CECL and IFRS 9 Modeling in MATLAB

Measuring Lifetime Expected Credit Losses

MathWorks®

Stuart Kozola, head of product management for computational finance and fintech at MathWorks, answers key questions about CECL to help finance institutions meet the new 2020 accounting standard.

Background

What is CECL?

The **new current expected credit loss** (CECL) method of accounting for loans and debt securities will be one of the biggest changes in the history of accounting for financial institutions. Created by the Financial Accounting Standards Board (FASB), CECL regulates credit loss reporting of financial institutions. It governs balance sheet accounting of loans, mortgages, and other credit instruments and primarily affects U.S. reporting. CECL incorporates a credit impairment charge on affected assets, resulting in higher provisioning expenses.

The new CECL method of accounting for loans goes into effect in 2020. The impending deadline will materially impact capital requirements and most bank operations at financial organizations. This Q&A with Stu Kozola helps prepare quantitative risk analysts and risk managers to meet these challenges with ***MATLAB***®.

What is IFRS 9?

IFRS 9 is the International Financial Reporting Standard that addresses the accounting of financial instruments, such as loans, mortgages, and other credit instruments. The directive has three core components: classification of instruments, impairment calculation, and hedge accounting.

CECL and IFRS 9 will impact financial organizations in similar but distinct ways. Organizations facing compliance for both standards will find similar requirements for data, modeling, and reporting.

Q: How does CECL change the way financial institutions measure and model loss?

A: CECL really challenges U.S. institutions in the way they measure loss. Rather than looking at history, you now need to forecast future economic conditions, and future internal conditions, and build that into credit loss measurement. In addition, you need to stretch that forecast and the forward view over the life of that asset, however long that might be. So one of the major challenges is having to layer a forecasting model instead of basing the architecture and the infrastructure on historical credit risk data, or including a very limited forward view.

Q: What are some of the changes that CECL introduces?

A: CECL introduces some new concepts, for example, there's the concept of **pooling**. How you pool similar loans or debt securities can have a very big impact on the number that you ultimately end up reporting as your CECL allowance. Also, CECL is a lifetime loss model, and it looks to the contractual life of the asset and the expected cashflows over the lifetime of the contract. You're not allowed to go beyond the contractual life of the asset when measuring your credit losses. These are some challenges and changes in the way financial firms develop models.

Q: How do IFRS 9 and CECL relate to one another?

A: The IFRS 9 standard is already in effect, while CECL is pending. IFRS 9 and CECL were originally born of an effort by the FASB and the International Accounting Standards Board (IASB) to join forces and create a converged credit loss standard. That convergence ultimately failed, and the two boards went their separate ways. But there is still a great deal of similarity between the two standards. The IFRS 9 standard, which was the predecessor of the CECL model, did not have the concept of reversion or the concept of SAB 102 built into it. Under IFRS 9, your healthy book has a 12-month probability of default model, and your unhealthy book, like CECL, has a lifetime model. CECL provides more options for modeling when compared to IFRS 9. For example, IFRS 9 requires you to do discounting no matter what, whereas CECL leaves the door open to essentially choose your methodology around discounting, choose how you discount, and choose different methods of discounting if you apply discounting to different pools or portfolios.

Q: What challenges does CECL pose for financial organizations?

A: CECL poses challenges to financial organizations to figure out how to best forecast losses out in the future. That's going to be a big thing, and these models will be scrutinized by regulators and auditors to justify the forecasting approach and the assumptions used. The CECL standard is not prescriptive in how to do this, and that's where there will be significant challenges to adopt reasonable and sound forecasting methodologies that both the organization and regulators agree on.

Q: What makes a model CECL compliant?

A: There's not a lot of guidance on what makes a model CECL compliant, so it's important to examine the criteria for determining CECL compliance.

The guidance says in order to be CECL compliant, the model needs to be relevant and include information from past events, the historical experience, whatever is happening today, and what you expect to happen in the future. Of

course, the big uncertainty is forecasting into the future. How to do that in a defensible way is really the biggest challenge finance teams face.

