MOODY'S ANALYTICS

ANALYSIS

February 2019

Prepared by

Sohini Chowdhury
Sohini.Chowdhury@moodys.com
Director

Cristian deRitis
Cristian.deRitis@moodys.com
Senior Director

Deconstructing Scenario Weights for CECL

Introduction

Current Expected Credit Loss is the new accounting standard issued by the Financial Accounting Standards Board as Topic 326 in Accounting Standard Update 2016-13. CECL requires institutions to set aside reserves for the entire lifetime-expected losses for all eligible loans at the time that credit is originated or acquired. These estimates of lifetime losses need to be based on historical and current conditions as well as conditions expected in the future. For institutions that are considering incorporating future conditions using a probability–weighted multiple scenarios approach, the choice of scenario weights is critical. This paper presents the theoretical motivation behind these weights and suggests reasonable ways of choosing these weights in practice.

Deconstructing Scenario Weights for CECL

BY SOHINI CHOWDHURY AND CRISTIAN DERITIS

Current Expected Credit Loss is the new accounting standard issued by the Financial Accounting Standards Board as Topic 326 in Accounting Standard Update 2016-13. CECL requires institutions to set aside reserves for the entire lifetime-expected losses for all eligible loans at the time that credit is originated or acquired. These estimates of lifetime losses need to be based on historical and current conditions as well as conditions expected in the future.

While the requirement of historical and current information is clear, the CECL guidance does not mandate a specific approach for incorporating forward-looking views of the economy into loss forecasts. As a result, institutions have a wide range of options from which to choose, from simple qualitative overlays to full-blown simulations.[1] One common method being considered[2] involves the use of probability-weighted multiple forward-looking scenarios to incorporate losses from a baseline, upside and downside economy. This is also the method required by International Financial Reporting Standard 9, or IFRS 9, CECL's international counterpart. In this approach, the choices of scenarios and scenario weights are important considerations, since they can significantly influence the size, volatility and cyclicality of loss-reserve estimates. While the choice of scenarios has been discussed quite extensively in earlier reports,[3] less has been written about scenario weights. The focus of this paper is to present the theoretical motivation behind these weights and to suggest reasonable ways of choosing weights in practice.

The case for using multiple scenarios in CECL

CECL requires an institution's assessment of future credit losses to be based on its management's best forward-looking outlook given all available information. To the extent that losses are sensitive to the economy, this creates a need for estimates of future economic conditions or economic scenarios. If losses were linearly related to the economy, a single baseline "most likely" estimate of future conditions would produce an accurate estimate of the "most likely," or average, loss. But the nonlinearity of credit losses to the economy is well understood. This asymmetry is best captured by evaluating losses from multiple economic outcomes and then calculating the probability-weighted average.

In theory, using multiple scenarios should also help to mitigate some of the quarter-to-quarter volatility and cyclicality observed in losses estimated using a single scenario. This claim is based on two assumptions. The first is that no forecast of the economy, however good, is 100% accurate over a reasonable forecast horizon. Every forecast, therefore, is subject to revisions over time. That most professional forecasters did not anticipate the depth, timing and duration of the Great Recession makes this a very reasonable assumption. By considering multiple outcomes and assigning nonzero likelihoods to tail events, a multiple-scenario approach, theoretically, takes out some of the forecast uncertainty associated with a single-scenario approach.

The second assumption, of compression in economic performance at the tails, is theoretically sound but difficult to prove empirically. The idea is that the impact of a shock on the economy is correlated with the

