The Morningstar Rating™ Methodology

Morningstar Methodology Paper
July 26, 2007

Contents

Introduction 3

Morningstar Categories
- **Category Peer Groups** 5
- **Style Profiles** 5
- **Defining Fund Categories** 6

Theory
- **Expected Utility Theory** 7
- **Degree of Risk Aversion** 9
- **Morningstar's Formulation of Utility Theory** 10

Illustration of the Utility Function 13

Calculations
- **Overview** 18
- **Total Return** 19
- **Load-Adjusted Return** 21
- **Morningstar Return** 23
- **Morningstar Risk-Adjusted Return** 24
- **Morningstar Risk** 24

The Morningstar Rating: 3-, 5-, and 10-Year 25

Morningstar Return and Morningstar Risk Scores 27

The Overall Morningstar Rating
- **The Overall Morningstar Rating With No Category Changes** 28
- **The Overall Morningstar Rating With Category Changes** 29

Rating Suspensions 31

Frequently Asked Questions 32

Conclusion 34

Methodology Changes 35

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 domiciled or available for sale in Europe, Hong Kong, Singapore, Taiwan, and the United States.

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 through leverage or de-leverage before comparing them. Hence, a fund's score is not sensitive to its proportion of risk-free assets or its amount of leverage. 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.

Introduction (continued)

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 fund's excess returns are 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. Among other things, the rating is designed to convey a sense of how skillfully a fund has been managed. Its usefulness depends, in part, on which funds are compared to one another.

It can be assumed that the returns of major asset classes (domestic equities, foreign equities, domestic bonds, etc.) 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 (i.e., 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 investments that dominate their portfolios. Where holdings data are unavailable, prospectus or other information may be used to classify the fund.
- In general, a single return benchmark should form a valid basis for evaluating the returns for all funds in a single category (i.e., 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. [1]

[1] In contrast, the Morningstar Style Box™ is a snapshot based on the fund's most-recent portfolio. In some markets, the same names are used for both the Morningstar Category and the Morningstar Style Box (e.g., Large Blend). A fund's style box may differ from its category assignment when the style of its current holdings deviates from the fund's long-term investment strategy.

Theory

Expected Utility Theory

Morningstar Risk-Adjusted Return (MRAR) 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 value of a portfolio being considered and u(.) 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 $u'(.) > 0$.

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] $u(E[W]) > E[u(W)]$

From probability theory, it follows that this can be true only if u(.) is everywhere a concave function, so $u''(.) < 0$.

3) No particular distribution of excess returns is assumed.
Expected utility theory does not rely on any assumptions about whether a fund's returns are 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, i.e., that 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] $$RRA(W) = -\frac{Wu''(W)}{u'(W)}$$

By assuming that RRA is a constant value (i.e., 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] $$u(W) = \begin{cases} -\dfrac{W^{-\gamma}}{\gamma} & \gamma > -1, \gamma \neq 0 \\ \ln(W) & \gamma = 0 \end{cases}$$

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

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 then –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 (e.g., equal likelihood of –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 is expected to earn 2% on average, with an equal likelihood of –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 –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 –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. First, all returns are adjusted for the impact of sales loads. Second, 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 return on the risk-free asset (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 $1.

Hence:

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

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

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

where

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

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 give up to reduce risk. 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 to 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] $$u\left(1+ER^{CE}\left(\gamma\right)\right)=E\left[u\left(1+ER\right)\right]$$

Hence:

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

Morningstar defines Morningstar Risk-Adjusted Return, MRAR(γ), 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] $$MRAR(\gamma)=\left[\frac{1}{T}\sum_{t=1}^{T}\left(1+ER_t\right)^{-\gamma}\right]^{-\frac{12}{\gamma}}-1$$

where

ER_t	=	the geometric excess return in month t	$=\frac{1+LR_t}{1+RF_t}-1$
LR_t	=	load-adjusted 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] $$\mathrm{MRAR}(0)=\left[\prod_{t=1}^{T}\left(1+\mathrm{ER}_t\right)\right]^{\frac{12}{T}}-1$$

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 fund 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.

