on income and capital gains. Morningstar does not adjust for these taxes for MRAR and the Morningstar Rating, because one single tax rate does not reflect the experience of all investors.

Morningstar will adjust for taxes in the MRAR calculation if most investors in that fund qualify for the same tax treatment. For example, Morningstar adjusts the dividends paid by U.S. municipal-bond funds to reflect their exemption from U.S. federal taxes. For single-state municipal-bond funds, there is an additional adjustment for state taxes. Morningstar adjusts these dividends to an equivalent pretax level for the purpose of calculating risk-adjusted return. This adjustment will make a difference for funds that distribute income only versus funds that distribute a combination of income and capital gains.

Morningstar adjusts these municipal-bond dividends with the following formula:

$$[18] \text{TDiv}_i = \frac{\text{Div}_i}{(1 - t_{Si})(1 - t_{Fi})}$$

where

TDiv_i = tax-adjusted dividend per share at time i

Div_i = actual dividend per share at time i

t_{Si} = maximum state tax rate at time i (for single-state municipal-bond funds)

t_{Fi} = maximum federal tax rate at time i

For the purpose of calculating the total returns that are used to calculate MRAR, Morningstar uses TDiv_i in place of Div_i in equation X for U.S. municipal-bond funds. This tax adjustment is not part of the standard depictions of total return for these funds.

Morningstar Return

Next, Morningstar adjusts the fund's monthly returns for the risk-free rate. Because investors always have an option to invest at the risk-free rate, Morningstar measures only the amount by which fund returns have exceeded that risk-free rate. This adjustment also accounts for how the risk-free rate has changed over time.

For each historical month, Morningstar calculates the fund's geometric excess return over the risk-free rate.

$$[19] \quad ER_t = \frac{1 + TR_t}{1 + RF_t} - 1$$

where

ER_t = the geometric excess return for the fund for month t

TR_t = the total return for the fund for month t

RF_t = the total return for the risk-free rate for month t

Morningstar selects a risk-free rate that is appropriate for the investor, and this varies for different Morningstar offices around the world. The risk-free rate is selected based on the primary currency of the investment, rather than where the fund invests. The attached appendix includes the risk-free rates applied by currency.

The annualized geometric mean of these excess returns is known as Morningstar Return.

$$[20] \quad \text{Morningstar Return} = \left[\prod_{t=1}^T (1 + ER_t) \right]^{\frac{12}{T}} - 1$$

where

T = number of months in the period (for example, three, five, or 10 years)

Morningstar Risk-Adjusted Return

Next, Morningstar adjusts for risk. As mentioned earlier, Morningstar uses expected utility theory to model how investors trade off return and risk. Morningstar Risk-Adjusted Return is the guaranteed return that provides the same level of utility to the investor as the specific combination of returns exhibited by the fund.

The formal equation for Morningstar Risk-Adjusted Return, equation [10], uses the parameter "gamma" to describe the model investor's sensitivity to risk. Morningstar sets that value equal to 2, so Morningstar Risk-Adjusted Return is calculated as follows:

$$[21] \quad \text{MRAR}(2) = \left[\frac{1}{T} \sum_{t=1}^T (1 + ER_t)^{-2} \right]^{\frac{12}{2}} - 1$$

The section inside the brackets determines the investor's average utility from this fund's monthly excess returns over 36, 60, or 120 months. Then, that level of utility is converted into a return by taking it to the power of $-1/2$. Lastly, Morningstar annualizes the result by taking it to the power of 12.

Morningstar Risk

Because MRAR is expressed as an annualized return, we can derive a risk component, Morningstar Risk, as the difference between Morningstar Return (adjusted for the risk-free rate) and MRAR (adjusted for the risk-free rate and risk). Morningstar Risk is always greater than or equal to zero.

The Morningstar Rating: Three-, Five-, and 10-Year

The Morningstar Rating is based on Morningstar Risk-Adjusted Return, using Morningstar Risk-Adjusted Return % Rank for funds in a category. Morningstar calculates ratings for the three-, five-, and 10-year periods, and then the overall Morningstar Rating is based on a weighted average of the available time-period ratings.

Three-, Five-, and 10-Year Ratings

Investments must have at least 36 continuous months of total returns in order to receive a rating. For each time period (three, five, and 10 years), Morningstar ranks all funds in a category using Morningstar Risk-Adjusted Return, and the funds with the highest scores receive the most stars. A fund's peer group for the three-, five-, and 10-year ratings is based on the fund's current category. That is, there is no adjustment for historical category changes in the three-, five-, and 10-year ratings.

Morningstar rates each share class of a portfolio separately because each share class has different fees and total return time periods available. However, the distribution of funds among the star ratings depends on the number of portfolios evaluated within the category rather than the number of share classes. This policy prevents multishare funds from taking up a disproportionate amount of space in any one rating level. Please refer to the [Morningstar Absolute Ranks, Percentile Ranks and Fractional Ranks methodology](#) document for more information on how Morningstar incorporates fractional weights into percentile ranks so that star ratings are based on distinct portfolios not distinct share classes.