Chart 1: Multitude of Scenarios to Represent All Outcomes

Schematic – Annualized U.S. GDP growth

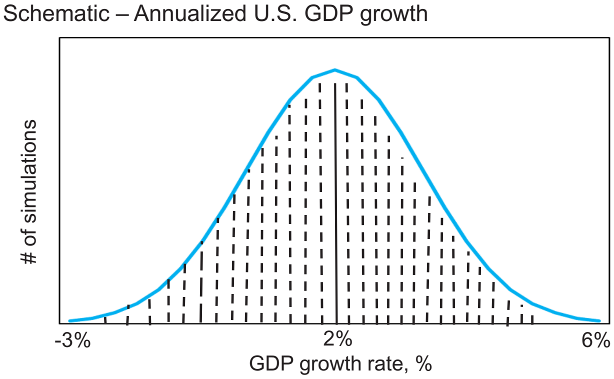

Source: Moody's Analytics

10k simulations over a 5-yr period

1 See A Practical Guide to Using Economic Forecasts for CECL Estimates

2 See the Moody's Analytics CECL Survey 2018 and page 13 of the KPMG CECL Survey 2018

3 See What's Reasonable and Supportable?

Chart 2: Only Two Scenarios to Represent All Outcomes

Schematic – Annualized U.S. GDP growth

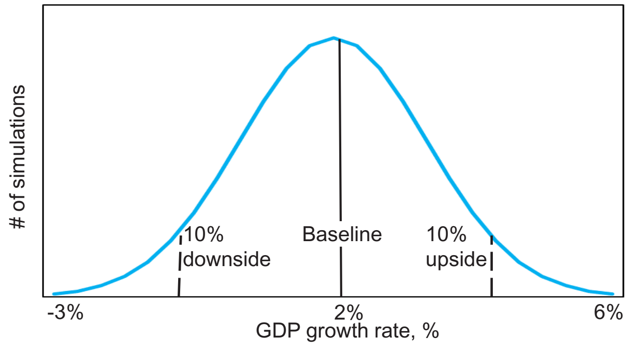

Source: Moody's Analytics

10k simulations over a 5-yr period

economy's position in the business cycle. More specifically, the farther away the economy is from its mean, the lower the impact of the shock. This is a corollary to the theory that the economy has a mean-reverting trend and therefore performance is bounded at both ends. On the upside, land, labor and capital resources put a ceiling on the economy's growth potential. For example, there is not much room for further improvement if the unemployment rate is close to its equilibrium historical low. On the downside, as economic conditions worsen, automatic stabilizers and policy responses kick in to prevent further deterioration in economic activity. In other words, there is less room to deteriorate further if the economy is already at a 12% unemployment rate versus a 5% unemployment rate. For CECL, this compression at the ends means that, over time, revisions to the upside and downside "tail" scenarios will be proportionately lower than revisions to the baseline "middle of the distribution" scenario.

In practice, whether or not multiple scenarios help to manage volatility and/or cyclicality will depend on the scenarios selected, the degree of compression at the tails, the probability weights, and the specific credit-loss models used.

Multiple scenarios as an approximation

CECL requires an estimate of the expected lifetime loss of each loan or cohort of similar loans. Theoretically, this is the average of the lifetime losses for each loan under all possible future conditions. To estimate this, an institution would simulate the paths of all the macroeconomic variables that impact its losses, and calculate the lifetime loss for each loan under each simulated path. The average of these losses would be the lenders' best estimate of its expected loss. Although theoretically sound, the approach of running simulations every month or every quarter is resource intensive and clearly not feasible for the majority of institutions.

Using a small set of outcomes/scenarios to approximate the set of all possible economic outcomes is a simpler and more workable alternative. This is the motivation behind using multiple scenarios in CECL. The probability weight assigned to any particular scenario in this approach represents the share of outcomes that are best approximated by that scenario. The probability weight is not the likelihood of that specific scenario or event occurring out of the infinite number of possible outcomes; that likelihood is zero.

As a motivating example, suppose a large number of scenarios, say 1,000, are chosen to represent the set of all possible economic outcomes. Chart 1 shows the schematic distribution of the growth rate of real GDP in the U.S. Each scenario, shown by a dashed vertical line, captures the small range of outcomes around it and is assigned a probability of 0.001. The solid vertical line represents the center of the distribution, or the baseline "most likely" growth rate of GDP.

