

the journal of
financial data science

**PORTFOLIO
MANAGEMENT
RESEARCH**

by S&P Global

Fall

2026

Volume

8

Number

4

pm-research.com

Confidence Revisited: The Distribution of Information

**Megan Czasonis, Mark Kritzman,
Fangzhong Liu, and David Turkington**




Megan Czasonis

Megan Czasonis is a Managing Director and Head of Portfolio Management Research at State Street Associates. Her team collaborates with academic partners to develop new research on asset allocation, risk management, and quantitative investment strategy. The Portfolio Management Research team delivers this research to institutional investors through indicators, advisory projects, thought leadership pieces, and interactive tools. Megan has published numerous journal articles, is a co-author of “Prediction Revisited: The Importance of Observation,” and her research received the 2022 Harry Markowitz Award for best paper in the *Journal of Investment Management*. Megan graduated Summa Cum Laude from Bentley University with a B.S. in Economics / Finance.



Mark Kritzman

Mark Kritzman, CFA is a Founding Partner of Windham Capital Management. He also teaches a graduate finance course at the MIT Sloan School. Mark served as a Founding Director of the International Securities Exchange and as an advisor to the Government Investment Corporation of Singapore (GIC). He has written more than 100 articles and is an author of nine books including *Prediction Revisited: The Importance of Observation*, *Asset Allocation: From Theory to Practice and Beyond*, *Puzzles of Finance*, and *The Portable Financial Analyst*. Mark has won numerous awards, including the Roger F. Murray Prize, the Peter L. Bernstein Award, the Harry M. Markowitz Award, and the CFA Institute Research Foundation's Vertin Award. In 2004, Mark was elected a Batten Fellow at the Darden Graduate School of Business Administration, University of Virginia. Mark has a BS in economics from St. John's University, and an MBA with distinction from New York University.



Fangzhong Liu

Fangzhong Liu is a senior quantitative researcher on the Portfolio Management Research team at State Street Associates. His research spans a broad range of topics, including portfolio construction, asset allocation, currency management, and advanced predictive methods. He holds a Bachelor of Science in Mathematics and Economics from the University of Southern California and a Master of Finance from the Massachusetts Institute of Technology.



David Turkington

David Turkington is Senior Managing Director and Head of State Street Associates, State Street Markets' decades-long partnership with renowned academics that produces innovative research on markets and investment strategy. David has published more than 50 research journal articles and 3 books featuring novel approaches to asset allocation, investment strategy, AI, and predictive analytics. His published research has received the Graham and Dodd Scroll Award from the *Financial Analysts Journal*, five Bernstein-Fabozzi/Jacobs-Levy Outstanding Article Awards from the *Journal of Portfolio Management*, the Peter L. Bernstein Award for best paper in an Institutional Investor journal, two Roger F. Murray First Prize awards for outstanding research presented at the Q Group seminars, and two Harry Markowitz awards for best paper in the *Journal of Investment Management*. He is a frequent presenter at industry conferences and serves on the editorial boards of the *Financial Analysts Journal* and *The Journal of Alternative Investments*. David graduated summa cum laude from Tufts University with a BA in mathematics and quantitative economics, and he holds the CFA designation.

Confidence Revisited: The Distribution of Information

Megan Czasonis, Mark Kritzman, Fangzhong Liu,
and David Turkington

Megan Czasonis

is a managing director at
State Street Associates in
Cambridge, MA.

mczasonis@statestreet.com

Mark Kritzman

is the chief executive
officer of Windham Capital
Management in Cambridge,
MA, and a senior lecturer
at the MIT Sloan School of
Management in Cambridge,
MA.

kritzman@mit.edu

Fangzhong Liu

is a vice president at
State Street Associates
in Cambridge, MA.

flu6@statestreet.com

David Turkington

is senior managing director
and head of State Street
Associates in Cambridge,
MA.

dturkington@statestreet.com

KEY FINDINGS

- Classical statistics assesses the confidence of a prediction by measuring error variance, estimation variance, or the sum of both. More complex machine learning models assess confidence using empirical methods, such as resampling or conformal prediction.
- Model-based measures of confidence, however, are likely to be unrealistically tight because they assume that the model is specified correctly and the predictive variables are measured without error.
- Relevance-based prediction, a model-free prediction technique, assesses the reliability of a prediction from the distribution of information that is used to form a prediction. It gives a more realistic assessment of an individual prediction's reliability because it does not rely on model correctness or the assumption that the predictive variables are error-free.

ABSTRACT

Classical statistics assesses the confidence of a prediction by measuring error variance, estimation variance, or the sum of both. More complex machine learning models assess confidence using empirical techniques, such as resampling or conformal prediction. These methods implicitly assume that the model is correctly specified and the predictive variables are measured without error. The authors describe a model-free approach to prediction, called *relevance-based prediction*, which does not rely on a model's correctness or assume that the predictive variables are error-free. It measures the uncertainty of each individual prediction based on the unique distribution of information that is used to form the prediction.

Data-driven predictions are usually formed by models, such as a linear regression model or a neural network. In classical statistics, there are three distinct concepts that relate to a prediction's confidence: Error variance reflects irreducible randomness in outcomes, estimation variance reflects noise in the prediction value that results from error-prone parameter estimates, and predictive variance is the sum of error variance and estimation variance, which determines the width of a prediction interval. These concepts can be derived analytically for a linear regression model. For more complex machine learning models, such as a neural network, it is usually impossible to determine these quantities analytically, but one may approximate them using empirical techniques, such as resampling or conformal prediction. In all these cases, for both classical statistics and machine learning, the estimated measures of confidence may be unrealistically tight because they assume that the model is correctly specified and the predictive variables are measured without error.