Illustration of the Utility Function

Morningstar uses a utility function to model how investors trade off return and risk. It follows the form:

[12] $$U = z + w(1+ER)^{-\gamma}$$

where

U	=	the model investor's utility or satisfaction from each monthly return
z	=	any number
w	=	any negative number
ER	=	the monthly geometric excess return for the fund (adjusted for loads and the risk-free rate)
γ	=	a parameter that expresses an investor's sensitivity to risk

As stated earlier, Morningstar sets $\gamma = 2$ to illustrate a model investor's sensitivity to risk. With $z = 2$ and $w = -1$, the utility function looks like this:

[13] $$U = 2 - (1+ER)^{-2}$$

Illustration (continued)

Investors have a certain level of utility or satisfaction for each level of monthly return. The utility function can be graphed like this:

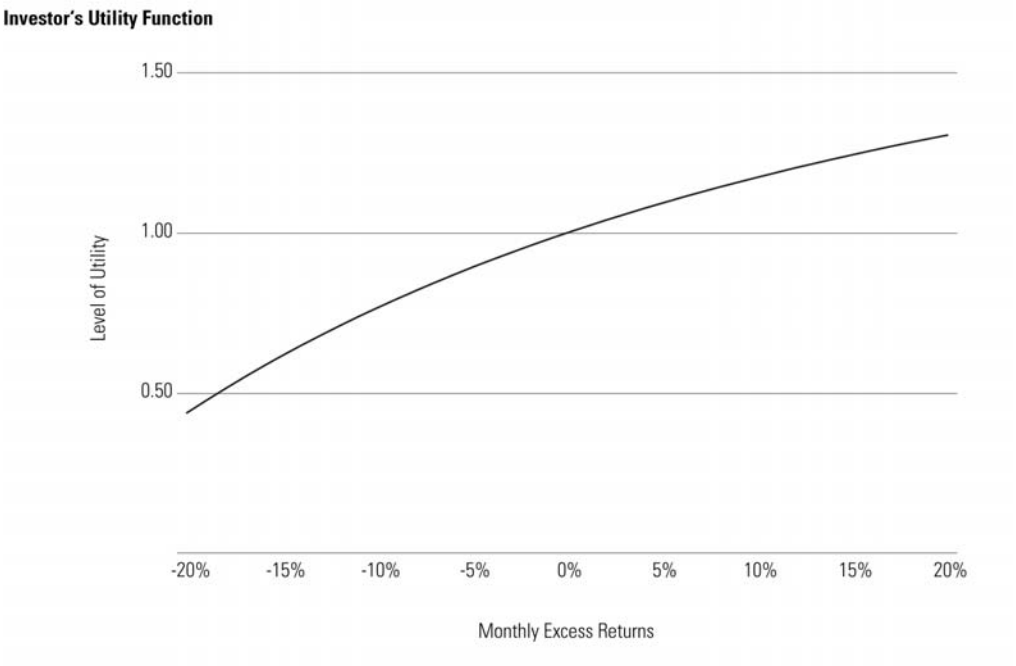

The model investor's utility function is concave, and there is decreasing marginal utility as returns increase. It is steeper for negative returns and starts flattening out for positive returns, and this puts more emphasis on downside variation.

Investors don't like to lose money and they experience more dissatisfaction with negative returns than satisfaction with positive returns. Moving from –20% to 0% produces a gain in utility of about 0.60, but moving from 0% to 20% only produces a gain in utility of about 0.30 in this example. When Morningstar averages the level of utility across many months for a fund, lower returns have much lower levels of utility or satisfaction and this brings down the average.

Illustration (continued)

If Morningstar used a higher value for γ to indicate a higher sensitivity to risk, the utility function would be even more concave, thereby giving greater penalties for downside variation. The dotted line below represents the utility function with $\gamma=3$.