Selecting scenario weights

In the example above, the 1,000 scenarios used to approximate the full distribution of GDP growth were weighted equally. But should the selected scenarios always be assigned equal weights? Consider an extreme case where the same distribution of GDP growth is represented by only two scenarios—the 10th percentile upside scenario and the symmetric (around the baseline) 90th percentile downside scenario. This is illustrated in Chart 2. Weighting the two outcomes equally would make sense if credit losses were a linear function of economic conditions. For example, if losses fell to half when the unemployment rate halved, and doubled when the unemployment rate doubled. However, it is well understood that this is not the case for most credit portfolios. Losses, in most instances, are disproportionately higher in a bad economy than they are lower in a good economy. In such a case, weighting the equidistant-from-the-baseline upside and downside scenarios equally will underestimate credit losses.

To see why this is the case, consider the distribution of loss rates for a typical consumer portfolio shown in Chart 3. In this particular example, the default probabilities are assumed to be functions of only GDP growth and the unemployment rate. The loss distribution (assuming zero recoveries) is obtained by estimating the losses from 1,000 simulated paths of these inputs.[4] The average and median loss rates are 8.5% and 7.3%, respectively. The convexity or asymmetry in losses is obvious from the long right tail of the distribution, especially prominent when superimposed against a normal distribution curve.

For CECL, institutions must estimate the expected loss number for every portfolio, every quarter. Some institutions are even preparing to run this exercise monthly to more closely monitor the changes in loss reserves. Under the simulation approach described above, institutions would have to simulate the full range of economic outcomes at least once every quarter and run the simulated forecasts through every loss model in their portfolio to produce the distribution of losses. This would be an extremely time-consuming and resource-intensive task, especially if done every month. A simpler and more practical alternative approach would be to use a handful of multiple scenarios to approximate the full range of simulated economic outcomes. Under this approach,

4 GDP and the unemployment rate are simulated using a vector autoregressive model described in Assigning Probabilities to Macroeconomic Alternative Scenarios

Chart 3: Long Right Tail in a Typical Loss Function

Distribution of lifetime loss rates, %, from 1,000 simulations

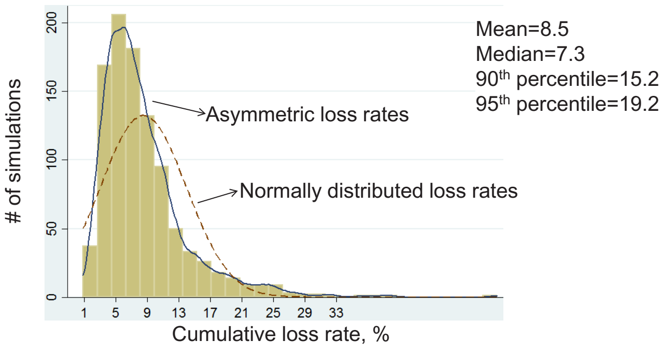

Source: Moody's Analytics

Chart 4: Discrete Probabilities - Midpoint Approach

Schematic – Annualized U.S. GDP growth, %

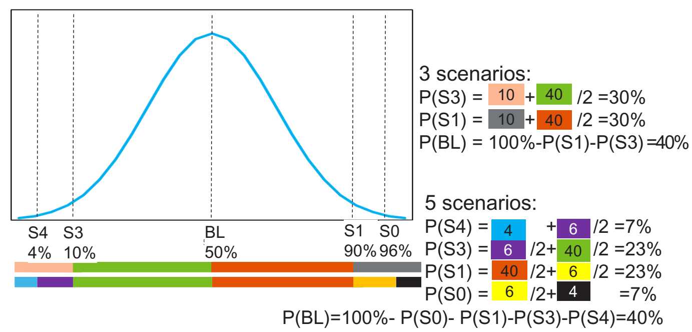

Source: Moody's Analytics

10k simulations over a 5-yr period

an institution estimates the lifetime losses from a handful of scenarios and weights these losses such that the weighted-average loss equals the estimated average theoretical lifetime loss (which is 8.5% in our example).