We propose a new framework for measuring prediction uncertainty based on a model-free prediction technique called *relevance-based prediction* (RBP). RBP gives a measure of prediction uncertainty that does not depend on the correctness of a parameterized model, because there is no model, nor does it assume that the predictive variables are error-free. RBP forms a prediction as a weighted average of observed outcomes, which allows us to see precisely how each observation informs the prediction. This transparency allows us to observe the distribution of information that goes into each prediction, whereby each observation corresponds to a solo prediction and a probability weight. RBP associates prediction uncertainty with the alignment of the information that goes into a prediction. If the information is mutually supportive, which would show up as a tight distribution, we should be more confident in the prediction. If it is mutually contradictory, as evidenced by a wide distribution, we should be less confident.

To build intuition for this approach, consider a simple nearest-neighbors prediction rule that also forms a prediction as a weighted average of outcomes. If we include only one nearest neighbor, its solo prediction equals its outcome, and its probability weight equals 1. If we include n nearest neighbors, their solo predictions are their respective outcomes and their weights are $1/n$ each. RBP is much more sophisticated and powerful than a nearest-neighbors rule. RBP weights reflect both how an observation extrapolates to the current prediction task and the importance of the observation. We must, therefore, transform RBP's original outcomes and weights into a set of solo predictions and modified weights that we can interpret as a distribution.

We proceed as follows. First, we describe partial sample regression, which is a building block of RBP. It forms a prediction as a weighted average of observed outcomes in which the weights come from a precisely defined and theoretically justified statistic called *relevance*. Linear regression predictions emerge as a special case of partial sample regression. We show how to construct distributions of the information that is used to form individual predictions, and we provide simulations to support our analysis. We then introduce the notion of *fit*, which quantifies the prevalence of useful patterns in a dataset pertaining to one specific prediction. Fit enables us to compute a composite prediction from a prediction grid that comprises many combinations of observations and predictive variables. Each cell in the prediction grid gives a partial sample regression prediction and an associated measure of reliability. The grid's composite prediction is formed as a reliability-weighted average of each cell's prediction. We describe how to form distributions of information aggregated from all the cells that form the grid's composite prediction, and we give an empirical illustration of solo distributions derived from grid prediction. Finally, we compare fit and the distribution of solo predictions with an analogy to a signal-to-noise ratio and then conclude with a summary.¹

PARTIAL SAMPLE REGRESSION

Partial sample regression forms a prediction as a weighted average of observed outcomes in which the weights are based on a statistic called *relevance*. Relevance is composed of similarity and informativeness, which are both measured as Mahalanobis distances, as shown in Equations 1 through 4.

$$r_{it} = \text{sim}(x_i, x_t) + \frac{1}{2}(\text{info}(x_i, \bar{x}) + \text{info}(x_t, \bar{x})) \quad (1)$$

¹The descriptions of partial sample regression and relevance-based prediction throughout this article follow closely language used from Czasonis et al. (2022a, 2022b, 2023, 2025a, 2025b, 2025c), but they are modified to fit the context of the current discussion.

$$\text{sim}(x_i, x_t) = -\frac{1}{2}(x_i - x_t)\Omega^{-1}(x_i - x_t)' \quad (2)$$

$$\text{info}(x_i, \bar{x}) = (x_i - \bar{x})\Omega^{-1}(x_i - \bar{x})' \quad (3)$$

$$\text{info}(x_t, \bar{x}) = (x_t - \bar{x})\Omega^{-1}(x_t - \bar{x})' \quad (4)$$

In Equations 1 through 4, x_i is a vector of the values of K predictive variables for a prior observation, x_t is a vector of the values of the predictive variables for a specific prediction task, $\bar{x} = \mathbf{1}_N' \mathbf{X} \mathbf{N}^{-1}$ is the average of the predictive variables across all observations, Ω^{-1} is the inverse of the covariance matrix of all the observations of the variables, and $'$ denotes matrix transpose. All else being equal, observations that are like the prediction circumstances but different from average circumstances are more relevant than those that are not.

This definition of relevance is not arbitrary. We know from information theory that the information contained in an observation is the negative logarithm of its likelihood.² We also know from the central limit theorem that the relative likelihood of an observation from a multivariate normal distribution is proportional to the exponential of a negative Mahalanobis distance. Therefore, the information contained in a point on a multivariate normal distribution is proportional to a Mahalanobis distance. If we define weights in terms of relevance, as shown in Equation 5, which admits the relevance-weighted average of every prior outcome in the full sample of observations, the result is precisely equivalent to the prediction given by linear regression.³

$$w_{it, \text{linear}} = \frac{1}{N} + \frac{1}{N-1} r_{it} \quad (5)$$

Partial sample regression assumes that we can produce a more reliable prediction by censoring observations that are less relevant than a chosen threshold, which leads to the following definition of prediction weights.

$$w_{it, \text{partial}} = \frac{1}{N} + \frac{\lambda^2}{n-1} (\delta(r_{it}) r_{it} - \phi \bar{r}_{\text{sub}}) \quad (6)$$

$$\delta(r_{it}) = \begin{cases} 1 & \text{if } r_{it} \geq r^* \\ 0 & \text{if } r_{it} < r^* \end{cases} \quad (7)$$

$$\lambda^2 = \frac{\sigma_{r, \text{full}}^2}{\sigma_{r, \text{retained}}^2} = \frac{\frac{1}{N-1} \sum_{i=1}^N r_{it}^2}{\frac{1}{n-1} \sum_{i=1}^N \delta(r_{it}) r_{it}^2} \quad (8)$$

In Equations 5 through 8, $n = \sum_{i=1}^N \delta(r_{it})$ is the number of observations that are fully retained, $\phi = n/N$ is the fraction of observations in the retained sample, and $\bar{r}_{\text{sub}} = \frac{1}{n} \sum_{i=1}^N \delta(r_{it}) r_{it}$ is the average relevance value of the observations in the retained sample. It is important to note that $w_{it, \text{partial}}$ depends crucially on the prediction

² Shannon showed that information is an inverse logarithmic function of probability, which is a key insight from his comprehensive theory of communication. See Shannon (1948).