To be CECL compliant, organizations will need to estimate (measure) and *record immediately* all *lifetime* expected credit losses (ECL) for receivables, loans, held-to-maturity, and available-for-sale debt securities. The top 3 criteria for CECL compliance are:

- **Relevant information about past events**. In retail credit risk, there's usually enough data to create a lot of good historical data to defend your hypotheses about what the data is telling you. Corporate risk is a little harder—you might only have corporate ratings, right? What do you do with credit ratings? There are ways to model that, but if it's all you have, you don't have a lot of historical data other than relying on the credit rating to tell you what the future value of that might be.
- **Current conditions.** Current conditions are readily accessible and known. They are the current market state and portfolio the organization holds. The current conditions form the basis for the beginning of the forecast into the future.
- **Reasonable and supportable forecasts.** The definition of reasonable and supportable forecasts is a bit fuzzy. How do you come up with a macro- and even microeconomic forecast that you can defend to a statistician, economist, or maybe even the regulator or auditor who's going to be looking into this? Each may have different views or expectations. We can look at how to model that in a methodical way that can be extended to support multiple viewpoints about the future.

Q: What kinds of models can I use to support this?

A: There are many different types of modeling approaches for CECL compliance. They generally divide into two classes, depending on whether they account for lifetime adjustments directly in the modeling approach.

Class 1: Loss-Rate or Collective Evaluation Approaches

In this class are models that are typically pooled and collectively modeled using a loss-rate method. These are often the simplest to perform and to explain. **Loss-rate or collective evaluation models** include **static pools, vintage analysis,** and **provision matrix models**. In essence, you assign probabilities of loss based on historical information of similar assets—the "pool"—and then calculate the estimated loss across each pool, adjusting for expected economic situations in the future. The challenge is really justifying the probability number that you assign for the historical pool and the future economic scenario; there's not a lot of guidance in the regulations on how you justify it. They say, "There's a probability. It's up to you to defend it and actually justify how that includes forward-looking information." The plus side is that these are typically very simple models and easy to implement, so they're very easy to do. The downside is that they take a lot of documentation, a lot of effort to really document the probability estimate you put on it. How did you come up with that? What was the forecast? What are your criteria behind it? So, really, you're defending that one number, so that's the downside. These models don't build in a lot of that justification already without the document trail to support it.

Class 2: Path-Dependent Time-Based Models

Models in the second class are much more detailed. They're more complex and typically tend to be more computationally intensive, but integrate the lifetime aspect and provide the ability to understand the forecast sensitivity to different economic and portfolio parameters. With **discounted cash flow (DCF)**, you actually model the prepayments that are happening through time. So it's very detailed, and you have very detailed assumptions, but they're a lot easier to justify because you have all of that built into the model itself. The same level of detail is required when using **probability of default models, credit rating migration or transitions**, and **roll analysis**. It's a pretty standard way to justify variables and factors and how they influence the forecast to economists and even statisticians in a defensible way. **Regression** is one of the models that you'll see employed as well, for its simplicity in integrating economic factors that you want to include for forecasting purposes.

One advantage over the first class of model is that this class enables you to include credit ratings to come up with the probability of default for your different scenarios.

The other advantage is that the lifetime element is typically built into this modeling approach. Lifetime is built into discounted cash flows. You're forecasting all your cash flows out into the future. The **probability of default** and **loss-given default** are easy to move into a lifetime forecast as well. But they do tend to have a little bit more math behind them, a little more computation, so they're a little more intensive for people who don't have the software systems and the data in place to get the historical estimates and move them forward. It can be kind of daunting.

Q: How does this intersect with macroeconomic forecasting? How you can do forecasting into the future?

A: We can talk about coming up with a sustainable and defensible model in terms of forecasting out into the future. The example I would use is to model the U.S. economy with the variables shown in Figure 1. There's gross domestic product. There's consumption. There are wages, and then there's the federal fund rate. So how do you go about building a model for this? This is the question we often get. We need to build a model that can help us forecast what the economy will do in the future and that lets us look at different labor rates, wages, and similar parameters that impact the counterparties condition in our portfolio. Most economists will use a dynamic stochastic general equilibrium (DSGE) model—a very big buzzword for a very nonlinear model. This can be too complex for most needs. In the example that I'm walking through, we'll use a simple model in MATLAB. It's a linear vector autoregressive model with cointegration terms, so it's called a vector error correcting (VEC) model. But the way to think of this is that it's not very complicated and does a reasonable job forecasting. The vector stands for multiple dimensions, so I have multiple variables being estimated at the same time that have dependence on the others. It's nothing more than autoregressive over multiple dimensions with an adjustment for whether or not the terms are cointegrated.

