
Morningstar Quantitative Rating™

Pillar Methodology

Morningstar Quantitative Research

23 January 2020

Version 2.1

Content

- 1 Introduction
- 2 Philosophy of the Ratings
- 3 Rating Descriptions
- 4 Methodology Overview
- 4 Pillar Rating Methodology
- 7 Final Rating Methodology
- 7 Model Accuracy
- 11 Conclusion
- 11 References

- Appendix A
- 12 Random Forest

- Appendix B
- 15 Quantitative Parent Pillar
- 16 Quantitative People Pillar
- 17 Quantitative Process Pillar

- Appendix C
- 18 Input Data FAQ

Introduction

Morningstar has been conducting independent investment research since 1984. Traditionally, our approach has been to provide analyst-driven, forward-looking, long-term insights to help investors better understand investments. Morningstar has one of the largest independent manager research teams in the world, with more than 100 analysts globally covering more than 3,700 unique funds.

The Morningstar Analyst Rating™ for funds (the Analyst Rating) provides a forward-looking evaluation of how these funds might behave in a variety of market environments to help investors choose superior funds. It's based on an analyst's conviction in a fund's ability to outperform its peer group and/or relevant benchmark on a risk-adjusted basis through a full market cycle of at least five years.

The number of funds that receive an Analyst Rating is limited by the size of the Morningstar analyst team. To expand the number of funds we cover, we have developed a machine-learning model that uses the decision-making processes of our analysts, their past ratings decisions, and the data used to support those decisions. The machine-learning model is then applied to the "uncovered" fund universe and creates the Morningstar Quantitative Rating™ for funds (the Quantitative Rating), which is analogous to the rating a Morningstar analyst might assign to the fund if an analyst covered the fund. These quantitative ratings predictions make up what we call the Morningstar Quantitative Rating. With this new quantitative approach, we can rate nearly 6 times more funds in the global market.

Only open-end funds and exchange-traded funds that don't currently have an Analyst Rating and are in a category that Morningstar currently rates are eligible to receive a quantitative rating. With the introduction of the Morningstar Quantitative Rating, we're extending a useful analytic tool to thousands of additional funds, providing investors with much greater breadth of coverage from the independent perspective they have come to know and trust from Morningstar.

Authors

Lee Davidson, CFA
Head of Quantitative Research
+1 312-244-7341
lee.davidson@morningstar.com

Madison Sargis
Associate Director of Quantitative Research
+1 312-244-7352
madison.sargis@morningstar.com

Timothy Strauts
Director of Quantitative Research
+1 312-384-3994
timothy.strauts@morningstar.com

Daniel Kleeman
Quantitative Analyst
+1 312 244 7336
daniel.kleeman@morningstar.com

Philosophy of Morningstar Quantitative Rating™ for Funds

Morningstar has been producing differentiated investment research since 1984. Although Morningstar research has expanded to equity, corporate credit, structured credit, and public policy, our roots are in the world of mutual funds. Traditionally, our approach has been to provide analyst-driven, forward-looking, long-term insights alongside quantitative metrics to further understanding of the investment landscape. We have developed a way to combine our analyst-driven insights with our robust fund data offering to expand fund analysis beyond the capabilities of our manager research staff. With this new development, we are able to cover 6 times more funds in the global market through empirical methods that are based on the proprietary ratings our analysts are already assigning to funds.

In general, there are two broad approaches that we could have chosen to expand our analyst-driven rating coverage in a quantitative way: Either automate the analyst thought process without regard for output similarity; or, replicate the analyst output as faithfully as possible without regard for the analyst thought process. Attempting to mechanically automate a thought process introduces tremendous complexity, so we opted to build a model that replicates the output of an analyst as faithfully as possible.

Replicating the Analyst Rating was a desirable goal because Morningstar has demonstrated throughout its history that the recommendations of its analysts provide value to investors. Therefore, at the outset, it seemed plausible that if a statistical model could be created that replicated the analysts' decision-making process, then there stood a decent chance it would produce valuable results as well. Indeed, based on live results since June 2017, this is exactly what we have found.

Perhaps the most obvious benefit to investors of the quantitative set of ratings is the breadth of coverage and frequency of update. Our quantitative coverage universe is many times the size of our analyst-covered universe, and growing. It is limited only by our access to the necessary input data. Additionally, the Morningstar Quantitative Rating has the unique advantage of maintaining a monthly update cycle. Each fund's rating is refreshed on a frequency unsustainable by a fund analyst.

Of course, no rating system — quantitative or analyst — is valuable without empirical evidence of its predictive ability. We have rigorously tested the performance, accuracy, and stability of the Quantitative Rating. This document includes numerous studies performed on the ratings, and we will continue to enhance our methodologies over time to improve performance.

Morningstar Quantitative Rating Descriptions

The Quantitative Ratings are composed of the Morningstar Quantitative Rating™ for funds, Quantitative Parent Pillar, Quantitative People Pillar, and Quantitative Process Pillar. A high-level description of each rating is found below. The statistical model is described in the Overview Methodology section on page 4. The pillar rating methodology begins on the same page.

- ▶ **Morningstar Quantitative Rating™ for funds:** Comparable to Morningstar's Analyst Ratings for open-end funds and ETFs, which are the summary expression of Morningstar's forward-looking analysis of a fund. The Analyst Rating is based on the analyst's conviction in the fund's ability to outperform its peer group and/or relevant benchmark on a risk-adjusted basis over a full market cycle of at least five years. Ratings are assigned on a five-tier scale with three positive ratings of Gold, Silver, and Bronze; a Neutral rating; and a Negative rating. Morningstar calculates the Quantitative Rating using a statistical model derived from the Analyst Rating our fund analysts assign to open-end funds.