³ See Czaronis et al. (2023) for proof of this result.

circumstances x_t . Relevance is reassessed for each prediction circumstance, which further affects the identification of the retained subsample and introduces nonlinear conditional dependence of the prediction $\hat{y}_t$ on the prediction circumstances x_t . The scaling factor λ^2 compensates for a bias that would otherwise result from relying on a small subsample of highly relevant observations. In the case of linear regression, $n = N$ and $\lambda^2 = 1$. Lastly, note that the weights always sum to 1.⁴

The prediction given by partial sample regression is shown in Equation 9.

$$\hat{y}_t = \sum_i \left(\frac{1}{N} + \frac{\lambda^2}{n-1} (\delta(r_{it})r_{it} - \phi\bar{r}_{sub}) \right) y_i \quad (9)$$

We can equivalently express the prediction in terms of deviations from the full sample $\bar{y}$, as shown in Equation 10.

$$\hat{y}_t - \bar{y} = \frac{\lambda^2}{n-1} \sum_i \delta(r_{it})r_{it}(y_i - \bar{y}) \quad (10)$$

Even though Equation 10 uses $\bar{y}$ based on the full sample, only noncensored retained observations cause the prediction to tilt away from the full sample average. The prediction's tilts emanate only from the subsample where $\delta(r_{it}) = 1$. This feature allows us to construct a predictive distribution because we ignore censored observations. We cannot ignore censored observations in Equation 6 because the term $\phi\bar{r}_{sub}$ implicitly manufactures the effect of the full sample $\bar{y}$.

As we have shown previously,⁵ we can express λ^2 equivalently as follows:

$$\lambda^2 = \frac{\sigma_{r,full}^2}{\sigma_{r,retained}^2} = \frac{info(x_t)}{\sigma_{r,retained}^2} = \frac{info(x_t)}{\frac{1}{n-1} \sum_j \delta(r_{jt})r_{jt}^2} = \frac{(n-1)info(x_t)}{\sum_j \delta(r_{jt})r_{jt}^2} \quad (11)$$

These equivalences allow us to cancel out $n - 1$ in Equation 10, giving us a new equation.

$$\hat{y}_t - \bar{y} = \frac{info(x_t)}{\sum_j \delta(r_{jt})r_{jt}^2} \sum_i \delta(r_{it})r_{it}(y_i - \bar{y}) \quad (12)$$

SOLO PREDICTION

By removing $n - 1$ from Equation 10, we are now able to make a prediction from a single observation, $solo(i)$, which has deviations from $\bar{y}$, as shown in Equations 13 and 14.

$$\hat{y}_{t,solo(i)} - \bar{y} = \frac{info(x_t)}{r_{it}^2} r_{it}(y_i - \bar{y}) \quad (13)$$

$$\hat{y}_{t,solo(i)} - \bar{y} = \frac{info(x_t)}{r_{it}} (y_i - \bar{y}) \quad (14)$$

⁴ See Czaronis et al. (2023) for proof of this result.

⁵ See Czaronis et al. (2022a).

The solo prediction scales the outcome deviation ($y_i - \bar{y}$) by a simple factor, which is intuitive. If the observation is equal to x_t , the scaling factor will equal 1 and the prediction will be formed from the outcome that occurred in the identical circumstance. All else being equal, as the informativeness of the current prediction circumstance x_t increases, the prediction amplifies the impact of the deviation for observation i .⁶ If observation i is highly relevant, the deviation will be scaled back. If the observation is not very relevant, it will be scaled up. This scaling is necessary to express the observed outcome in the same expected scale as the prediction task. The scaling factor also contains a positive or negative sign, so it has the potential to flip the direction of the deviation. The solo prediction is technically undefined when relevance equals exactly zero because the observation is orthogonal to the task and cannot be used to inform it.

As we show next, we can combine solo predictions to form the actual prediction, thereby accounting for multiple observations. Observations that we censor using $\delta(r_{it}) = 0$ receive zero weight in the composite sum, thus nullifying the impact of their extreme solo predictions. Solo predictions that are undefined because relevance equals exactly zero also receive zero weight in the composite sum, and it is appropriate to treat these instances as zero contributions to the composite prediction.

CONTRIBUTION OF A SOLO PREDICTION

Consider the contribution, $contr(i)$, of a single observation to a multi-observation prediction's deviation from the unconditional average:

$$\hat{y}_{t,contr(i)} - \bar{y} = \frac{info(x_t)}{\sum_j \delta(r_{jt}) r_{jt}^2} r_{it} (y_i - \bar{y}) \quad (15)$$

If we express contribution as a weight ξ_i times that observation's solo prediction:

$$\hat{y}_{t,contr(i)} - \bar{y} = \xi_i (\hat{y}_{t,solo(i)} - \bar{y}) \quad (16)$$

Then

$$\xi_i = \frac{\hat{y}_{t,contr(i)} - \bar{y}}{\hat{y}_{t,solo(i)} - \bar{y}} \quad (17)$$

This substitution cancels out many terms, leading to

$$\xi_i = \frac{r_{it}^2}{\sum_j \delta(r_{jt}) r_{jt}^2} \quad (18)$$

The weight placed on the solo prediction of a retained observation ($\delta(r_{it}) = 1$) is equal to its squared relevance as a fraction of the sum of the squared relevance of all the other retained observations. Clearly these weights sum to 1, they are all nonnegative, and the weights of the censored observations in this context are all 0.