Investor's Utility Function

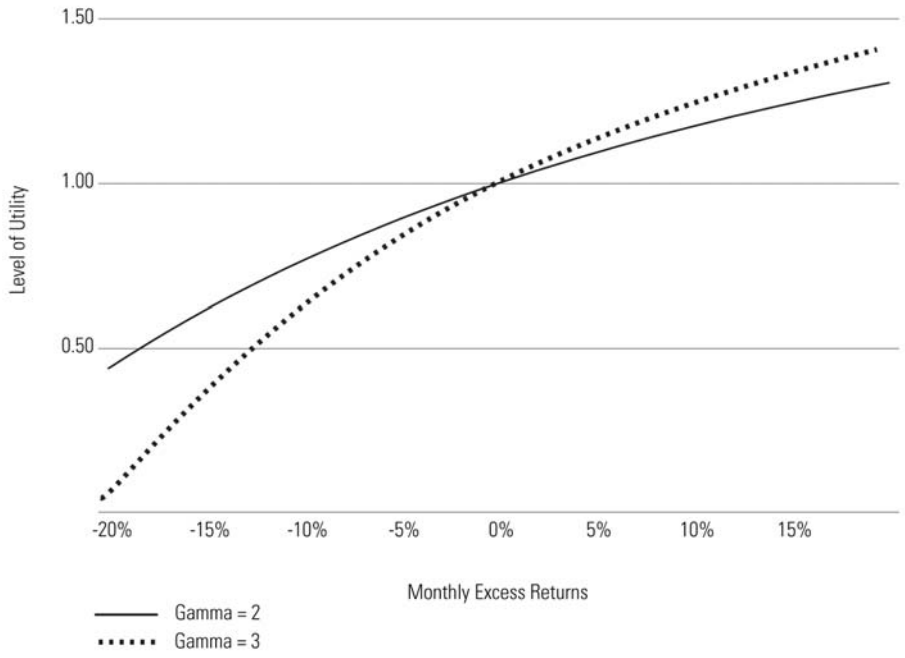

Illustration (continued)

Morningstar plots the fund's monthly geometric excess returns on this utility curve and then determines the average level of utility provided by this particular combination of returns. Morningstar then determines the level of monthly return (with no risk) that provides that same level of utility as the risky fund. This is called the certainty equivalent geometric excess return, which is the monthly expression of Morningstar Risk-Adjusted Return.

The average level of utility is converted into a return with the following equations, which are just versions of equations [12] and [13]. Because these equations subtract z and divide by w (the variables we previously added and multiplied), any values can be chosen for z and any negative values can be chosen for w.