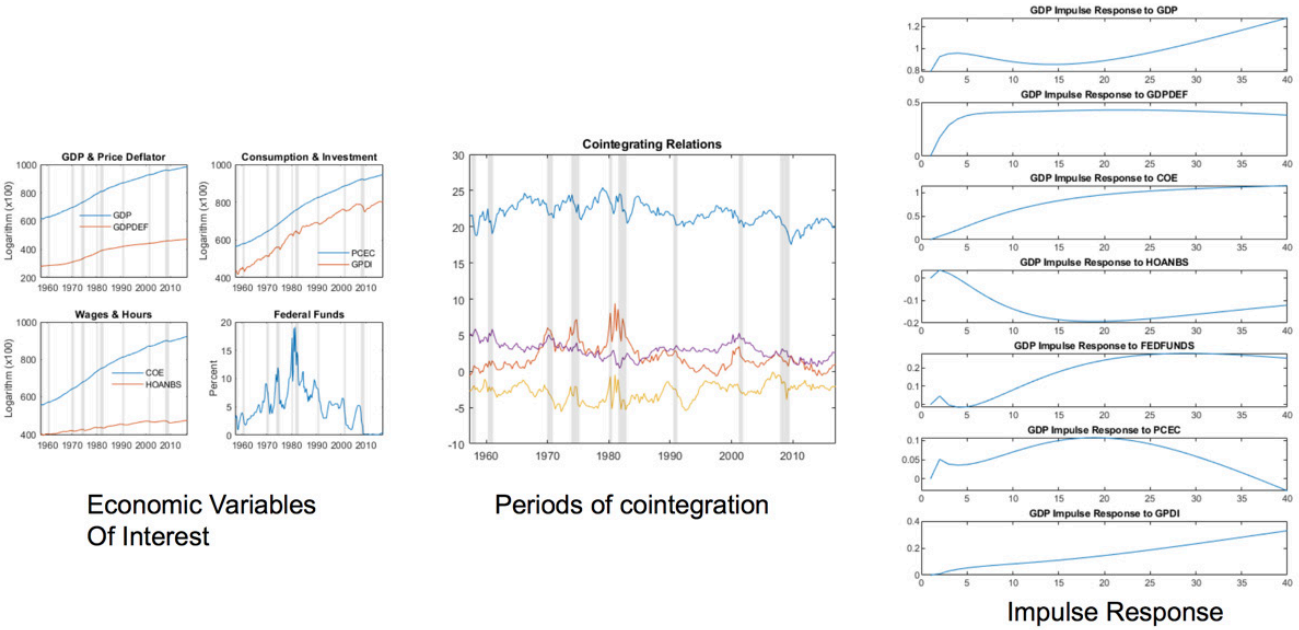

Figure 1. Reasonable and supportable forecasts from economic scenarios in MATLAB.

Q: What is cointegration?

A: When I studied the math behind cointegration, I liked the great example of a drunk and his dog. The story goes that a drunk is walking home from the bar with his dog, and the dog is on an elastic leash. So if you look at the paths, the dog is going to wander off and smell everything but can only move as far from his owner as the leash allows. The drunk is going to meander about a mean direction of travel. Both the dog and the drunk are going in the same average direction. They don't necessarily look correlated at times—the dog moves one way, the drunk the other, but they are bounded in their movements (the length of the leash). They're going different ways, and then they come back together. It's that kind of expanding and contracting like a rubber band between the two. That's what cointegration measures. It's a measure of how well they trend together with allowance for moving apart, whereas correlation is a single number that doesn't vary with time the way cointegration does. The cointegration relations shows how well the movements in different variables are bounded by this elastic band. In the figure for cointegrating relations, the movements tend to be stable over time and have similar behaviors during periods of recession (the dark bands in Figure 1) that can be used for forecasting undesirable economic conditions.

Q: How do impulse functions factor in here?

A: Impulse response functions are useful for saying, "If I want to change one variable, what does that do to the overall model output?" An impulse response is nothing more than an input-output ratio, and it represents the model dynamics to a sudden change in the variable (a spike, if you will). It gets its name from an impulse, rather than a step change in the variable. Impulse responses enable you to completely characterize the dynamics of your system. From these responses, you can learn a lot about the system, especially if you can't test the system directly. For example, is the model stable? Do things tend to co-move? What's the level of impact to the response for a unit change in an input parameter?

In this case, the impulse response tends move up to an asymptote, which tells you that it's a stable system (and will not explode on you). So my model is pretty stable. These are all things you can include in your validation of the model and the sensitivity of your model to input parameters can be easily estimated.

