

# What is Considered a Good Value for MAPE

Authored by  
**stats writer**

December 14, 2025

## RECOMMENDED CITATION

stats writer (2025). *What is Considered a Good Value for MAPE*. PSYCHOLOGICAL SCALES. Retrieved from <https://scales.arabpsychology.com/?p=107448>

One of the most essential and frequently utilized metrics for quantifying the forecasting accuracy of a statistical model is the **Mean Absolute Percentage Error**, commonly abbreviated as **MAPE**. This metric provides a clear, scalable measurement of error, expressing the deviation between predicted outcomes and actual observations as a percentage. Its simplicity and intuitive interpretability have cemented its position as a cornerstone in fields like supply chain management, finance, and demand forecasting.

Understanding MAPE goes beyond simply knowing the acronym; it requires appreciating how percentage-based errors stabilize comparisons across data sets with varying scales. Because the error is normalized by the actual value, MAPE offers a relative measure of performance. A model that achieves a low MAPE suggests high reliability in its predictive capabilities, offering business leaders and analysts confidence in making operational decisions based on the forecasts generated.

However, while widely popular, MAPE is not without its controversies and nuances. Effective application of MAPE demands an understanding of its mathematical foundation, its limitations, and, most importantly, the contextual benchmarks required to determine what truly constitutes a "good" value in a specific industrial or commercial setting. Without this deeper insight, a MAPE score is just a number lacking actionable meaning.

## Deconstructing the MAPE Formula

To fully appreciate MAPE, we must examine its mathematical construction. The formula dictates how the absolute error is weighted against the actual observation before being averaged across the entire sample size. This structure is critical because it ensures that errors are always treated as positive deviations, and they are always presented relative to the magnitude of the underlying data point.

The standard mathematical formula used to calculate **MAPE** is expressed as follows:

$$\text{MAPE} = (1/n) * \sum(|\text{actual} - \text{forecast}| / |\text{actual}|) * 100$$

This formula is broken down into several essential components, each playing a vital role in determining the final error percentage:

$\Sigma$  - This is the standard mathematical symbol for summation, indicating that we must aggregate the percentage errors calculated for every single observation in the sample.

**n** - Represents the total number of observations, or the **sample size**, used in the evaluation period. The division by 'n' ensures that the result is the mean (average) of the calculated percentage errors.

**actual** - Refers to the true, observed data value for a specific period. This is the benchmark

against which the prediction is measured.

**forecast** - Represents the predicted data value generated by the forecasting model for that same period.

The simplicity of MAPE's final interpretation is perhaps its greatest strength. If a model yields a MAPE value of 8%, it conveys that, on average, the predictions deviated from the actual outcomes by 8%. This percentage format is immediately understandable to non-technical stakeholders, making it an invaluable communication tool in decision-making processes across various departments.

## The Contextual Nature of a "Good" MAPE Score

When initiating a forecasting project, one of the most persistent and critical questions posed by stakeholders is: ***What is a good value for MAPE?*** The realistic, albeit unsatisfying, answer is: **It depends entirely on the context.** Unlike many statistical measures where predefined thresholds exist, MAPE is fundamentally a relative metric. While the objective is always to minimize the error--meaning, the lower the MAPE value, the better the model performance--there is no universal, one-size-fits-all threshold that delineates "good" from "bad."

The determination of an acceptable MAPE requires consideration of several interconnected factors, primarily related to the environment in which the forecast is applied and the volatility inherent in the data being predicted. A 10% MAPE might be considered exceptional in a highly unpredictable market, while the same score might signify failure in a stable, predictable operational context. Therefore, analysts must refrain from comparing their MAPE against arbitrary external benchmarks and instead focus on meaningful internal and industry-specific comparisons.

The two overarching factors that critically influence the interpretation and acceptability of a MAPE value are the specific demands and volatility profiles associated with the industry, and how the model's performance compares against simpler, established baseline methods. Evaluating performance relative to these internal standards is far more informative than chasing an abstract ideal.