[14] $$ER = \left[\frac{1}{w}(U - z)\right]^{-\frac{1}{\gamma}} - 1$$

[15] $$ER = [-(U - 2)]^{-\frac{1}{2}} - 1$$

Investor's Utility Function

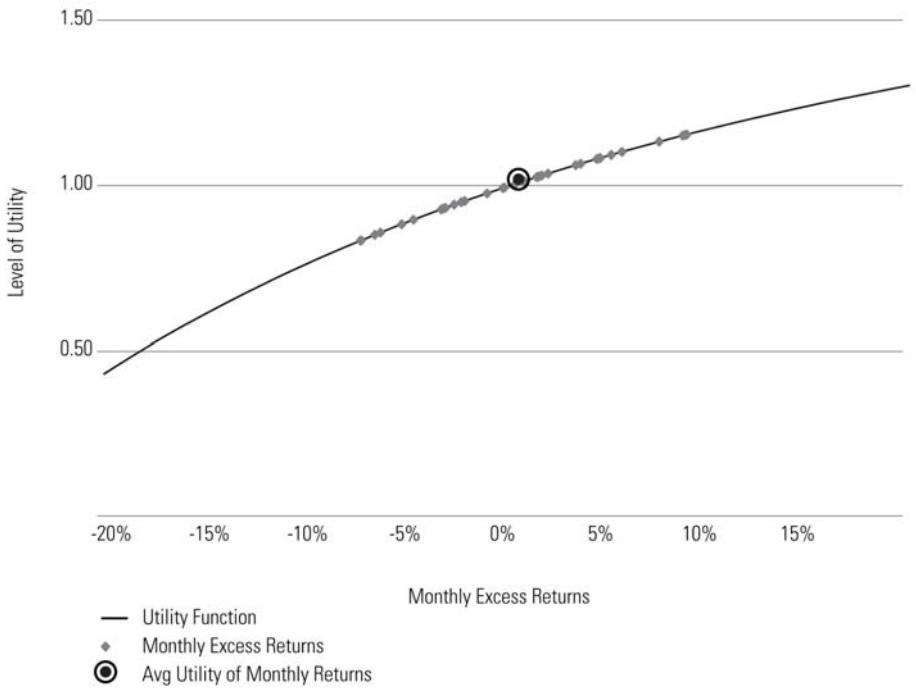

Illustration (continued)

The following graph illustrates the monthly excess returns of the sample fund from the previous page. The arithmetic average of these monthly returns is 0.92% and the geometric average monthly return is 0.81% per month. The average level of utility for this fund's 36 monthly excess returns is 1.012, and the certainty equivalent geometric excess return is 0.60% per month. The certainty equivalent return is less than the arithmetic or geometric averages because we set $\gamma > 0$. This investor is demanding a risk premium of 0.32% per month to invest in the risky portfolio.

Volatile Monthly Returns and Certainty Equivalent

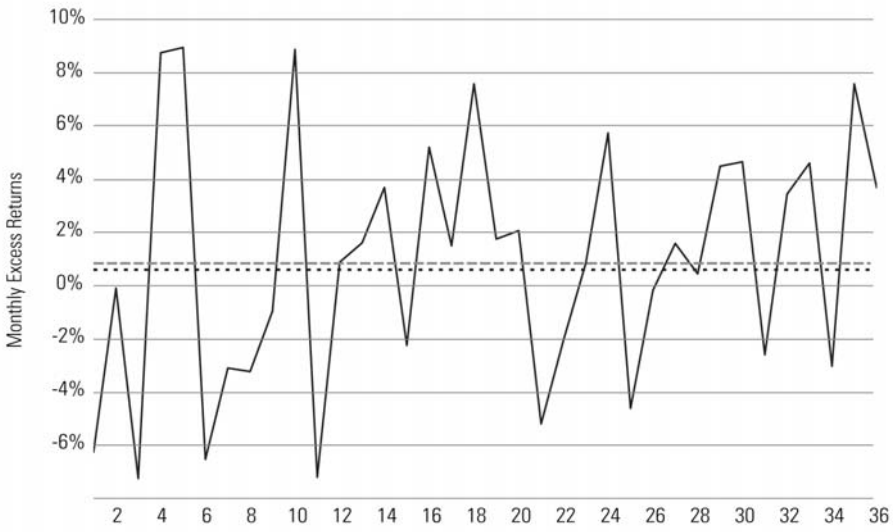

			Monthly	Annual
——	Monthly Excess Returns			
– – – –	Average Excess Return	Morningstar Return	0.81%	10.20%
·····	Certainty Equivalent	Morningstar Risk-Adjusted Return	0.60%	7.40%

The monthly 0.81% geometric excess return is annualized as a Morningstar Return of 10.20%. The 0.60% certainty equivalent is annualized as a Morningstar Risk-Adjusted Return of 7.40%. Morningstar Risk is the difference between Morningstar Return and MRAR, or 2.80% in this case. Morningstar Risk represents how much return the model investor is willing to give up to eliminate risk in the portfolio.

Calculations

Overview

There are four steps to calculate Morningstar Risk-Adjusted Return (MRAR). 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) Load-Adjusted Return: Adjust monthly total returns for sales loads.[2]
3) Morningstar Return: Calculate or collect monthly total returns for the appropriate risk-free rate. Adjust load-adjusted returns for the risk-free rate to get Morningstar Return.
4) 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.