- ▶ **Quantitative Parent Pillar:** Comparable to Morningstar's Parent Pillar ratings, which provide Morningstar's analyst opinion on the stewardship quality of a firm. Morningstar calculates the Quantitative Parent Pillar using an algorithm designed to predict the Parent Pillar rating our fund analysts would assign to the fund. The Quantitative Rating is expressed as High, Above Average, Average, Below Average, or Low.

- ▶ **Quantitative People Pillar:** Comparable to Morningstar's People Pillar ratings, which provide Morningstar's analyst opinion on the fund manager's talent, tenure, and resources. Morningstar calculates the Quantitative People Pillar using an algorithm designed to predict the People Pillar rating our fund analysts would assign to the fund. The Quantitative Rating is expressed as High, Above Average, Average, Below Average, or Low.

- ▶ **Quantitative Process Pillar:** Comparable to Morningstar's Process Pillar ratings, which provide Morningstar's analyst opinion on the fund's strategy and whether the management has a competitive advantage enabling it to execute the process and consistently over time. Morningstar calculates the Quantitative Process Pillar using an algorithm designed to predict the Process Pillar rating our fund analysts would assign to the fund. The Quantitative Rating is expressed as High, Above Average, Average, Below Average, or Low.

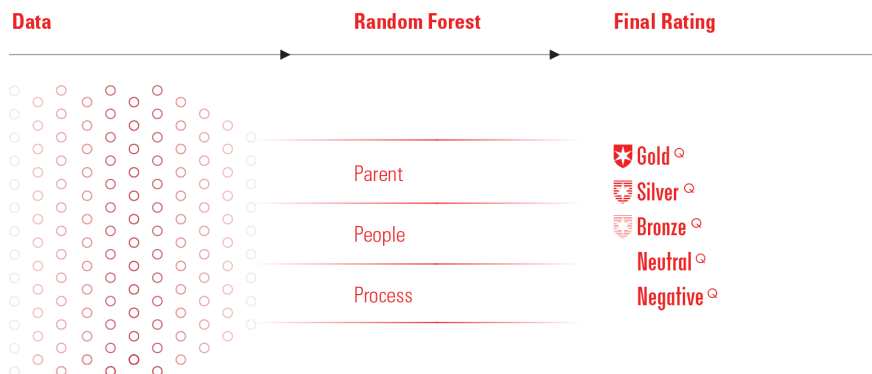
Overview of the Quantitative Rating Methodology

The Quantitative Rating consists of a series of seven individual models working in unison that were designed to provide a best approximation for the Analyst Rating on the global universe of open-end funds and ETFs. Visually, you can think of the estimation as a two-step process. First, we estimate the pillar ratings for each fund, and then we estimate the overall rating.

To estimate the pillar ratings, we chose a machine-learning algorithm known as a "random forest" to fit a relationship between the fund's pillar ratings and its attributes. For each pillar, two random forest models were estimated that seek to determine the probability that fund will be rated Positive or Negative, respectively. Since there are three pillars, we estimated six individual random forest models to answer these questions and produce six probabilities (two per pillar). Then, at the pillar level, we aggregate these probabilities to produce one overall pillar rating.

After the pillar ratings are estimated, we needed to aggregate them into an overall fund rating. In order to do this, we apply the analyst ratings framework. The final result is the Morningstar Quantitative Rating™ for funds. For more information on the Analyst ratings, see References section.

Exhibit 1 Representation of a Morningstar Quantitative Rating Methodology



Source: Morningstar, Inc.

Morningstar Quantitative Rating — Pillar Rating Methodology

The pillar ratings represent the foundation of the Analyst Rating. For the Quantitative Rating, the pillar ratings were estimated using a series of random forest models and rated on a 1 to 5 scale, and labeled as High, Above Average, Average, Below Average, or Low.

In order to estimate the pillar ratings, data was collected for the funds that analysts have currently assigned pillar ratings. In total, 180-plus attributes and 10,000-plus rating updates were considered in

order to train the random forest model. After numerous iterations, only the attributes most crucial to classifying each pillar rating were retained.

Each pillar rating is estimated using a combination of two random forest models. First, a model is estimated that seeks to distinguish funds based on whether that fund's pillar rating would be rated Positive, defined as High or Above Average. Second, a different model is estimated that seeks to distinguish funds based on whether that fund's pillar rating would be rated Negative, defined as Low or Below Average. Each model puts out probability scores that the fund would be Positive or Negative. By combining these two probabilities via a weighted summation, a more robust estimator is achieved.

$$\textit{Estimated Pillar Rating} = \frac{\textit{Prob(Positive)} + [1 - \textit{Prob(Negative)}]}{2}$$

The output for these pillar ratings will, therefore, be on a scale of 0 to 1. The closer to 1 a fund's estimated pillar rating is, the more likely it is that the true pillar rating is High. Similarly, the closer to 0 a fund's estimated pillar rating is, the more likely that the true pillar rating is Low. After the ratings were computed, thresholds were assigned that tended to correspond to natural distinctions between the five rating options for each pillar.

The intuition underlying this method is subtle, yet important. First, the weighted summation captures information about a fund along two dimensions—the likelihood that a fund's pillar is High and the likelihood that a fund is not Low. In practice, this has the result of classifying many Average pillars as decidedly not High and not Low.