Suppose the Moody's Analytics baseline scenario and the 10th percentile upside and downside scenarios are chosen to approximate the set of possible economic outcomes. A rough calculation shows that in order to match the theoretical average loss rate of 8.5%, the loss rates from the baseline, upside and downside scenarios must be weighted by 20%, 40% and 40%, respectively, assuming equal weights on the two alternative scenarios. Such a small weight on the baseline is not intuitive. The reason these weights are not very reasonable is because the loss distribution is skewed heavily to the right. The disproportionate number of "large" losses means that to have the weighted-average loss match the theoretical-average loss, the downside scenario must be assigned a higher weight than the upside scenario.

Another iteration exercise shows that the following combination of asymmetric weights also produces a weighted-average loss of 8.5% and is more reasonable: 55% on the baseline scenario, 10% on the upside scenario, and 35% on the downside scenario. This is not to suggest that an institution should always use unequal weights. In this particular example involving a highly skewed loss function, unequal weights on the upside and the downside scenarios are more reasonable and supportable than equal weights. In cases where the loss functions are more symmetric, equal weights could certainly work well. Given how similar the baseline and upside scenarios are in the latest forecast vintages, we could drop the upside scenario from this exercise. With only two scenarios to approximate the range of all possible economic outcomes, a 65-35 weighting of the baseline and the downside scenario produces the same weighted-average lifetime loss; the downside scenario represents the worst 35% outcomes, and the baseline the rest. However, as economic conditions worsen, the baseline will move closer to the downside scenario and further away from the upside scenario. The upside scenario will then cease to be immaterial. It is for this reason that we recommend that institutions use a mix of the baseline, upside and downside scenarios at all times.

Scenario weights in practice

In practice, it will be difficult (and possibly overkill) to choose scenario weights to mimic the theoretical loss distribution. For one, different portfolios will have different loss distributions, and choosing portfolio-specific scenario weights will be difficult to interpret and communicate to stakeholders. Therefore, keeping in mind operational efficiency, we suggest the following practical methods for weighting scenarios. Institutions can use these weights as a starting point and modify them based on their specific views of the economy and loss models. For example, although the upside and downside scenarios receive equal weight in our recommended approaches, asymmetric weights would be more appropriate for more convex loss functions, as previously demonstrated.

Continuing with the example in which we approximate the set of all possible economic outcomes using the Moody's Analytics 50th percentile baseline and the 10th percentile upside and downside scenarios, there are at least three ways of assigning discrete probability weights to the three scenarios:

1. **Midpoint approach: 30-40-30**

This approach assumes that the 10% upside/downside scenario represents all outcomes that fall to the right/left of it and half the outcomes that lie between itself and the next adjacent scenario (which in this case is the baseline). As a result, the upside and downside scenarios each receive a weight of

Table 1: Moody's Analytics Scenarios

Scenario	Description	Percentile
S0	Very Strong Near-Term Growth	4%
S1	Stronger Near-Term Growth	10%
BL	Baseline	50%
S2	Slower Near-Term Growth	75%
S3	Moderate Recession	90%
S4	Protracted Slump	96%