		Adjust for loads ↓	Adjust for risk-free rate ↓	Adjust for risk ↓	
	Total Return 3 Yr	**Load-Adj Return 3 Yr**	**Morningstar Return 3 Yr**	**Morningstar Risk-Adj Return 3 Yr**	**Morningstar Risk 3 Yr**
Focus Fund A	34.43	31.80	29.40	26.37	3.03
Growth Fund Inv	23.30	23.30	21.06	17.11	3.95

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

	Total Return % Rank 3 Yr	**Load-Adj Return % Rank 3 Yr**	**Morningstar Return % Rank 3 Yr**	**Morningstar Risk-Adj Return % Rank 3 Yr**	**Morningstar Rating**
Focus Fund A	10	15	15	15	4
Growth Fund Inv	27	27	29	35	3
Equity Fund	92	93	93	93	1

For example, the load adjustment impacted the Focus fund's rating the most, moving it from the 10th to the 15th percentile. The Growth fund's rating was most affected by the adjustment for risk, which moved it from the 29th to the 35th percentile. The Equity fund's low returns (percentile rank 92) led to its low rating.

[2] In some markets, Morningstar sets a cap on loads in the calculation of MRAR but uses the maximum published load in the calculation of the standard load-adjusted return. In these markets, the above comparisons of annualized returns and percentile ranks may not clearly illustrate the size of the adjustments for loads and the risk-free rate. See page 21 for more information.

Total Return

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

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

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] $$TR_c = \prod_{t=1}^{T} (1 + TR_t) - 1$$

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 (e.g., 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 pre-tax 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] $$TDiv_i = \frac{Div_i}{(1 - t_{Si})(1 - t_{Fi})}$$

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 D_i in equation [16] for U.S. municipal bond funds. This tax adjustment is not part of the standard depictions of total return for these funds.

Load-Adjusted Return

Next, Morningstar adjusts the fund's total return for any front loads, deferred loads, or redemption fees that were in our records during the month-end under consideration.

- Morningstar uses the maximum front load for the fund and the deferred loads and redemption fees appropriate for the time period being measured (three-, five-, or 10-year).
- Morningstar uses the higher deferred load if two straddle a time period. For example, if the fund charges a 4% deferred load to investors that have held the fund for 2-3 years and 3% for 3-4 years, Morningstar will use the 4% load for the three-year rating.
- Per industry standards, the deferred load is applied to the lesser of the initial investment or the ending amount.
- The fund's historical load structure is not considered for its current rating. Morningstar does not prorate the loads based on the historical months when the load was actually in effect.
- Morningstar does not account for temporary load waivers. The fund must permanently remove the load before Morningstar will remove the load from that fund's rating.

In markets where the published sales load does not reflect the experience of any investors and instead represents a legal maximum that would only be charged under unusual circumstances, Morningstar may cap loads to a more realistic level that better represents the experience of the retail investor. The caps are applied only when calculating the load-adjusted returns that are used in MRAR; the maximum published load is still used for calculating the standard load-adjusted return. In Europe, Morningstar caps front and deferred loads, using a 5% maximum load for equity funds and a 3% maximum load for bond funds.

A fund's load-adjusted return is calculated as follows:

[19] $$LR_c = (1+TR_c)(1-F)(1-R) - D(1-F)\frac{\min(P_0,P_T)}{P_0} - 1$$

where

LR_c	=	cumulative load-adjusted return for three, five, or ten years
TR_c	=	cumulative total return for three, five, or ten years
F	=	maximum front load, expressed as a decimal
R	=	redemption fee for the time period being measured, expressed as a decimal
D	=	deferred load for the time period being measured, expressed as a decimal
P_0	=	beginning-of-period NAV per share
P_T	=	end-of-period NAV per share

Morningstar calculates a monthly adjustment factor for loads, *a*, as follows:

[20] $$a = \left(\frac{1 + LR_c}{1 + TR_c} \right)^{\frac{1}{T}}$$

where

T	=	number of months in the period (e.g., three, five, or 10 years)

Then, Morningstar adjusts the monthly total returns with this factor:

[21] $$LR_t = a\left(1 + TR_t\right) - 1$$

where

LR_t	=	the load-adjusted total return for month t
TR_t	=	the total return for month t

The adjustment factor and the load-adjusted monthly returns are different for the three-, five-, and ten-year periods for T=36, 60, and 120, respectively, because the loads are spread over different amounts of time.