Furthermore, by using two models to estimate a pillar rating, we are able to distinguish between data points that are important to each model individually. It makes intuitive sense that the data points that might indicate to an analyst to rate a fund High could be different from those that are used to rate a fund Low. By adding in that flexibility, we dramatically improved our estimation. Empirically, several pillar models exhibited significant overlap in data points used to estimate each model, but that did not always hold.

Smoothing Algorithm

After raw pillar ratings have been computed, we implement a smoothing algorithm to reduce inter-month volatility. This algorithm takes the average of the current raw pillar rating and the two prior months' raw pillar ratings to create a three-month moving average. The three-month moving average was chosen to balance the desire to reduce unnecessary volatility of ratings month to month, but also allow the ratings to be adaptable to significant changes at the fund, such as a manager change.

People and Process Pillar Business Logic

After smoothing, we implement a business rule to ensure that People and Process Pillar ratings do not change depending on the share class, and People scores do not vary within portfolio management teams. Technically, each fund share class will have their own People and Process Pillar ratings produced by the model, but we want to ensure that these are consistent for the same fund. To ensure this, we

implement an asset-weighted average of raw People and Process Pillar ratings across share classes with the weights determined by share-class-level net assets. In the case where net assets are not available, share-class-level ratings will be equally weighted. To ensure the People Pillar rating is applied consistently to a team, we create manager-level scores by averaging the People Pillar ratings of the funds they manage. We then roll back up People Pillar ratings for funds by averaging the manager scores, weighted by tenure. For funds who do not report the manager names, this logic is not applied. Instead, for actively managed funds without a reported manager, we restrict their ability to receive a High or Low People Pillar rating.

The final raw Pillar ratings, after smoothing, asset-weighting, and adjusting for teams, are saved as the pillar rating estimate for the current month for each fund share class.

In the case where an analyst has rated a fund belonging to the same strategy, all other funds under that same strategy identifier will inherit the People and Process Pillar rating assignments as determined by the analyst. This ensures that the analyst view is leveraged whenever available to ensure consistency between the Analyst Rating and Quantitative Rating systems when it comes to the People and Process Pillars.

For index products, the analyst team assigns the same People Pillar to all products linked to a firm. To mimic, the quantitative system will assign the same People score to all index products within an asset class to a firm. We decided to further filter by asset class to leave room for some variation in subject matter expertise. Similarly, the analyst fund assigns the same Process Pillar to all index products tracking the same benchmark. The quantitative system applies the same logic by averaging all raw Process Pillar ratings tied to a primary prospectus benchmark.

Parent Pillar Business Logic

In the same spirit, we implement one final business rule. In the case where there is an Analyst Rating for the Parent Pillar of a fund for a particular branding entity, we will suppress the Quantitative Parent Pillar for all funds from that particular branding entity and default to the analyst opinion. In this way, we ensure consistency of opinion between analyst and quant rating systems when it comes to the Parent Pillar.

Pillar Threshold

For those pillars where an analyst rating is not available, pillar labels (High, Above Average, Average, Below Average, or Low) will be assigned according to a static threshold to the raw pillar ratings using a symmetric distribution of 10%, 22.5%, 35%, 22.5%, and 10%.

- ▶ If raw pillar rating ≤ 0.10 , then 1 — Low
- ▶ If $0.10 < \text{raw pillar rating} \leq 0.325$, then 2 — Below Average
- ▶ If $0.325 < \text{raw pillar rating} \leq 0.675$, then 3 — Average
- ▶ If $0.675 < \text{raw pillar rating} \leq 0.90$, then 4 — Above Average
- ▶ If raw pillar rating > 0.90 , then 5 — High

Calculating the Quantitative Rating

The final step in the Quantitative Rating applies the Analyst Rating framework to assign ratings. Once the pillar scores have been determined, a fund receiving a Morningstar Quantitative Rating uses the same rating calculation as the Morningstar Analyst Rating.

We provide a short description of their ratings logic below. For the full detailed methodology, the Morningstar Analyst Rating methodology document is found in the References section.

Funds are first organized into investment supergroups, defined by the global set of funds that invest in a common pool of securities. We calculate the distribution of performance for each investment group, including statistics for each distribution. To ascertain the expected forward-looking performance for a fund, we calculate the weighted sum of its pillar scores, scale it by the width of the distribution it falls under, and subtract its fee. To achieve a rating, funds are sorted by their expected performance into rating groups defined by their Morningstar Category and their active or passive status. Each rating group is further split into a Morningstar Medalist-eligible and Medalist-ineligible cohort based on whether we predict each fund will outperform its benchmark and category average. The top 15% of Medalist-eligible share classes in a rating group are given a Gold rating, the next 35% Silver, and the bottom 50% a Bronze rating. The top 70% of Medalist-ineligible funds are given a Neutral rating and the bottom 30% a Negative rating.

Rating Stability

To increase the rating stability for funds near the breakpoints, we implement a buffering system. Between Negative/Neutral and Neutral/Bronze, the buffer is 2%. Between Bronze/Silver and Silver/Gold, the buffer is 1%. A fund near the rating thresholds must move past the buffer before the rating changes. For example, a fund below the 15.0 percentile will need to move to the 17.0 percentile before the rating upgrades from Negative to Neutral. Similarly, a fund above the 15.0 percentile will need to move below the 13.0 percentile before being downgraded from Neutral to Negative.