⁶ Each prediction contains both a signal and noise. As the impact is scaled up, so too is the noise, which leads to more uncertainty about the prediction.

BUILDING A DISTRIBUTION OF SOLO PREDICTIONS

Rewriting the weight formula so that the weights are all positive allows us to calculate a distribution of solo predictions. Each solo prediction is scaled and signed as a conceptually viable prediction that receives a weight ξ_i in the actual partial sample regression prediction. Together, these solo prediction deviations and weights define a histogram or distribution of relative likelihood.

Censored observations receive zero weight, so they are censored from the distribution as they are from the prediction. They have no effect irrespective of their solo prediction values. Observations with very low relevance may result in unrealistic solo predictions, but their probability weights in the histogram will be extremely small. Rather than showing unrealistic outliers in a distribution that extrapolates noise to extremes, a graphical solution for this issue is to show cumulative tail bars on both sides of the distribution, limiting the display to a reasonable range. This distribution will only reflect retained observations, just as the prediction itself only reflects retained observations, setting aside the computation of $\bar{y}$.

It is important to distinguish this distribution from the distributions implied by model-based estimates of error variance and estimation variance. The distribution of solo predictions is a distribution of the information that is used to form each individual prediction $\hat{y}_{t,solo(i)}$. It, therefore, allows us to measure prediction uncertainty in a way that captures nuanced details about how each prediction is formed that would be obscured by a distribution that is derived from summary statistics or empirically resampled predictions. Moreover, unlike classical measures of confidence that are based on the same theoretical distribution for all predictions, the distributions of solo predictions are specific to each individual prediction task.

SIMULATIONS OF THE DISTRIBUTION OF SOLO PREDICTIONS

Linear Regression

We now simulate a synthetic data sample of 100,000 observations to illustrate the distributions of $\hat{y}_{t,solo(i)}$. First, we randomly draw values of a single predictive X variable from a standard normal distribution. Next, we form outcomes Y as the values of the predictive variable plus the values of a simulated error term, which also comes from a standard normal distribution.


To provide a point of reference with classical statistics, we show the distribution of $\hat{y}_{t,solo(i)}$ for predictions given by linear regression, which is the only model-based prediction for which we observe the distribution of $\hat{y}_{t,solo(i)}$. We can observe this distribution owing to the equivalence of linear regression with RBP in the special case in which partial sample regression is applied to all the observations.

The left panel of Exhibit 1 shows the distribution of the solo predictions for the prediction task of $x_t = 1$, and the right panel shows the distribution for the task $x_t = 2$. The gray bars show the distribution of the solo predictions given by each observation, and the orange lines show the classical distribution of potential outcomes implied by the error variance estimated from the regression residuals. Both are probability density functions that integrate to 1.

This example highlights the differences between RBP's distribution of solo predictions, classical error variance, and classical estimation variance. In this experiment, the classical estimation variance of the prediction is extremely small because the noise in 100,000 observations mostly cancels out in the prediction estimates. In other words, the predictions $\hat{y}_t$ are nearly certain to be close to their

EXHIBIT 1

Solo Distributions and Classical Error Variance Distributions



NOTE: RBP = relevance-based prediction.

theoretical values given this large sample. We do not show the distributions implied by the estimation variance in Exhibit 1, because they would be extremely pointed narrow distributions near the average prediction. By contrast, the classical error variance, which creates the distributions shown in orange, does not become so narrow, because no matter how reliable the overall prediction value becomes, actual outcomes include the influence of the error term that has an irreducible variance. Likewise, the distribution of solo predictions does not converge to a narrow outcome, because the information contributed by each observation contains noise from the randomness in the data sample.

Exhibit 1 has several interesting takeaways. First, the distributions of the solo predictions when the value of the predictive variable equals 1 and 2 are both reasonably close to the classical error distribution, which confirms that they are sensible. Second, the distributions of solo predictions are different from each other, which highlights the fact that the solo predictions depend crucially on the prediction circumstances. The classical error distribution, by contrast, has the same variance irrespective of the value of the predictive variable. It is based on the average squared residuals across all predictions and does not recognize that some predictions can be more reliable than others. It is worth noting that if we average the variance of the solo prediction distributions across all prediction tasks x_t in the sample, the average of all these distributions will converge to the classical error variance for the special case in which there is a single predictive variable. See the Appendix for an explanation of this result.

Partial Sample Regression

We now illustrate distributions of $\hat{y}_{t,solo(i)}$ for predictions given by partial sample regression based on simulated data from a more complex data-generating process. Partial sample regression is based on the premise that we can form more reliable predictions from subsamples of relevant observations than the full sample of observations. This is because in many prediction circumstances, nonrelevant observations contradict or obscure useful patterns that would otherwise allow us to form a more reliable prediction.

To illustrate the information distributions that come from partial sample regression, we draw simulated observations from two distinct regimes whose observations for X collectively form a bimodal distribution. One of the regimes, which we refer to as the *majority regime*, occurs 75% of the time. Its predicted outcomes are equal to the values of X drawn from a normal distribution with a mean of 5 and a standard deviation of 3, plus an error term that is normally distributed with a mean of 0 and a standard deviation of 3. The second regime, which we refer to as the *minority regime*, occurs 25% of the time. Its predicted outcomes are equal to two times the values of X drawn from a normal distribution with a mean of -5 and a standard deviation of 2, plus an error term that is normally distributed with a mean of 0 and a standard deviation of 3.


Exhibit 2 shows boxplots demarcating the median, 25th, 75th, 10th, and 90th percentiles, along with the actual prediction (average), for predictions of a majority regime circumstance and a minority regime circumstance given several censoring thresholds. This exhibit reveals that the distributions as summarized by these boxplots are not far apart from each other for the majority regime shown in Panel A, which makes sense because the majority regime is highly prevalent and, therefore, well represented by most subsamples of observations.