Source: Moody's Analytics

Chart 5: Discrete Probabilities - Average Approach

Schematic – Annualized U.S. GDP growth, %

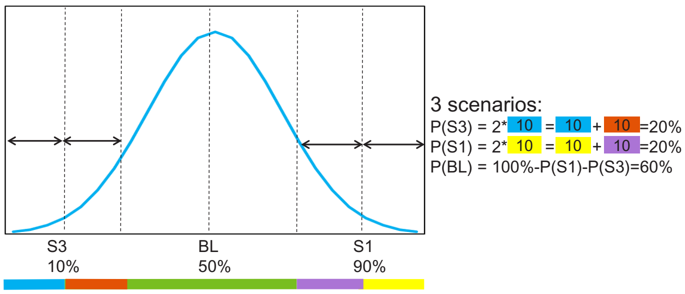

Source: Moody's Analytics 10k simulations over a 5-yr period

Chart 6: Discrete Probabilities - Percentile Approach

Schematic – Annualized U.S. GDP growth, %

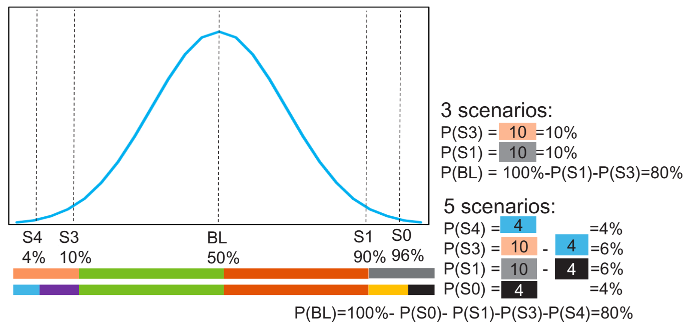

Source: Moody's Analytics 10k simulations over a 5-yr period

30%, and the baseline receives the residual 40%. The advantage of this approach is that it can be extended to any number of scenarios. See Chart 4 for a schematic representation. Table 1 lists the scenarios.

2. Average approach: 20-60-20

An alternative approach assumes that a given scenario stands in the middle of the probability interval it approximates. So, the 10% downside scenario represents all outcomes between the 0% and 20% quantiles and is assigned a weight of 20%. The 10% upside scenario also receives a 20% weight because it represents all the outcomes between the 80% and 100% quantiles. The baseline receives the residual 60% weight. Unlike the midpoint approach, the average approach breaks down when applied to more than three scenarios. See Chart 5.

3. Percentile approach: 10-80-10

Under this approach, the 10th percentile upside/downside scenario is assumed to represent only the outcomes that fall to the left/right of it. So the upside and downside scenarios each receive a weight of 10%, and the baseline receives the residual 80%. This is the most extreme view because it assigns the lowest possible reasonable weight on the tail scenarios. An even lower weight would not be reasonable. With more than three scenarios, each scenario represents the edge of its range. See Chart 6.

Dynamic scenario weights

The Moody's Analytics off-the-shelf scenarios, S0-S4, are dynamic in nature because they are anchored to the baseline and are updated every month in tandem with the baseline. If the baseline turns pessimistic, so do the upside and downside scenarios. Put differently, the specific paths of the economic indicators in the scenarios are not constant but are conditional on the economy's position in the business cycle (see Chart 7). What is constant, instead, is the relative position of the scenario in the distribution of all possible outcomes. The 10th percentile downside scenario remains so throughout the economic cycle, although the paths of the economic indicators change. For CECL, this means that the probability weights on the scenarios do not need to be updated every month to accommodate a new starting point for the economy; the monthly updates to the scenarios take care of this. Of course, an institution may update the weights if the changing environment causes management to have a different view that is not reflected in the scenarios.

The lack of specific guidance around scenario selection for CECL means that institutions have other options too. They can, for example, select scenarios with constant severities and update the corresponding probability weights every quarter. Moody's Analytics currently produces only one constant-severity scenario, where the scenario severity is independent of current business cycle conditions. In this particular scenario, the peak unemployment rate always apexes at 15% within the next eight quarters, regardless of the economy's starting position. This is shown in Chart 8. Moody's Analytics will be adding other constant-severity scenarios to its off-the-shelf scenario repertoire in 2019, covering a range of peak unemployment rates.

In conjunction with this suite of scenarios, Moody's Analytics will also publish every month the estimated likelihood of each of these scenarios. For example, the 10% peak unemployment rate constant-severity scenario will be associated with a probability that the economy reaches