Model Accuracy

The Morningstar Quantitative Rating model is constructed to mimic the rating assignment behavior of our manager research staff. While we believe that forecasting out-of-sample future performance is the most important aspect for investors, we have tested the accuracy of Quantitative Pillar Ratings in their ability to match the Analyst Pillar Ratings.

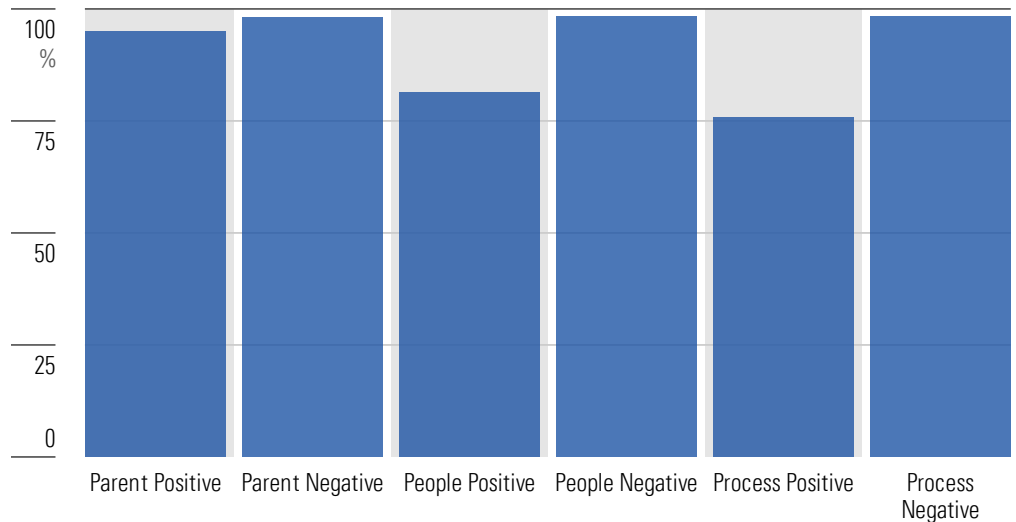
To evaluate the accuracy, we look at both the Positive and Negative directions. In the case of Negatives, we look at how well we can match an analyst assigning a Below Average or Low pillar score. Our accuracy here is near perfect, consistently above 95.0% for all three pillars. This makes sense as the attributes of poorly analyst-rated funds are obvious: high fees, manager turnover, and bad performance.

In the case of Positives, we look at how well we can match an analyst assigning an Above Average or High pillar score. Our accuracy varies by pillar. With an accuracy score of 95.0%, it is unsurprising that Parent is the best-performing of the Positive pillars. Input data points about firms are the most widely available in the model.

Out of all the tests, the Positive People and Process ratings have the lowest accuracy, at 81.4% and 75.7%, respectively. However, there are very few instances of large disparities between the Analyst Pillar Ratings and the Quantitative Pillar Ratings. When the two systems disagree, it is often because the Quantitative system skews to Average.

Overall, we are happy with the precision of the Quantitative Rating as we balance the desires to increase accuracy, avoid overfitting, and achieve strong future performance.

Exhibit 2 Percentage Accuracy Between Quantitative Pillar Ratings and the Analyst Pillar Ratings

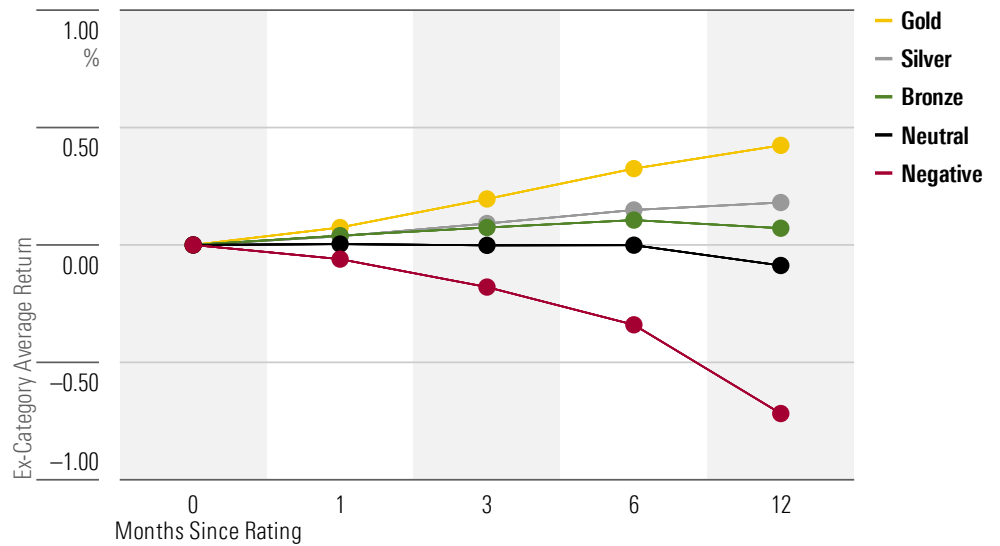


Source: Morningstar, Inc. Data as of March 31, 2019.

Performance

Since the June 2017 launch of the Morningstar Quantitative Ratings, we have closely monitored their performance. The system continues to pick clearly over- and underperforming funds. In Exhibit 3, we present an event study of the ex-category average returns of the rating calls since launch. When we adjust for style and risk, we see a pronounced separation between funds. Over a 12-month period, we find Gold-rated funds outperform their category average by 0.42%. The biggest performance gap is between the Negative and Neutral funds. Neutral-rated funds perform at the category average while Negative-rated funds underperform their peers by 0.72%.