The distributions for the minority regime in Panel B are wider because uncertainty in X is magnified as a result of its greater contribution to the Y outcomes. In addition, the more focused subsamples of relevant observations give notably tighter distributions. This occurs because partial sample regression carefully curates the observations to predict uncommon outcomes, censoring the irrelevant information from the majority regime.

This insight raises the question of how to determine the optimal subsample of relevant observations for a given prediction task, which leads us to the broader formulation of relevance-based prediction.

EXHIBIT 2

Boxplots of Predictions of Outcomes for Different Censoring Thresholds



RELEVANCE-BASED PREDICTION

To focus attention on the distribution of $\hat{y}_{t, \text{solo}(t)}$, we illustrated it using partial sample regression, which is a simplified and less refined version of RBP. RBP extends partial sample regression in two important ways: It applies a principled way for determining the optimal subsample of relevant observations for each prediction task based on a statistic called *fit*, and it incorporates a prediction grid to account for the codependence of observations and predictive variables on the specific circumstances of each individual prediction task.

Fit

Fit quantifies the prevalence of useful patterns in a dataset for the purpose of forming a specific prediction. It reveals how much confidence we should have in a prediction, separate from the confidence we have in the overall prediction routine.⁷ Thus, it also provides a principled way to evaluate the relative efficacy of alternative calibrations before a prediction is rendered. Mathematically, fit measures the degree of alignment between weights and outcomes and may be expressed in terms of products for pairwise observations, as shown in Equation 19, or equivalently as the squared correlation between weights and outcomes, as shown in Equation 20. We should also note that all the formulas we have thus far considered for weights rely only on relevance, which in turn relies only on the x_t s, the x_t , and the $\bar{x}$. They do not use any of the information from outcomes. To determine fit, however, we must consider outcomes (the y_t s).

$$fit_t = \frac{1}{(N-1)^2} \sum_i \sum_j z_{w_{it}} z_{w_{jt}} z_{y_i} z_{y_j} \quad (19)$$

$$fit_t = \rho(w_t, y)^2 \quad (20)$$

In Equation 19, $z_{y_i} = \frac{y_i - \bar{y}}{\sigma_y}$ and $z_{w_{it}} = \frac{w_{it} - 1/N}{\sigma_{w_{it}}}$. Note that for partial sample weights, stronger patterns in the retained subsample may increase fit, whereas the presence of near-zero weights for censored observations will dampen fit. Fit evaluates the fundamental trade-off associated with focused subsamples. Like relevance, fit is not arbitrary. In the case of linear regression with $n = N$, the informativeness-weighted average fit across all prediction tasks in the observed sample equals R^2 .⁸

$$R^2 = \frac{1}{N-1} \sum_{t=1}^N info(x_t, \bar{x}) fit_t \quad (21)$$

Censoring observations that fall below a relevance threshold is more effective to the extent that there is asymmetry between the fit of the weights formed from the retained subsample of observations and the fit of the weights formed from the complementary set of censored observations. We measure asymmetry between the fit of the retained and censored subsamples, as shown in Equation 22. The (+) superscript designates weights formed from the retained observations, and the (−) superscript designates weights formed from the censored observations. Asymmetry recognizes the benefit of censoring nonrelevant observations that contradict the

⁷ For evidence of fit's ability to distinguish more reliable predictions from less reliable predictions, see Czaronis et al. (2025b, 2025c).

⁸ See Czaronis et al. (2022a) for proof of this result.

predictive relationships that exist among the relevant observations. This assessment also inherently considers the relative sample sizes of the two subsamples.

$$asymmetry_t = \frac{1}{2} (p(w_t^{(+)}, y) - p(w_t^{(-)}, y))^2 \quad (22)$$

To calculate adjusted fit, we add asymmetry to fit and multiply this sum by K , the number of predictive variables included in the prediction, as shown in Equation 23. Multiplication by the number of predictive variables allows us to compare predictions based on different numbers of predictive variables. Adjusted fit recognizes that we are more likely to observe a spurious relationship from prediction weights based on just one or a few variables than we are based on a collection of many variables.

$$adjusted\ fit_t = K(fit_t + asymmetry_t) \quad (23)$$

Grid Prediction

Grid prediction employs a grid in which the columns represent different combinations of predictive variables, and the rows represent subsamples of observations determined by different relevance thresholds. Each cell contains a prediction and an associated adjusted fit. The assessment of reliability using adjusted fit occurs before the prediction is rendered and the subsequent outcome is known. Grid prediction forms a composite prediction as a reliability-weighted average of the predictions from all possible calibrations. Equation 24 defines reliability weights, ψ_θ , as the adjusted fit for a parameter calibration, θ , divided by the sum of all adjusted fits across all parameter calibrations.

$$\psi_\theta = \frac{adjusted\ fit_\theta}{\sum_{\tilde{\theta}} adjusted\ fit_{\tilde{\theta}}} \quad (24)$$

Equation 25 describes the composite grid prediction.

$$\hat{y}_{t,grid} = \sum_{\theta} \psi_\theta \hat{y}_{t,\theta} \quad (25)$$

Note that each cell's prediction is a linear function of observations, and the grid prediction is a linear function of each cell's prediction. Therefore, we can express the grid prediction in terms of composite weights applied to each observation, as shown in Equation 26. Composite weights are important because they preserve the transparency of each observation's contribution to the current prediction task, and they allow us to calculate fit from composite weights as a final gauge of the grid prediction's reliability.

$$w_{it,grid} = \sum_{\theta} \psi_\theta w_{it,\theta} \quad (26)$$

ILLUSTRATION OF DISTRIBUTIONS OF SOLO PREDICTIONS FROM RBP'S PREDICTION GRID

As an illustration, we now apply grid prediction to forecast the future one-year change in the effective federal funds rate. We use two predictive variables:

- X1: Employment, measured as the annual change in nonfarm payrolls
- X2: Inflation, measured as the annual change in personal consumption expenditures, excluding food and energy

We consider all three combinations of the predictive variables: X1 and X2, X1 by itself, and X2 by itself. We consider four subsamples of observations in the grid based on censoring thresholds of 0% (full sample), 20%, 50%, and 80%. For each prediction, we filter on both relevance and similarity.⁹ Therefore, the prediction grid has 24 cells (3 variable combinations times 4 subsamples times 2 censoring criteria). For our training data, we use nonoverlapping calendar year observations from 1961 to 2023.