Morningstar Return

Next, Morningstar adjusts the fund's monthly load-adjusted 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.

[22] $$ER_t = \frac{1+LR_t}{1+RF_t} - 1$$

where

ER_t	=	the geometric excess return for the fund for month t
LR_t	=	the load-adjusted 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 where the investment is available for sale and the primary currency of the investment, rather than where the fund invests.

- United States: 3 month Treasury Bill
- Europe: Different categories are paired with different risk-free rates, including the constant-maturity versions of the Merrill Lynch 1-month LIBOR index with currencies CHF, DKK, EUR, GBP, JPY, NOK, and SEK and the 3 month Treasury Bill.

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

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

where

T	=	number of months in the period (e.g., 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 (MRAR) 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 two, so Morningstar Risk-Adjusted Return is calculated as follows:

[24] $$\mathrm{MRAR}(2) = \left[\frac{1}{T} \sum_{t=1}^{T} \left(1 + \mathrm{ER}_t\right)^{-2} \right]^{-\frac{12}{2}} - 1$$

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 loads and the risk-free rate) and MRAR (adjusted for loads, the risk-free rate, and risk). Morningstar Risk is always greater than or equal to zero.

[25] Morningstar Risk = Morningstar Return – Morningstar Risk-Adjusted Return

The Morningstar Rating: 3-, 5-, and 10-Year

The Morningstar Rating is based on Morningstar Risk-Adjusted Return. 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.

The following items are needed to calculate the Morningstar Rating for funds:

- Morningstar Risk-Adjusted Return for all funds in a category for three, five, and 10 years
- An identifier or method to indicate which funds are share classes in the same portfolio
- Historical category assignments for all funds in the sample
- A table that describes the degree of similarity between any two pairs of categories, ranging from zero (highly dissimilar) to one (identical)

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 loads, 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 multi-share funds from taking up a disproportionate amount of space in any one rating level.

Let:

n	=	the total number of distinct portfolios being evaluated
n_1	=	the number of distinct portfolios that receive one star
n_2	=	the number of distinct portfolios that receive two stars
n_3	=	the number of distinct portfolios that receive three stars
n_4	=	the number of distinct portfolios that receive four stars
n_5	=	the number of distinct portfolios that receive five stars

Morningstar sets n_1 through n_5 so that the distribution of funds across the rating levels forms an approximate bell curve:

n_1	=	10% of n
n_2	=	22.5% of n
n_3	=	35% of n
n_4	=	22.5% of n
n_5	=	10% of n

Morningstar assigns fractional weights to those funds that are merely different share classes of the same underlying portfolio. For example, if the same portfolio can be bought in the form of five share classes and all five have performance for the time period being measured, each share class is counted as one-fifth of a fund. This allows Morningstar to rate share classes on a framework that is based on the number of portfolios.

If a category has 60 portfolios with data available for a particular time period, Morningstar will assign a five-star rating to *the equivalent of* six portfolios (10% of the sample). This is rarely six single-share funds. Instead, it is usually a combination of single-share funds and multi-share classes (that have the best risk-adjusted performance) whose fractional weights add up to the equivalent of six distinct portfolios.

Three-year star ratings are assigned as follows:

1) Assign a fractional weight to each fund, based on the number of other share classes in the same portfolio that have at least three years of data available.
2) Sort all funds in the category by three-year MRAR in descending order.
3) Starting with the highest MRAR, add up the funds' fractional weights until n_5 is reached but not exceeded. These funds receive five stars.
4) Continue counting off funds until the total number reaches but does not exceed n_5+n_4. The additional funds receive four stars.
5) Continue counting off funds until the total number reaches but does not exceed $n_5+n_4+n_3$. The additional funds receive three stars.
6) Continue counting off funds until the total number reaches but does not exceed $n_5+n_4+n_3+n_2$. The additional funds receive two stars.
7) The remaining funds receive one 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 Scores

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 one to five, and these scores are typically expressed as word labels in Morningstar products.