Exhibit 3 Morningstar Quantitative Rating™ for Funds Event Study Since Launch



Source: Morningstar, Inc. Data as of March 31, 2019.

Stability

Finally, we see that the Quantitative Ratings are quite stable through time. Below we show the transition figures since the June 2017 launch. When a fund receives a Negative rating, we would expect that it has only a 0.24% probability of receiving a Bronze rating a year later and a 0.04% probability of receiving a Gold or Silver rating. Similarly, when a fund receives a Gold rating, we would expect that the fund has only a 2.41% probability of receiving a Neutral rating a year later and a 0.01% probability of receiving a Negative rating. In other words, we tend to stick to our guns when rating funds using the Quantitative Rating. We expect the modifications effective October 2019 to further improve the stability of the ratings system.

Exhibit 4 Morningstar Quantitative Rating™ for Funds Stability Transition Matrix: 1 Month

Time	Time + 1 Month					%
	Gold ^Q	Silver ^Q	Bronze ^Q	Neutral ^Q	Negative ^Q	
Gold ^Q	92.16	7.47	0.34	0.03	0.00	
Silver ^Q	3.70	86.61	9.31	0.38	0.00	
Bronze ^Q	0.06	6.42	84.89	8.64	0.00	
Neutral ^Q	0.00	0.08	2.35	96.05	1.52	
Negative ^Q	0.00	0.00	0.00	5.50	94.50	

Source: Morningstar, Inc. Data as of March 31, 2019.

Exhibit 5 Morningstar Quantitative Rating™ for Funds Stability Transition Matrix: 12 Months

Time	Time + 12 Months					%
	Gold ^Q	Silver ^Q	Bronze ^Q	Neutral ^Q	Negative ^Q	
Gold ^Q	57.19	28.11	12.28	2.41	0.01	
Silver ^Q	14.80	41.77	29.13	14.20	0.09	
Bronze ^Q	3.08	19.14	37.35	39.74	0.68	
Neutral ^Q	0.23	2.99	10.69	76.92	9.18	
Negative ^Q	0.00	0.04	0.24	33.12	66.60	

Source: Morningstar, Inc. Data as of March 31, 2019.

Conclusion

The Morningstar Quantitative Rating™ for funds is intended to be predictive of future alpha, and performance studies have affirmed that it is, in fact, performing as intended. For additional details, please refer to the Morningstar Quantitative Rating™ for funds FAQ document or feel free to contact us.

We expect that, over time, we will enhance the Quantitative Rating to improve performance. We will note methodological changes in this document as they are made.

References

Examining the Morningstar Quantitative Rating System. 2017.

<https://www.morningstar.com/content/dam/marketing/shared/research/foundational/878656-ExaminingMorningstarQuantRating.pdf>

Morningstar Analyst Rating for Funds Methodology Document. 2011.

https://www-prd.morningstar.com/content/dam/marketing/shared/research/methodology/778136_Morningstar_Analyst_Rating_for_Funds_Methodology.pdf

Morningstar Quantitative Equity Ratings Methodology. 2012.

<http://corporate.morningstar.com/US/documents/MethodologyDocuments/FactSheets/QuantitativeEquityCreditRatingsMeth.pdf>

Stekhoven, Daniel J., and Peter Bühlmann. "MissForest—non-parametric missing value imputation for mixed-type data." *Bioinformatics* 28.1 (2011): 112-118.

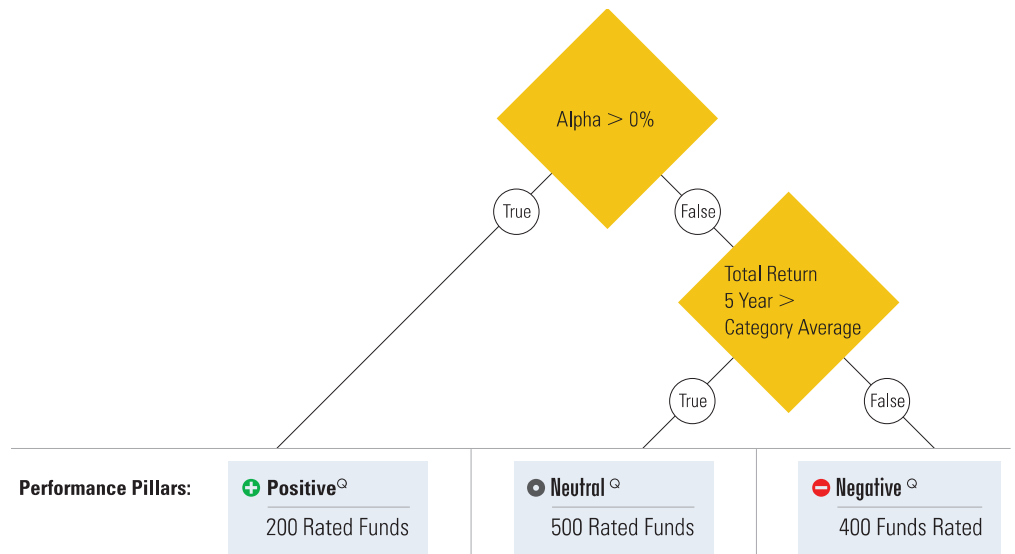
Appendix A: Random Forest

A random forest is an ensemble model, meaning its end prediction is formed based on the combination of the predictions of several submodels. In the case of a random forest, these submodels are typically regression or classification trees (hence the "forest" in "random forest"). To understand the random forest model, we must first understand how these trees are fit.