We highlight two predictions: the one-year rate change for 2023 and the one-year rate change for 2025. For each prediction, our inputs reflect the employment and inflation values for the prior year. For 2023, RBP predicted that the effective federal funds rate would increase by 90 bps compared with an actual increase of 120 bps. For the calendar year that ended in December 2025, RBP predicted the effective federal funds rate would decrease by 30 bps compared to an actual decrease of 75 bps.

The composite predictions that are given by the grid are fit-weighted averages of the individual cell predictions. And each cell's prediction is a weighted average of the solo predictions from the individual observations that are used by each cell. We construct a distribution of the information given by pooling the solo predictions across all the cells in the grid, which we show in Exhibit 3. The dotted lines in Exhibit 3 show the aggregate predictions, which are the means of each distribution.

Exhibit 3 demonstrates how the grid preserves and assembles the information from each observation across all the grid cells to form composite distributions of all the information used to form each prediction. It also reveals that these distributions differ from prediction to prediction and that they are distinctly different from classical distributions based on a model's summary statistics. From first principles, we should expect the distributions of information given by the grid to have greater resolution than distributions composed from a single combination of variables and observations because the grid distributions capture a much broader set of distinct solo predictions based on different combinations of variables.

EXHIBIT 3

Distribution of Solo Predictions from Grid Prediction



⁹It is often helpful to censor the observations on just similarity as well as relevance, though we always form the prediction as a relevance-based weighted average. Moreover, we need not choose which censoring criterion to use. We let adjusted fit determine the best censoring criterion.

THE CONNECTION OF FIT AND SOLO PREDICTIONS

We described two distinct but related measures of a prediction's reliability: fit and the distribution of solo predictions, which is the distribution of information that is used to form a prediction. Both give insight into the reliability of a prediction, but in different ways. Fit measures the strength of the signal within a dataset for a particular prediction. It quantifies the alignment between the observation weights given by the X variables and the outcomes given by the y variable. The distribution of the solo predictions gives a visual description of the noise around the prediction. Although we can compute the variance of the solo predictions, as we show in the Appendix, its greater value is that it reveals how the noise is shaped. Fit can be thought of as the signal for the prediction, and the distribution of solo predictions can be thought of as the noise associated with the signal.

We can also think of a solo prediction as a single observation's vote on what the prediction should be. If the votes are alike, we know that there is strong agreement across the observations. If the votes are widely dispersed, we know that there is considerable disagreement. If one vote is extremely large, we know that a single observation strongly influences the prediction. If the votes are bimodal, we know that the average prediction is unlikely to occur and that we need to consider the circumstances that bifurcate the distribution.

We would never be able to observe the solo predictions given by a regression model except in the special case in which it converges to RBP (full sample, all variables), or a neural network, which is hopelessly opaque. RBP gives full visibility into the alignment of observation weights and outcomes, which is fit, and into the distribution of the observations' solo predictions. Models, by contrast, only yield summaries of the nuanced information revealed by RBP, and summaries, without accompanying transparency into the requisite details, can be very misleading.

SUMMARY

We first compared model-based assessments of confidence with RBP's information-based measure of prediction uncertainty. Model-derived distributions tell us about the range of prediction errors that we should expect and the range of estimated predictions that we should expect if we repeated the prediction with different data. We should have less confidence in a prediction if these distributions are wide than if they are tight. These model-based methods, however, assume that the model is correctly specified and that the predictive variables are measured without error. By contrast, RBP assesses a prediction's reliability on the basis of the consistency of the information that is used to form each individual prediction. RBP tells us that we should have more trust in a prediction if it is formed from information that is mutually supportive than if it is formed from information that is mutually contradictory.

We then introduced the notion of partial sample regression, which is a model-free prediction technique that forms a prediction as a weighted average of observed outcomes in which the weights are based on a precise and theoretically justified statistic called relevance. We explained that the premise of partial sample regression is to form a prediction from a subsample of relevant observations, which in most cases gives a more reliable prediction than the full sample of observations.

We then showed how the model-free nature of partial sample regression enables us to form single-observation predictions called solo predictions. These solo predictions are the building blocks that aggregate to the prediction given by all the observations that are used by partial sample regression to form the prediction.

We evaluated the distribution of solo predictions on simulated distributions, assuming that both the observations and errors come from a standard normal distribution, for a special case of partial sample regression that uses the full sample of observations. This special case gives the same prediction as linear regression, which enabled us to compare our information distributions with the classical distribution that comes from the stylized assumptions of linear regression. This comparison showed how the distributions of information differ from each other and from the classical distribution of possible values.

We then evaluated distributions of solo predictions in which the observations were drawn from a bimodal distribution composed of a majority regime and a minority regime. We summarized these distributions with boxplots. We showed that the boxplots were relatively similar for predictions of outcomes associated with the majority regime but different for outcomes associated with the minority regime. This distinction highlighted the fact that partial sample regression carefully curates the observations to predict unusual outcomes.