Chart 7: S3 Severity Conditional on Current Conditions

Unemployment rate, %, forecast vintages of S3 scenario

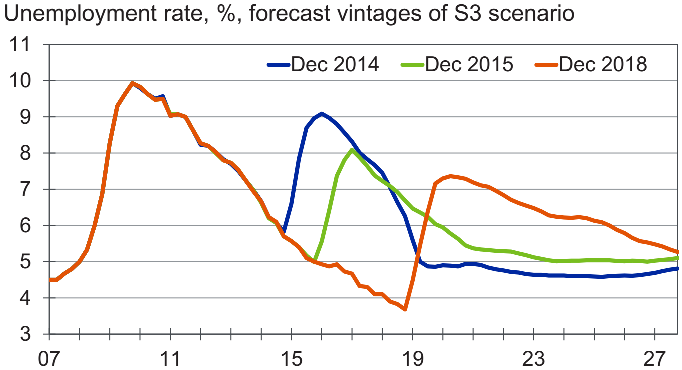

Sources: BLS, Moody's Analytics

Chart 8: Constant Severity Across Vintages

Unemployment rate, %, forecast vintages of constant severity scen

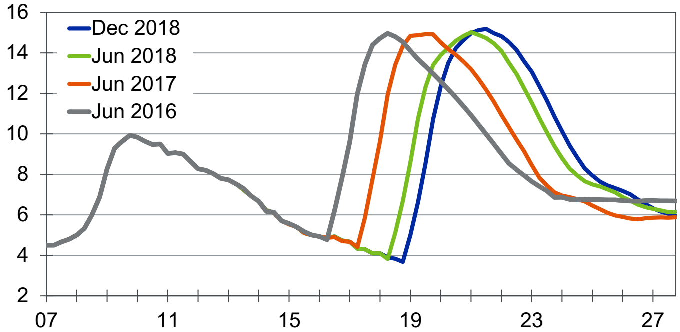

Sources: BLS, Moody's Analytics

Chart 9: Estimated Recession Odds

Probability of recession in 12 mo, %, using only…

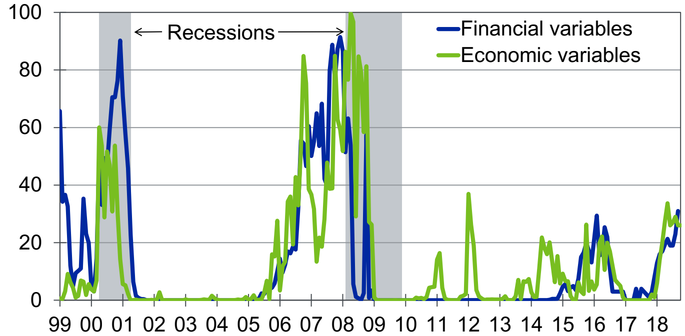

Source: Moody's Analytics

a peak unemployment rate of 10% or higher over the next two years. This probability will be updated each month to take into account the economy's current conditions; the probability will increase as current conditions worsen and vice versa. These probabilities will be estimated using a model similar to the models Moody's Analytics uses today to estimate the probability of recession over the next 12 months[5]. Chart 9 shows the recession probabilities estimated from two different models—one using only economic variables and one using only financial variables.

There are certain advantages of using constant-severity scenarios with changing probability weights in CECL compared with the more common approach of using dynamic scenarios with fixed probabilities. First, the former concept is a more intuitive one: The scenarios, or events, remain unchanged from month to month, and only the likelihood of their occurrence changes. Second, using constant-severity scenarios clamps down the upper and lower bounds of the institution's lifetime-expected losses. This, in theory, lowers quarter-to-quarter volatility and may make loss reserves less procyclical.

However, the constant-severity scenarios come with their own drawbacks. The first is that they lack a narrative. For example, the unemployment rate could peak at 10% for a number of reasons, none of which are explored. If an institution wants scenarios with narratives, it should use the existing Moody's S0-S4 scenarios, which come with monthly updated narratives. Second, the estimated scenario probabilities could suffer from noise in the data and extrapolation given that the unemployment rate has rarely reached 10% and hasn't hit 15% since the Great Depression. Finally, the fixed unemployment rate bounds are not set in stone and will need to change if the economy breaches them. For example, if the economy comes close to a 15% unemployment rate, Moody's Analytics may need to introduce a new scenario with a peak unemployment rate greater than 15%. These points only stress what most institutions understand by now: the need to regularly monitor the loss-reserve estimation process, whether using dynamic scenarios or constant scenarios.

Conclusion

There is no single acceptable combination of scenarios and scenario weights to be used in CECL. Institutions should use the principles-based nature of the guideline to their benefit to design procedures that accurately reflect their best estimate of future losses while addressing concerns of forecast accuracy and volatility. Understanding the theoretical construct behind scenario weights is a step in the right direction. Ultimately, the choice of weights will also be empirically driven, depending upon factors like the scenarios selected, the historical loss data, and the loss function. The best way to understand the impact of different scenario assumptions and weights on loss reserves is through testing and experimentation. Only by running the macroeconomic scenarios through the loss models can institutions understand the sensitivities and address any concerns.