Regression Trees

A regression tree is a model based on the idea of splitting data into separate buckets based on your input variables. A visualization of a typical regression tree is shown in Exhibit 6. The tree is fit from the top down, splitting the data further into a more complex structure as you go. The end nodes contain groupings of records from your input data. Each grouping contains records that are similar to each other based on the splits that have been made in the tree.

Exhibit 6 Sample Regression Tree With Dummy Data



Source: Morningstar, Inc.

How are splits determined?

As you can see, the tree is composed of nodes that then are split until they reach terminal nodes that no longer split. Each split represents a division of our data based on a particular input variable, such as alpha, or total return five-year versus the category average (Exhibit 6). The algorithm determines where to make these splits by attempting to split our data using all possible split points for all of the input variables, and chooses the split variable and split point to maximize the difference between the variance of the unsplit data and the sum of the variances of the two groups of split data as shown in the following function.

$$VarDiff = \frac{\sum(y - \bar{y}_{presplit})^2}{N_{presplit}} - \left[\frac{\sum(y - \bar{y}_{left})^2}{N_{left}} + \frac{\sum(y - \bar{y}_{right})^2}{N_{right}} \right]$$

Intuitively, we want the split that maximizes the function because the maximizing split is the one which reduces the heterogeneity of our output variable the most. That is, the companies that are grouped on each side of the split are more similar to each other than the pre-split grouping.

A regression or classification tree will generally continue splitting until a set of user-defined conditions has been met. One of these conditions is the significance of the split. That is, if the split does not reduce heterogeneity beyond a user-defined threshold, then it will not be made. Another condition commonly used is to place a floor on the number of records in each end node. These conditions can be made more or less constrictive in order to tailor the model's bias-variance trade-off.

How are the end-node values assigned?

Each tree, once fully split, can be used to generate predictions on new data. If a new record is run through the tree, it will inevitably fall into one of the terminal nodes. The prediction for this record then becomes the arithmetic mean of the output variable for all of the training set records that fell into that terminal node.


Aggregating the Trees

Now that we understand how trees are fit and how they can generate predictions, we can move further in our understanding of random forests. To arrive at an end prediction from a random forest, we first fit N trees (where N can be whatever number desired — in practice, 100 to 500 are common values) and we run our input variables through each of the N trees to arrive at N individual predictions. From there, we take the simple arithmetic mean of the N predictions to arrive at the random forest's prediction.

A logical question at this point is: Why would the N trees we fit generate different predictions if we give them the same data? The answer is: They wouldn't. That's why we give each tree a different and random subset of our data for fitting purposes (this is the "random" part of "random forest"). Think of your data as represented in Exhibit 7.

Exhibit 7 Sample Random Forest Data Representation

● Subsets ○ Random Data



Source: Morningstar, Inc.

A random forest will choose random chunks of your data, including random cross-sectional records as well as random input variables, as represented by the highlighted sections in Exhibit 7, each time it attempts to make a new split. While Exhibit 7 shows three random subsets, the actual random forest model would choose N random subsets of your data, which may overlap, and variables selected may not be adjacent. The purpose of this is to provide each of your trees with a differentiated data set, and thus a differentiated view of the world.

Ensemble models use a "wisdom of crowds" type of approach to prediction. The theory behind this approach is that many "weak learners," which are only slightly better than random at predicting your output variable, can be aggregated to form a "strong learner" so long as the weak learners are not perfectly correlated. Mathematically, combining differentiated, better-than-random, weak learners will always result in a strong learner or a better overall prediction than any of your weak learners individually. The archetypal example of this technique is when a group of individuals is asked to estimate the number of jellybeans in a large jar. Typically, the average of a large group of guesses is more accurate than a large percentage of the individual guesses.

Random forests can also be used for classification tasks. They are largely the same as described in this appendix except for the following changes: Slightly different rules are used for the splitting of nodes in the individual tree models (Gini coefficient or information gain), and the predictor variable is a binary 0 or 1 rather than a continuous variable. This means that the end predictions of a random forest for classification purposes can be interpreted as a probability of being a member of the class designated as "1" in your data.

Appendix B: Pillar—Quantitative Models

Quantitative Parent Pillar Model

What are the Quantitative Parent Pillar threshold values?

The threshold values for the Parent Pillar are set using a symmetric distribution: 10%, 22.5%, 35%, 22.5%, and 10%. The breakpoints for the labels are below:

- ▶ If raw pillar rating ≤ 0.10 , then 1—Low
- ▶ If $0.10 < \text{raw pillar rating} \leq 0.325$, then 2—Below Average
- ▶ If $0.325 < \text{raw pillar rating} \leq 0.675$, then 3—Average
- ▶ If $0.675 < \text{raw pillar rating} \leq 0.90$, then 4—Above Average
- ▶ If raw pillar rating > 0.90 , then 5—High

What variables are used in each of the random forest models (Positive and Negative)?

Each model's variables and their ranked relative importance are shown below. We see that Average Morningstar Rating Overall and Average Net Expense Ratio Rank are the most important input to the Parent Positive Pillar and Parent Negative Pillar model, respectively.