Morningstar sets the distribution of funds across the rating levels, assigning three-year star ratings as follows:

1. All funds in the category are sorted by three-year MRAR % Rank in descending order.
2. Starting with the highest MRAR % Rank, those funds with a rank that meets but does not exceed 10% receive a 5-star rating.
3. Funds with a rank that meets but does not exceed 32.5% receive a 4-star rating.
4. Funds with a rank that meets but does not exceed 67.5% receive a 3-star rating.
5. Funds with a rank that meets but does not exceed 90% receive a 2-star rating.
6. The remaining funds receive 1 star.

If the data are available, five-year ratings are assigned using 60 months of data and 10-year ratings are assigned using 120 months of data.

Morningstar Return and Morningstar Risk Rating

Morningstar uses the same bell curve and rating procedure above to assign scores for Morningstar Return and Morningstar Risk for three, five, and 10 years. Funds are scored from 1 to 5, and these scores are typically expressed as word labels in Morningstar products.

Exhibit 3 Morningstar Risk and Return Rating Legend

Score	Percent	Word Label
5	Top 10%	High
4	Next 22.5%	Above Average
3	Next 35%	Average
2	Next 22.5%	Below Average
1	Bottom 10%	Low

Note that the word label High is generally good for Morningstar Return and Low is generally good for Morningstar Risk.

The Morningstar Return Score and Morningstar Risk Rating are helpful when funds have the same rating and similar MRARs but different levels of risk. All of the funds below got 3 stars and Average return, but they took very different levels of risk to achieve that rating.

Exhibit 4 Different Levels of Risk Achieve The Same Rating

Name	Morningstar Rating 3 Yr	Morningstar Risk-Adj Return 3 Yr	Morningstar Return 3 Yr	Morningstar Return Rating 3 Yr	Morningstar Risk 3 Yr	Morningstar Risk Rating 3 Yr
Fund 1	3	23.45	26.48	Average	3.03	Above Avg
Fund 2	3	23.15	25.05	Average	1.91	Average
Fund 3	3	22.52	24.29	Average	1.78	Below Avg
Fund 4	3	21.28	23.84	Average	2.57	Average
Fund 5	3	20.72	21.64	Average	0.92	Low

Ratings Curve and Ratings Overlay

Morningstar peer groups at times contain multiple legal structures and vehicle types. In many markets there are a number of different legal vehicle types that are ready substitutes for one another despite technical differences. For example, in Europe many retail funds are structured as Fond common du placement, a partnership structure, rather than SICAV, or open-end investment company; the differences are highly technical and will rarely have an impact on the investor's choice. Therefore, we consider both legal structures open-end funds. Similarly, Morningstar now considers exchange-traded funds to be direct substitutes for open-end funds. This treatment has been in place in many regions since 2007, but was first applied to the United States in 2016.

In some situations, fund products are managed similarly, but have a wrapper or structure that makes them poor substitutes. In these cases, Morningstar calculates the MRAR for the primary category peer group and assigns star ratings based on the above ranking calculations. The highest MRAR for each rating is mapped as a breakpoint for other investment types. The MRAR of the substitute product is then mapped to the primary group's breakpoints and assigned a star rating based on an overlay. This is applied to limited distribution collective U.S. investment trusts, as well as extended performance for open-end funds.

The Overall Morningstar Rating

An overall star rating for each fund is based on a weighted average (rounded to the nearest integer) of the number of stars assigned to it in the three-, five-, and 10-year rating periods.

Exhibit 5 Overall Rating Weightings

Months of Total Returns	Overall (Weighted) Morningstar Rating
36–59	100% three-year rating
60–119	60% five-year rating 40% three-year rating
120 or more	50% 10-year rating 30% five-year rating 20% three-year rating

For example, the weighted average of the ratings below is 2.5, and this rounds up to an overall rating of 3 stars.

Exhibit 6 Example Overall Calculation

Period	Rating	Weight %	Contribution
10-year	3	50	1.5
Five-year	2	30	0.6
Three-year	2	20	0.4
Total			2.5

While the 10-year overall star rating formula seems to give the most weight to the 10-year period, the most recent three-year period actually has the greatest impact because it is included in all three rating periods.

Rating Suspensions

Over many years, Morningstar has observed shifts of funds across categories. In general, the category changes do not represent a change in the character of the fund. Generally, short-term bond funds do not become mid-cap-growth funds, nor do U.S. large-blend funds become diversified emerging-markets funds. However, Morningstar reserves the right to suspend a fund's rating if it has undergone a significant change in investment strategy. In these cases, it is misleading to compare the fund's prior performance with its current category. Morningstar also reserves the right to suspend the rating for an actively managed fund that held 100% cash for more than a year after its inception date. (This policy does not apply to money market funds, which are unrated.)

The process for reviewing suspensions is as follows:

- ▶ A fund becomes eligible for a suspension if it has changed broad asset classes or if it has an exceptionally long period of time in cash. A broad asset class is a collection of similar categories, for example, international stock, taxable bond, or balanced. Broad asset classes are defined in the different category systems that are in place in different markets around the world.
- ▶ If a fund is eligible for a suspension, the local research team will review the situation and determine if the suspension should take place. Not all broad asset-class changes will require a rating suspension. For example, a conservative-allocation fund moving to a bond category is not a significant enough change to merit a suspension.