We next introduced RBP, which depends crucially on fit. Fit quantifies the prevalence of useful patterns in a dataset for forming a prediction. We showed how fit gives advance guidance of a prediction's reliability, which enables us to identify the best subsample of observations for a given prediction task. We discussed the connection of fit to solo predictions within the context of a signal-to-noise ratio, with fit representing the signal and the distribution of solo predictions representing the noise associated with the signal. We then extended this concept to the choice of predictive variables. We discussed how the choice of both observations and predictive variables is codependent on the unique circumstances of each prediction task. We described the prediction grid, which considers a vast number of combinations of observations and predictive variables. The columns of the grid represent different combinations of predictive variables, and the rows represent different subsamples of observations based on different censoring thresholds. Each cell in the grid has a prediction and an associated fit. The grid forms a composite prediction as a fit-weighted average of the predictions across all the grid cells. It, therefore, produces a composite prediction that is diversified across many calibrations, but in a way that bends toward those combinations of observations and predictive variables that give more reliable predictions.

We then presented a simple illustration of grid prediction in which we predicted annual changes in the effective federal funds rate. We demonstrated that because each cell's prediction is a weighted average of solo predictions and because the grid's composite prediction is a weighted average of all the cell predictions, we can construct distributions of all the information in the grid that is used to form a composite prediction.

Finally, we described how fit and the distribution of solo predictions are related within the context of a signal-to-noise ratio. Fit measures the strength of the signal that comes from the alignment of the observation weights and outcomes within a dataset, and the distribution of solo predictions represents the noise associated with the signal.

In conclusion, we do not assert that RBP's measure of prediction uncertainty is necessarily better than a model-based assessment of confidence. It does, however, give a completely novel perspective about a prediction's reliability, and it does so in a way that does not assume perfect accuracy of the model and error-free observation of the predictive variables. We, therefore, recommend it as a valuable complement to conventional measures of confidence.

APPENDIX

SOLO DISTRIBUTION VARIANCE VS. LINEAR REGRESSION

In the following analysis, we derive an expression for the average variance of the RBP solo prediction distributions across all prediction tasks in the case of a linear regression and we relate it to the classical variance of the error term from a linear model.

Let us assume without loss of generality that $\bar{y} = 0$. The following expressions give the solo predictions and their associated weights.

$$\hat{y}_{t,solo(i)} = \frac{info(x_t)}{r_{it}} y_i \quad (A-1)$$

$$\xi_i = \delta(r_{it}) \frac{r_{it}^2}{\sum_j \delta(r_{jt}) r_{jt}^2} \quad (A-2)$$

Let us calculate the mean of the distribution as a probability-weighted sum:

$$\mu_{t,solo} = \sum_i \xi_i \hat{y}_{t,solo(i)} \quad (A-3)$$

$$\mu_{t,solo} = \sum_i \delta(r_{it}) \frac{r_{it}^2}{\sum_j \delta(r_{jt}) r_{jt}^2} \frac{info(x_t)}{r_{it}} y_i \quad (A-4)$$

$$\mu_{t,solo} = \frac{info(x_t)}{\sum_j \delta(r_{jt}) r_{jt}^2} \sum_i \delta(r_{it}) r_{it} y_i \quad (A-5)$$

$$\mu_{t,solo} = \hat{y}_t \quad (A-6)$$

Now, calculate the variance of the distribution:

$$\sigma_{t,solo}^2 = \sum_i \xi_i (\hat{y}_{t,solo(i)} - \hat{y}_t)^2 \quad (A-7)$$

$$\sigma_{t,solo}^2 = \sum_i \xi_i \left(\frac{info(x_t) y_i}{r_{it}} - \hat{y}_t \right)^2 \quad (A-8)$$

$$\sigma_{t,solo}^2 = \sum_i \delta(r_{it}) \frac{r_{it}^2}{\sum_j \delta(r_{jt}) r_{jt}^2} \left(\frac{info(x_t)^2 y_i^2}{r_{it}^2} + \hat{y}_t^2 - \frac{2info(x_t) \hat{y}_t y_i}{r_{it}} \right) \quad (A-9)$$

$$\sigma_{t,solo}^2 = \left(\frac{\sum_i \delta(r_{it}) info(x_t)^2 y_i^2}{\sum_j \delta(r_{jt}) r_{jt}^2} + \frac{\hat{y}_t^2 \sum_i \delta(r_{it}) r_{it}^2}{\sum_j \delta(r_{jt}) r_{jt}^2} - \frac{2info(x_t) \hat{y}_t \sum_i \delta(r_{it}) r_{it} y_i}{\sum_j \delta(r_{jt}) r_{jt}^2} \right) \quad (A-10)$$

$$\sigma_{t,solo}^2 = \frac{info(x_t)^2 \sum_i \delta(r_{it}) y_i^2}{\sum_j \delta(r_{jt}) r_{jt}^2} + \hat{y}_t^2 - \frac{2info(x_t) \hat{y}_t}{\sum_j \delta(r_{jt}) r_{jt}^2} \quad (A-11)$$

$$\sigma_{t,solo}^2 = \frac{info(x_t) \left(\frac{1}{N-1} \right) \sum_j r_{jt}^2 \left(\frac{1}{n-1} \right) \sum_i \delta(r_{it}) y_i^2}{\left(\frac{1}{n-1} \right) \sum_j \delta(r_{jt}) r_{jt}^2} + \hat{y}_t^2 - \frac{2 \left(\frac{1}{N-1} \right) \sum_j r_{jt}^2 \hat{y}_t}{\frac{1}{n-1} \sum_j \delta(r_{jt}) r_{jt}^2} \quad (A-12)$$

$$\sigma_{t,solo}^2 = \lambda^2 info(x_t) \sigma_{y,partial}^2 - (2\lambda^2 - 1) \hat{y}_t^2 \quad (A-13)$$

$$\sigma_{t,solo}^2 = \lambda^2 info(x_t) \sigma_{y,partial}^2 - (2\lambda^2 - 1) \sigma_w^2 \sigma_y^2 fit_t \quad (A-14)$$