## Analyzing Industry-Specific Benchmarks

The acceptable level of MAPE varies dramatically across different sectors. This variation stems from fundamental differences in product demand volatility, pricing strategies, market maturity, and the presence of external, unpredictable factors. For instance, companies operating in highly stable industries, such as essential utilities or certain staple consumer goods, often experience predictable demand patterns that fluctuate minimally over time.

In these stable environments, where demand remains relatively constant and pricing adjustments

are infrequent, forecasting models can achieve remarkably low MAPE values, often falling under 3%. When a company has tightly controlled variables and robust historical data, a high MAPE in this context (e.g., 15%) would signal a severe failure in the chosen forecasting methodology or underlying data quality.

Conversely, industries characterized by high volatility, rapid technological change, intense promotional activity, or seasonal spikes--such as retail fashion, technology gadgets, or high-growth startups--naturally present a greater challenge to predictive modeling. When companies constantly run promotions and specials, consumer purchasing behavior becomes highly erratic. Consequently, predicting demand with pinpoint accuracy is significantly harder, leading to higher, yet still acceptable, MAPE scores. In these dynamic environments, a MAPE in the range of 10% to 20% might be considered standard, or even excellent, depending on the historical difficulty of the predictions.

Therefore, when establishing a performance target for a new forecasting system, analysts must first research the common benchmarks accepted within their specific industry sector. This comparative analysis helps set realistic expectations for the modeling team and provides a sound foundation for justifying the chosen performance threshold to executive leadership.

### **The Importance of Baseline Comparison Models**

A far more rigorous and statistically meaningful approach than seeking an arbitrary external benchmark is to compare the MAPE of your sophisticated model against the performance of simple, established baseline forecasting methods. This internal comparison provides evidence of whether the complexity introduced by the new model--such as incorporating machine learning algorithms or external regressors--actually translates into a tangible and valuable improvement in forecasting accuracy.

If a complex, resource-intensive model yields a MAPE of 7%, but a basic baseline model based on historical averages achieves a MAPE of 7.2%, the marginal gain (0.2 percentage points) is likely not significant enough to warrant the operational effort and cost associated with maintaining the complex model. The goal is not just to get a low MAPE, but to achieve a MAPE that demonstrates clear, demonstrable improvement over simpler alternatives.

When implementing any new forecasting project, analysts should always calculate the MAPE for at least two standard, simple methods to establish a robust performance floor. If the new, advanced model fails to achieve a statistical significance improvement over these baselines, it raises serious questions about the model's utility and suggests that the investment in its development may have been misplaced.

## Key Simple Forecasting Benchmarks

Establishing a baseline typically involves implementing and calculating the MAPE for two classic, yet surprisingly effective, methodologies: the Average Forecasting Method and the Naïve Forecasting Method. These models require minimal computational power and serve as excellent yardsticks for measuring true predictive skill.

There are two well-known simple forecasting models:

**The Average Forecasting Method:** This approach operates on the principle of simplicity and stability. This forecast model predicts that the value for the next upcoming period will be equal to the average of all previously observed historical periods. While seemingly overly simplistic, this method is remarkably robust when dealing with data series that exhibit stability--meaning, data whose statistical properties do not change over time. It helps smooth out minor fluctuations and often performs surprisingly well, making it a powerful baseline comparison.

**The Naïve Forecasting Method:** This is arguably the simplest method available. This model predicts that the value for the next upcoming period will be identical to the value observed in the immediately preceding period. Again, although this method is quite simple, it tends to work surprisingly well for short-term forecasts and serves as a powerful null hypothesis. If a sophisticated model cannot outperform this basic assumption, it lacks statistical significance.

By measuring the MAPE of a new, complex model against both the Average and Naïve forecasts, practitioners gain critical context. Only if the new model's MAPE is substantially lower should the new model be considered truly useful or implemented in a live production environment. If the improvement is marginal, the simpler, cheaper baseline model is usually the preferred operational choice.