Morningstar will suspend the rating after the strategy change and will mark that suspension date in our systems. Three years after the suspension date, the fund will be eligible for a three-year rating and overall rating. Then, as the fund accumulates 5 and 10 years of performance in the new style, Morningstar will add the five- and 10-year ratings. Morningstar will not suspend percentile ranks or other category comparisons.

Suspended Structures

From time to time Morningstar analysts become aware of novel vehicle structures that are not significant or unique enough to generate the launch of a new Morningstar Category or product universe; however, comparison to other funds in their assigned category would be inappropriate. Some examples are funds only offered to restricted investors (including internal master/feeder), negotiated fee share classes, exemptions on daily dealing, exchange-traded notes, among others. These investments are frequently in our databases with a category assignment of an otherwise rated category, but we have excluded the specific operating attributes from rated products. .

Conclusion

The Morningstar Rating measures how funds have performed on a risk-adjusted basis against their category peers. It gives investors the ability to quickly and easily identify funds that are worthy of further research. The Morningstar Rating is calculated for three years, five years, and 10 years, and the overall rating is a weighted average of the time-period ratings.

Morningstar Risk-Adjusted Return is calculated based on expected utility theory, a framework that recognizes that investors are risk-averse and willing to give up some portion of expected return in exchange for greater certainty of return. Morningstar calculates risk-adjusted return by adjusting total return for the risk-free rate and risk. **MM**

Appendix 1: Risk-Free Rates Applied

Currency	Risk-Free Rate
Australian Dollar	RBA Bank accepted Bills 90 Days
Brazilian Real	Brazil CDI
Canadian Dollar	CIBC WM 91 Day Treasury Bill_CA
Chilean Peso	Chile PDBC 30 Day Monthly
Chinese Yuan Renminbi	RMB 3 month Lump-Sum Deposit
Danish Krone	BofAML DKK LIBOR 1 Mon CM
Euro	BofAML Euro LIBOR 1 Mon CM
Hong Kong Dollar	BofAML HKD LIBOR 1 Mon CM
Indian Rupee	FBIL MIBOR Overnight
Indonesian Rupiah	SBI Rates 3 month
Japanese Yen	BofAML Japanese Yen LIBOR 1 Mon CM
Malaysian Ringgit	BofAML MYR LIBOR 1 Mon CM
Mexican Peso	Cat 80%IMF Mexico T-Bill&20%Pip CETES
New Israeli Shekel	Israel T-Bill 3 Month
New Zealand Dollar	NZ 90 Day Treasury Bills Issue Rate
Norwegian Krone	BofAML NOK LIBOR 1 Mon CM
Singapore Dollar	BofAML SGD LIBOR 1 Mon CM
South African Rand	JIBAR 1 Month
South Korean Won	KBP CD
Swedish Krona	BofAML SEK LIBOR 1 Mon CM
Swiss Franc	BofAML CHF LIBOR 1 Mon CM
Thai Baht	Thailand 91 day T-bill
U.K. Pound Sterling	BofAML GBP LIBOR 1 Mon CM
U.S. Dollar	USTREAS T-Bill Auction Ave 3 Mon

Appendix 2: Methodology Changes

The following is a timeline of significant methodology changes to the Morningstar Rating. In addition to the changes noted below, Morningstar has occasionally added new categories to the category structure in each market, and those changes can also have an impact on the fund ratings.

Date	Description
Oct. 31, 2016	Removed load adjustment from star ratings calculations in the United States and Europe. Removed load-waived hypothetical share classes from the ratings. Merged U.S. exchange-traded fund and open-end funds as a single population for comparison. Removed similarity matrix for overall rating with category changes. Merged methodology with the Ibbotson Stars in Japan.
Dec. 31, 2013	Merged with the existing star rating methodology for Canada-domiciled funds.
Oct. 31, 2006	Europe: Introduced five-year and 10-year ratings and an overall rating that is based on a weighted average of the three-year, five-year, and 10-year ratings. Started to apply deferred loads and redemption fees to risk-adjusted return calculation.
Sept. 30, 2006	Released a new version of the methodology document that is more appropriate for a global audience. The U.S. calculations did not change. Also, the document was revised in order to offer more explanations on certain topics and to clarify the order of the calculations. Added rating suspension policy.
July 31, 2006	United States: Minor changes implemented. Removed the function that rounded variables n1-n5 (the rating breakpoints) to integers. Changed the logic for assigning ratings to look for all funds up to but not exceeding each breakpoint, instead of all funds reaching or just exceeding each breakpoint.
June 30, 2002	United States: Implemented significant enhancements to the rating, including category peer groups, fractional weights for multishare funds, category change adjustment, and more-robust risk-adjustment process (Morningstar Risk-Adjusted Return).
March 31, 2001	Europe: Introduced three-year (=overall) Morningstar Rating based on category peer groups.
1985	United States: Introduced Morningstar Rating (three-year, five-year, 10-year, overall) based on broad asset-class peer groups.