Exhibit 8 Ranked Importance Input Variable for the Quantitative Parent Pillar Model

	+ Positive	- Negative
% Funds Team Managed	9	8
% Funds With a Management Change TTM	22	22
Asset-Weighted Manager Tenure	15	7
Average Actual Management Fee Rank	2	14
Average Net Expense Ratio Rank	10	1
Average Max Management Fee Rank	3	13
Average Prospectus Operating Expense Ratio Rank	19	6
Average Manager Tenure	8	16
Average Morningstar Rating 3 Yr	6	19
Average Morningstar Rating 5 Yr	14	10
Average Morningstar Rating 10 Yr	12	5
Average Morningstar Rating Overall	1	15
Average Number of Months Since Manager Change	16	18
Average Portfolio Transparency TTM	20	20
Firm Age	5	2
Retention 5 Yr	4	12
Risk-Adjusted Success Ratio 3 Yr	11	11
Risk-Adjusted Success Ratio 5 Yr	13	17
Risk-Adjusted Success Ratio 10 Yr	18	3
Success Ratio 3 Yr	7	9
Success Ratio 5 Yr	21	21
Success Ratio 10 Yr	17	4

Source: Morningstar, Inc. Data as of March 31, 2019.

Quantitative People Pillar

What are the Quantitative People Pillar threshold values?

The threshold values for the People Pillar are set using a symmetric distribution: 10%, 22.5%, 35%, 22.5%, and 10%. The breakpoints for the labels are below:

- ▶ If raw pillar rating < 0.10, then 1—Low
- ▶ If .10 < raw pillar rating ≤ 0.325, then 2—Below Average
- ▶ If .325 < raw pillar rating ≤ 0.675, then 3—Average
- ▶ If 0.675 < raw pillar rating ≤ 0.90, then 4—Above Average
- ▶ If raw pillar rating > 0.90, then 5—High

What variables are used in each of the random forest models (Positive and Negative)?

Each model's variables and their ranked relative importance are shown below. We see that Number of Months Since Management Change and Manager Excess Return 5 Yr are the most important inputs for the People Positive Pillar and People Negative Pillar model, respectively.

Exhibit 9 Ranked Importance of Input Variables for the Quantitative People Pillar Model

	+ Positive	- Negative
Actual Management Fee Rank	10	4
Alpha 10 Yr - Category Average	3	—
Asset-Weighted Manager Tenure	7	—
Average Manager Tenure	12	9
Average Morningstar Rating 5 Yr	4	—
Average Morningstar Rating 10 Yr	—	3
Average Morningstar Rating Overall	—	8
Index Fund	16	14
Max Management Fee Rank	11	6
Manager Excess Return 3 Yr	9	10
Manager Excess Return 5 Yr	2	1
Manager Experience	6	5
Manager Investment	—	13
Manager Investment - 1 Million	14	—
Net Expense Ratio Rank	13	11
Number of Months Since Management Change	1	7
Retention 5 Yr	8	2
Success Ratio 5 Yr	5	—
Team Size	15	12

Source: Morningstar, Inc. Data as of March 31, 2019.

Quantitative Process Pillar Model

What are the Quantitative Process Pillar threshold values?

The threshold values for the Process Pillar are set using a symmetric distribution: 10%, 22.5%, 35%, 22.5%, and 10%. The breakpoints for the labels are below:

- ▶ If raw pillar rating < 0.10, then 1—Low
- ▶ If .10 < raw pillar rating ≤ 0.325, then 2—Below Average
- ▶ If .325 < raw pillar rating ≤ 0.675, then 3—Average
- ▶ If 0.675 < raw pillar rating ≤ 0.90, then 4—Above Average
- ▶ If raw pillar rating > 0.90, then 5—High

What variables are used in each of the random forest models (Positive and Negative)?

Each model's variables and their ranked relative importance are shown below. We see that Alpha 10 Yr - Category Average and % Assets in Top 10 Holdings are the most important inputs for the Process Positive Pillar and Process Negative Pillar model, respectively.

Exhibit 10 Input Variable Importance for the Quantitative Process Pillar Model

	+ Positive	- Negative
% Assets in Top 10 Holdings	—	1
Alpha 10 Yr - Category Average	1	—
Alpha 5 Yr - Category Average	6	7
Alpha 5 Yr - Category Benchmark	2	8
Asset-Weighted Manager Tenure	8	—
Average Actual Management Fee Rank	—	14
Average Max Management Fee Rank	—	6
Average Morningstar Rating 10 Yr	3	—
Average Morningstar Rating Overall	—	3
Average Net Expense Ratio Rank	—	9
Average Prospectus Operating Expense Ratio Rank	—	4
Down Capture 3 Yr - Category Average	15	18
Firm Age	11	17
Index Fund	17	—
Information Ratio 5 Yr - Category Average	—	15
Information Ratio 5 Yr - Category Benchmark	4	—
Manager Excess Return 3 Yr	14	5
Manager Excess Return 5 Yr	5	2
Manager Experience	7	10
Max Management Fee Rank	—	12
Net Expense Ratio Rank	—	16
Number of Months Since Management Change	9	13
Retention 5 Yr	13	11
Risk-Adjusted Success Ratio 5 Yr	12	—
Team Size	16	19
Trailing 5 Yr Return Rank	10	—

Source: Morningstar, Inc. Data as of March 31, 2019.

Appendix C: Input Data FAQ

Are all the input variables used in each pillar model?