## Critical Limitations and Drawbacks of MAPE

Although MAPE is easy to calculate and interpret, its reliance on division by the actual value introduces significant mathematical and practical drawbacks that analysts must be aware of when interpreting the results. These limitations primarily revolve around situations where the actual values are close to zero or involve very low volumes of data.

The two major potential drawbacks to using MAPE include:

**Undefined Error for Zero Actual Values:** Since the formula to calculate absolute percent error is  $|\text{actual}-\text{forecast}| / |\text{actual}|$ , this means that it will be **undefined** if any of the actual values are zero. This issue is particularly problematic in forecasting scenarios where intermittent demand or new product introductions lead to periods of zero sales or inventory movement, rendering the overall MAPE calculation invalid or requiring complex workarounds like adjusted MAPE (AMAPE) or

symmetric MAPE (SMAPE).

**Bias with Low Volume Data:** MAPE tends to produce highly inflated error percentages when dealing with low-volume data or actual values close to zero (even if non-zero). For example, if the actual demand for some item is 2 units and the forecast is 1 unit, the value for the absolute percent error will be  $|2-1| / |2| = 50\%$ , which makes it seem like the forecast error is quite high, despite the forecast only being off by one unit. Because the error is normalized by a tiny denominator, small deviations translate into dramatic percentage errors, potentially leading to misleading performance assessments, particularly for slow-moving inventory items.

These inherent biases mean that MAPE tends to penalize positive errors (under-forecasting) more heavily than negative errors (over-forecasting) and often exaggerates the severity of errors in low-demand environments. Therefore, analysts must exercise caution and potentially shift to alternative metrics when dealing with highly sparse or intermittent data.

## Alternative Error Metrics: MAD and RMSE

Due to the limitations inherent in MAPE, especially when dealing with data series containing zero or near-zero observations, practitioners often turn to alternative measures of forecast error that operate based on absolute scale or squared errors rather than percentages. These alternatives offer a more stable measure of error magnitude across different data scales and volumes.

Potential alternatives to MAPE include the **Mean Absolute Deviation (MAD)** and the **Root Mean Squared Error (RMSE)**.

**Mean Absolute Deviation (MAD):** MAD calculates the average magnitude of the errors without considering their direction. It is the average of the absolute differences between the forecast and actual values. Unlike MAPE, MAD is expressed in the same units as the data itself (e.g., units, dollars), making it highly intuitive for operational teams, but less useful for comparing accuracy across different time series that have vastly different scales.

**Root Mean Squared Error (RMSE):** RMSE is often favored in statistical modeling because it squares the error before averaging them. By squaring the error, RMSE heavily penalizes large errors, meaning a model that produces a few very large mistakes will receive a significantly higher RMSE than a model that produces many small errors, even if the Mean Absolute Deviation is similar. Like MAD, RMSE is also scale-dependent, but its emphasis on large deviations makes it an excellent choice when avoiding large forecasting misses is a primary organizational goal.

The decision of which metric constitutes the "best" measure of accuracy depends entirely on the business objective. If the priority is communication and ease of interpretation across different datasets, MAPE is dominant, provided the data does not contain zero actuals. If the priority is minimizing the impact of severe outliers, **RMSE** is often the superior choice.

## Conclusion: Establishing Meaningful MAPE Targets

In summary, determining what qualifies as a "good" value for MAPE is a process of contextual calibration, not arbitrary adherence to predefined global standards. A truly meaningful MAPE target is one that is informed by the unique volatility profile of the industry and rigorously validated against established, simple forecasting benchmarks.

Analysts should prioritize two main evaluation strategies: first, understanding industry tolerance for error (e.g., a 5% MAPE is excellent in stable industries, but poor for utilities); and second, ensuring the chosen complex model delivers a demonstrably lower MAPE than simple Naïve or Average forecasting methods. If a sophisticated model does not outperform these baselines, its utility is questionable, regardless of its absolute score. By adopting this rigorous, relative approach, MAPE transforms from a confusing number into a powerful, actionable metric for continuous model improvement.

What is Considered a Good RMSE Value?