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 Score are helpful when funds have the same rating and similar MRARs but different levels of risk. All of the funds below got three stars and "Average" return, but they took very different levels of risk to achieve that rating.

Name	Morningstar Rating 3 Yr	Morningstar Risk-Adj Return 3 Yr	Morningstar Return 3 Yr	Morningstar Return Score 3 Yr	Morningstar Risk 3 Yr	Morningstar Risk Score 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

The Overall Morningstar Rating With No Category Changes

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. If the fund in question has been in its current category over the entire evaluation period, the weights are:

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 three stars.

Period	Rating	Weight	Multiply
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.

The Overall Morningstar Rating With Category Changes

If the fund's investment style and category have changed in the past, Morningstar modifies the weights for the Overall Rating to reflect the average degree of similarity between the current category and the fund's historical categories. Slightly more weight is placed on the time periods when the fund was in the current category.

Morningstar has identified the degree of similarity between any two pairs of categories, within each local market. The degree of similarity can range from zero (highly dissimilar) to one (identical). This data is available upon request from Morningstar's local offices. Here are some examples of the degrees of similarity between category pairs, using the U.S. category system.

Category 1	Category 2	Degree of Similarity
Large Blend	Large Value	0.50
Mid-Cap Growth	Small Blend	0.25
Large Growth	Moderate Allocation	0.25
Specialty Technology	Specialty Communications	0.25

Generally, categories that are adjacent on the Morningstar Style Box (e.g., Large Blend with Large Value, Large Growth, and Mid-Cap Blend) have a degree of similarity of 0.50, while categories that are diagonal (Large Blend to Mid-Cap Value and Mid-Cap Growth) have a degree of similarity of 0.25. For many category pairs, the degree of similarity is zero.

In the following formulas, s denotes the number of months back in time with s=1 meaning the current month, s=2 meaning the previous month, etc. (If a category change happened mid-month, the month-end assignment is used for the entire month.)

Let:

Ds	=	the degree of similarity between the fund's category in month 1 and the fund's category in month s.

The average degree of similarity for the three-year period is:

[26] $$\overline{D}_3 = \frac{\sum_{s=1}^{36} D_s}{36}$$

The average degree of similarity for the five-year period is:

[27] $$\overline{D}_5 = \frac{\sum_{s=1}^{60} D_s}{60}$$

The average degree of similarity for the 10-year period is:

[28] $$\overline{D}_{10} = \frac{\sum_{s=1}^{120} D_s}{120}$$

When there are five years of data available, the three- and five-year ratings are combined with the following weights:

[29] $$W_5 = \frac{0.60\overline{D}_5}{0.40\overline{D}_3 + 0.60\overline{D}_5}$$

[30] $$W_3 = \frac{0.40\overline{D}_3}{0.40\overline{D}_3 + 0.60\overline{D}_5}$$

When there are 10 years of data available, the three-, five-, and 10-year ratings are combined with the following weights:

[31] $$W_{10} = \frac{0.50\overline{D}_{10}}{0.20\overline{D}_3 + 0.30\overline{D}_5 + 0.50\overline{D}_{10}}$$

[32] $$W_5 = \frac{0.30\overline{D}_5}{0.20\overline{D}_3 + 0.30\overline{D}_5 + 0.50\overline{D}_{10}}$$

[33] $$W_3 = \frac{0.20\overline{D}_3}{0.20\overline{D}_3 + 0.30\overline{D}_5 + 0.50\overline{D}_{10}}$$

Rating Suspensions

The degree of similarity adjustment in the Overall Rating is sufficient for most historical category changes. 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 to 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.)

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 editorial 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 five 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.

Why did a fund's rating change from one month to the next?