5 Ryan Sweet, "Seeing the Odds of a Recession Through Different Lenses," Moody's Analytics Regional Financial Review, August 2016.

About the Author

Sohini Chowdhury is a director and senior economist with Moody's Analytics, specializing in macroeconomic modeling and forecasting, scenario design, and market risk research, with a special focus on stress-testing and CECL applications. Previously, she led the global team responsible for the Moody's Analytics market risk forecasts and modeling services while managing custom scenarios projects for major financial institutions worldwide. An experienced speaker, Sohini often presents at client meetings and industry conferences on macroeconomic models, scenarios and CECL solutions. Sohini holds a PhD and a master's degree in economics from Purdue University, and a master's degree in applied statistics from West Chester University in Pennsylvania. Before joining Moody's Analytics in 2011, she taught economics at the University of Cincinnati.

Cristian deRitis is a senior director at Moody's Analytics, where he leads a team of economic analysts and develops econometric models for a wide variety of clients. His regular analysis and commentary on consumer credit, policy and the broader economy appear on the firm's Economy.com web site and in other publications. He is regularly quoted in publications such as The Wall Street Journal for his views on the economy and consumer credit markets. Currently he is spearheading efforts to develop alternative sources of data to measure economic activity more accurately than traditional sources of data.

Before joining Moody's Analytics, Cristian worked for Fannie Mae and taught at Johns Hopkins University. He received his PhD in economics from Johns Hopkins University and is named on two U.S. patents for credit modeling techniques.

About Moody's Analytics

Moody's Analytics provides financial intelligence and analytical tools supporting our clients' growth, efficiency and risk management objectives. The combination of our unparalleled expertise in risk, expansive information resources, and innovative application of technology helps today's business leaders confidently navigate an evolving marketplace. We are recognized for our industry-leading solutions, comprising research, data, software and professional services, assembled to deliver a seamless customer experience. Thousands of organizations worldwide have made us their trusted partner because of our uncompromising commitment to quality, client service, and integrity.

Concise and timely economic research by Moody's Analytics supports firms and policymakers in strategic planning, product and sales forecasting, credit risk and sensitivity management, and investment research. Our economic research publications provide in-depth analysis of the global economy, including the U.S. and all of its state and metropolitan areas, all European countries and their subnational areas, Asia, and the Americas. We track and forecast economic growth and cover specialized topics such as labor markets, housing, consumer spending and credit, output and income, mortgage activity, demographics, central bank behavior, and prices. We also provide real-time monitoring of macroeconomic indicators and analysis on timely topics such as monetary policy and sovereign risk. Our clients include multinational corporations, governments at all levels, central banks, financial regulators, retailers, mutual funds, financial institutions, utilities, residential and commercial real estate firms, insurance companies, and professional investors.

Moody's Analytics added the economic forecasting firm Economy.com to its portfolio in 2005. This unit is based in West Chester PA, a suburb of Philadelphia, with offices in London, Prague and Sydney. More information is available at www.economy.com.

Moody's Analytics is a subsidiary of Moody's Corporation (NYSE: MCO). Further information is available at www.moodysanalytics.com.

DISCLAIMER: Moody's Analytics, a unit of Moody's Corporation, provides economic analysis, credit risk data and insight, as well as risk management solutions. Research authored by Moody's Analytics does not reflect the opinions of Moody's Investors Service, the credit rating agency. To avoid confusion, please use the full company name "Moody's Analytics", when citing views from Moody's Analytics.

About Moody's Corporation

Moody's Analytics is a subsidiary of Moody's Corporation (NYSE: MCO). MCO reported revenue of $4.2 billion in 2017, employs approximately 11,900 people worldwide and maintains a presence in 41 countries. Further information about Moody's Analytics is available at www.moodysanalytics.com.