No. The input variables depend on the pillar model. For example, Manager Investment is only used within the People Negative model. The binary signal of investment helps the model sort between Negative-rated People scores. On the other hand, Manager Investment - 1 Million is only used within the People Positive model to help discern between positively rated People scores.

How do we normalize the input data?

After all data is calculated and collected, we cross-sectionally normalize the data by region to be mean zero and standard deviation 1. This puts everything into the same units (in terms of standard deviation), which makes the data a bit easier to interpret.

How do we assign regions?

In order to normalize by region, we need to know what funds belong to what regions. Countries are assigned to regions based on the Morningstar Region classification system. We assign funds to regions based on the fund's domicile, unless the fund's domicile is not contained within the set of Available for Sale countries. In that case, we choose an Available for Sale country depending on which of those countries belongs to the domicile with the most industrywide assets (for example, U.S. > emerging-markets Asia).

How do we handle missing data?

We use the MissForest algorithm to impute our data. In short, the training set utilizes the analyst pillar decision to impute a more informed value. In the prediction set, we iterate through the five potential pillar outcomes to estimate missing data. For more information on the imputation algorithm, the paper is listed in the References section.

How do we handle category changes?

Input data reflects information available at a given time. Therefore, historical data incorporates the fund's historical category. For performance-related metrics where we require a time series of a fund's category average performance or category index return, we use the monthly track record reflecting the fund's category for that specific month.

How do we handle multiple analyst-assigned Process ratings to passive products tracking the same benchmark?

We select the largest share class with assets under management and then apply the analyst assigned Process score to all other passive products tracking the same benchmark.

What data points are category-relative?

First, most data points will be calculated relative to the category (for example, category average alpha, success ratios, return ranks, beta, fee ranks, star ratings, and so on), but some will not (for example,

tenure, retention ratio, or number of holdings). We prefer to use category-relative data points where possible but tended to refrain when the data point was more operational in nature.

What currency do we use for calculating fund performance statistics?

To estimate fund performance, we convert all fund and index returns to U.S. dollars prior to running our regressions. This eliminates any effects due to the difference in currency return.

What does "average" stand for?

Average stands for an equally weighted average of all share classes given a branding ID.

When are the input data and ratings updated?

The input data and ratings are updated on the 15th day of each month.

When do new funds receive a rating?

A new share class or fund receives a rating when it has a full month of data present. For example, if a new fund is incepted on May 12 and the April production run completes on May 18, then the fund will not receive a rating for the month of April as it has no data for the month of April. Further, when the May production runs on June 18 the fund will not receive a rating for the month of May because the data for the month of May is not complete. The first rating the fund will receive will be a rating for the month of June when June production runs on July 18.

How is the smoothing algorithm implemented for new funds?

For the first two months that a share class or fund receives pillar ratings, the smoothing algorithm step does not occur. The smoothing algorithm requires three months of raw pillar ratings. When the share class or fund receives pillar ratings on the third month, the smoothing algorithm will operate by smoothing the raw pillar ratings of the third month with the first two months of raw pillar ratings. The smoothing algorithm step will occur every month after that.

Why are fee data points used as inputs to the People pillar estimation?

Here fees are directly related to how much a fund charges by managing money for clients, for two reasons. One, our model testing shows that fees do help explain the variance in the People Pillar rating. Two, fees empirically affect all pillars directly or indirectly.

Why do we use the input variables Percentage of Assets in Top 10 Holdings for the Process Pillar?**What is the effective relationship between the variable and the pillar rating?**

Percentage of Assets in Top 10 Holdings is a good indicator to measure how concentrated a fund's portfolio is. The higher the top-10 asset percentage, the more concentrated the portfolio. Such portfolios are implicitly taking on higher risk. The variable reflects a fund's investment philosophy and actual investment process.

Are there fee requirements for funds to receive a Morningstar Quantitative Rating?

Yes. In Europe, we restrict coverage based on the following fee levels:

- ▶ Allocation: Annual Fee < 0.19%
- ▶ Convertibles: Annual Fee < 0.30%
- ▶ Equity Funds: Annual Fee <= 0.22%
- ▶ Fixed-Income ex-Ultrashort Bond: Annual Fee < 0.15%
- ▶ Ultrashort Bond: Annual Fee < 0.10%



About Morningstar® Quantitative Research

Morningstar Quantitative Research is dedicated to developing innovative statistical models and data points, including the Quantitative Equity Ratings and the Global Risk Model.

For More Information

+1 312 244-7541

lee.davidson@morningstar.com



22 West Washington Street
Chicago, IL 60602 USA

©2020 Morningstar. All Rights Reserved. Morningstar's Credit Ratings & Research is produced and offered by Morningstar, Inc., which is not registered with the U.S. Securities and Exchange Commission as a Nationally Recognized Statistical Rating Organization ("NRSRO"). Unless otherwise provided in a separate agreement, you may use this report only in the country in which its original distributor is based. The information, data, analyses and opinions presented herein do not constitute investment advice; are provided solely for informational purposes and therefore are not an offer to buy or sell a security; and are not warranted to be correct, complete or accurate. The opinions expressed are as of the date written and are subject to change without notice. Except as otherwise required by law, Morningstar shall not be responsible for any trading decisions, damages or other losses resulting from, or related to, the information, data, analyses or opinions or their use. The information contained herein is the proprietary property of Morningstar and may not be reproduced, in whole or in part, or used in any manner, without the prior written consent of Morningstar. To order reprints, call +1 312-696-6100. To license the research, call +1 312-696-6869.