Note that the two terms in Equation A-14 have a parallel structure with conceptually related components. The first elements in each term relate to λ^2 ; the second elements represent the variance in relevance weights, uncensored and censored, respectively; and the third elements represent the variance in outcomes in the partial sample and full sample, respectively. The equation shows intuitive relationships. The variance of the solo distribution is fundamentally in the units of the variance of outcomes. All else equal, it increases with higher variance of partial sample outcomes for the current prediction, and with more extreme prediction circumstances (high $info(x_t)$). All else equal, it decreases when fit rises.

In the special case of a linear regression, where $\delta(r_{it}) = 1$ for all i , we have $\lambda^2 = 1$, $\sigma_w^2 = info(x_t)$ and $\sigma_{y,partial}^2 = \sigma_y^2$, so the formula reduces to

$$\sigma_{t,solo,linear}^2 = info(x_t) \sigma_y^2 (1 - fit_t) \quad (A-15)$$

Now let us take the average of $\sigma_{t,solo,linear}^2$ across all prediction tasks t in the observed sample of N (dividing by $N - 1$ rather than N , which makes little difference):

$$\frac{1}{N-1} \sum_t \sigma_{t,solo,linear}^2 = \frac{\sigma_y^2}{N-1} \sum_t info(x_t) (1 - fit_t) \quad (A-16)$$

$$\frac{1}{N-1} \sum_t \sigma_{t,solo,linear}^2 = \sigma_y^2 (K - R^2) \quad (A-17)$$

Equation A-17 characterizes the average variance of the solo distribution. Notably, it increases with the number of predictive variables, K . When there are more predictive variables, the noise in those variables creates more dispersion in the solo predictions, reflecting greater uncertainty in the information that underlies a linear regression prediction. For a sufficiently large sample, adding a pure noise variable to linear regression will have little impact on prediction values because the noise cancels out. But this noise variable will still increase the variance of the information that informs the prediction. For RBP grid predictions, the grid assigns less weight to cells that do not exhibit useful patterns for a prediction. As a result of this prediction logic, noise variables may contribute less to information variance than they would in a linear regression context.

Let us now consider how Equation A-17 relates to the classical range of outcomes for a prediction. For the special case of just one predictive variable ($K = 1$), we have

$$\frac{1}{N-1} \sum_t \sigma_{t,solo,linear}^2 = \sigma_y^2 (1 - R^2) = \sigma_\epsilon^2 \quad (A-18)$$

Equation A-18 reveals why the solo distributions from Exhibit 1 were close to the classical distributions with a variance of σ_ϵ^2 and why one of the distributions was narrower than the classical distribution while the other was wider. On average across all empirical prediction tasks, the solo distributions will converge to a variance of σ_ϵ^2 for the special case where $K = 1$.

For comparison, it is interesting to note how the variance of a solo distribution for a linear regression prediction from Equation A-15 ($\sigma_{\hat{y}_{t,linear}}^2 = \text{info}(x_t)\sigma_y^2(1 - \text{fit}_t)$) relates to the classical variance of a linear regression prediction $\hat{y}_{t,linear}$, which can be expressed as

$$\sigma_{\hat{y}_{t,linear}}^2 = \frac{1}{N-1} \text{info}(x_t) \sigma_\varepsilon^2 \quad (\text{A-19})$$

$$\sigma_{\hat{y}_{t,linear}}^2 = \frac{1}{N-1} \text{info}(x_t) \sigma_y^2 (1 - R^2) \quad (\text{A-20})$$

The classical statistic pertains to the prediction value, rather than the distribution of information contained in the solo distribution. Therefore, the classical statistic reflects an average over $N - 1$ observations that decreases its variance by comparison. Although it does depend on the prediction circumstances $\text{info}(x_t)$, it also implicitly assumes that the reliability of every prediction task is equal by using R^2 in every case, and it only provides the statistical summary of variance for the uncertainty. By contrast, RBP's solo distributions reflect prediction-specific fit, and they reveal the entire empirical distribution rather than just summaries.

REFERENCES

- Czaronis, M., M. Kritzman, and D. Turkington. 2022a. *Prediction Revisited: The Importance of Observation*. Wiley & Sons.
- . 2022b. “Relevance.” *Journal of Investment Management*, 20 (1).
- . 2023. “Relevance-Based Prediction: A Transparent and Adaptive Alternative to Machine Learning.” *The Journal of Financial Data Science*, 5 (1): 27–46.
- . 2025a. “Prediction with Incomplete Information.” *The Journal of Financial Data Science*, 7 (4): 8–21.
- . 2025b. “A Transparent Alternative to Neural Networks with an Application to Predicting Volatility.” *Journal of Investment Management*, 23 (3).
- . 2025c. “The Virtue of Transparency: How to Maximize the Utility of Data without Overfitting.” *The Journal of Financial Data Science*, 7 (2): 42–61.
- Shannon, C. 1948. “A Mathematical Theory of Communication.” *Bell System Technical Journal* 27 (July, October): 379–423, 623–656.

Disclaimer

This material is for informational purposes only. The views expressed in this material are the views of the authors, are provided “as-is” at the time of first publication, are not intended for distribution to any person or entity in any jurisdiction where such distribution or use would be contrary to applicable law, and are not an offer or solicitation to buy or sell securities or any product. The views expressed do not necessarily represent the views of Windham Capital Management, State Street Global Markets®, or State Street Corporation® and its affiliates.