The following are some of the most common reasons why a fund's rating may change:

- The fund recently changed its Morningstar category and is now being compared to a new peer group.
- The fund recently hit its five- or 10-year anniversary, so more time periods are being included in the overall rating.
- The fund is only three years old so volatility in the three-year rating produces volatility in the overall rating, without any longer periods to stabilize the rating.
- The fund's risk-adjusted return was very close to a breakpoint between rating levels and slipped slightly over the breakpoint from one month to the next.[3]
- The fund just rolled off a period of very strong performance three, five, or 10 years ago.

How can I predict my fund's rating?

There is no way to predict a fund's future rating with absolute certainty, but there are ways to evaluate how a fund ranks on a risk-adjusted basis against its peers.

- Calculate a custom-period MRAR: For example, calculate a 35-month MRAR for a fund that is one month shy of its three-year anniversary and compare that to the 35-month MRARs of all other funds in that category. Percentile ranking the population will provide a rough estimate of how the subject fund ranks. To be even more precise, provide fractional weights to multi-share funds to determine how the ratings would be distributed.
- Estimate returns: Fill in monthly returns prior to a fund's inception in order to obtain 36, 60, or 120 months of performance. For example, if a fund is 35 months old, estimate the fund's monthly return for the month prior to the fund's inception. Then, calculate a three-year MRAR for the fund and compare it to the three-year MRARs for the other funds in the category to determine how it would have been rated. These estimated monthly returns can be category average monthly returns or other estimates based on similar funds or older share classes.

[3] This can be investigated with Morningstar Direct™, Morningstar's institutional research platform, by pulling in all funds in a category and sorting by Morningstar Risk-Adjusted Return for the period in question.

Why did Fund 1 get a different rating than Fund 2?

- Check to see if the two funds are in the same category. The ratings are based on category peer groups and ratings between funds in two different categories will likely be different.
- If the two funds under consideration are share classes of the same portfolio, the different ratings are likely due to different load structures or different inception dates. If the inception dates are different, different time periods may be contributing to the Overall Ratings of these funds.
- The funds may have different loads or different levels of risk. Study the percentile ranks for the two funds for total return, load-adjusted return, Morningstar Return, and Morningstar Risk-Adjusted return to see why they got rated differently and which factors impacted the ratings the most (see page 18).

How does the fund's expense ratio affect the Morningstar Rating?

The expense ratio will impact the rating only to the extent that it is used in the calculation of net asset values and total return. Total return forms the basis for Morningstar Risk-Adjusted Return, and there is no further explicit adjustment for a fund's expense ratio.

Why didn't my fund get a rating?

There are a few reasons why a fund would not get a rating.

- The fund is not three years old yet.
- The fund's rating has been suspended.
- There are a few "convenience" categories that are not rated, because the group contains different funds with different risk factors. For example, Morningstar does not rate funds in the U.S. bear-market category or the Europe guaranteed category.
- If the category is small, the sample may not be robust enough to calculate a rating. Morningstar requires at least five distinct portfolios in a category for a specific time period in order to calculate that rating.

Why is Morningstar Risk so high for my fund?

Morningstar Risk tends to be high when the fund has experienced a lot of variability in its monthly returns and/or when it has some instances of strongly negative monthly returns.

While Morningstar Risk and standard deviation are not the same, there is often a high correlation between the two measures. Standard deviation measures variation both above and below the mean equally, but Morningstar Risk gives more weight to downside variation. If the fund has had both upside and downside variation, its percentile rank in category for standard deviation is usually very close to its percentile rank in category for Morningstar Risk.

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 sales loads, the risk-free rate, and risk.

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 impact the fund ratings.

Date	Description
October 31, 2006	Europe: Introduced 5y and 10y ratings and an Overall Rating that is based on a weighted average of the 3y, 5y, and 10y ratings. Started to apply deferred loads and redemption fees to risk-adjusted return calculation.
September 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 n_1-n_5 (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 multi-share 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 (3y, 5y, 10y, Overall) based on broad asset class peer groups.