Figure 2 shows the forecast of the U.S. economy over time, and what we show here is the simple linear model (red line) compared with the model published by the Congressional Budget Office, or CBO (blue line). Overall, the simple model does a good job forecasting, and results like these can be used to help justify the model to reviewers or regulators.

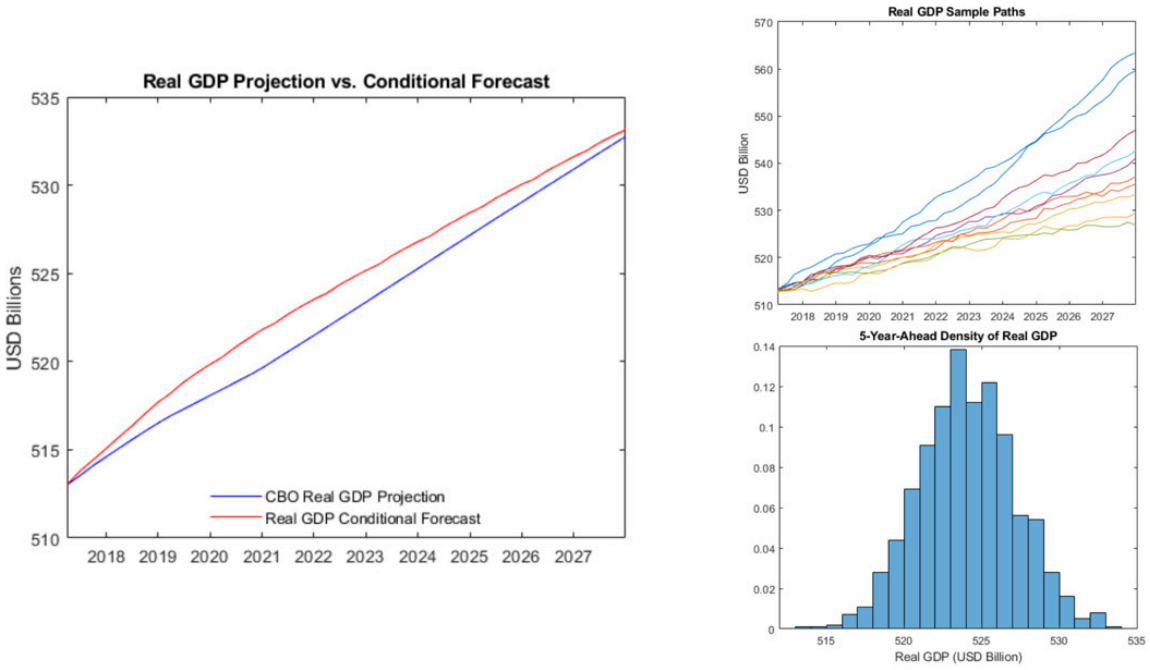

Figure 2. Modeling the U.S. economy in MATLAB.

Q: How do simple linear models compare with complex DSGE models?

A: A simple linear model is doing quite well compared with a very complex DSGE model. There is always a tradeoff in the model between simplicity and accuracy. The goal here isn't to be perfect, but to capture the trends. This simple model of the economy enables rapid sensitivity studies on parameters and the use of ***Monte Carlo simulations*** to perform what-if studies. This is particularly important for smaller organizations that don't have access to DSGE models or economists to support them in developing them.

Q: How do you come up with a probability of default model that forecasts out in the future based upon economic scenarios?

A: The whole goal is to come up with this number, expected loss, right? So, to illustrate how this works, let's talk about how to do that in the context of the **probability of default measure**. It's the one approach where you can actually use publicly published information such as credit ratings and convert them into a variety of measures for

probability of default. The equation for **expected loss** that we'll be looking at is the probability of default times loss-given default times the exposure at default. Each of these tends to be a distribution. If you multiply them, you get another distribution that you can look at and provide a confidence interval around losses for your portfolio.

Probability of default generally increases as the size of the counterparty decreases. Governments or sovereign entities would have the lowest probability of default, while small- and medium-sized enterprises (SMEs) or consumer credit would be on the higher end.

Probability of default and loss given default (LGD) are estimated externally to our portfolio holdings. They depend on the counterpart. For our portfolio, the risk exposure we'd see is driven by expected loss, and will help us to answer the question, "How much do we expect to lose?"

For those not familiar with expected loss, it's important to think about this qualitatively first. In Figure 3, the size of circle represents how much we could lose if the entity defaulted (LGD).

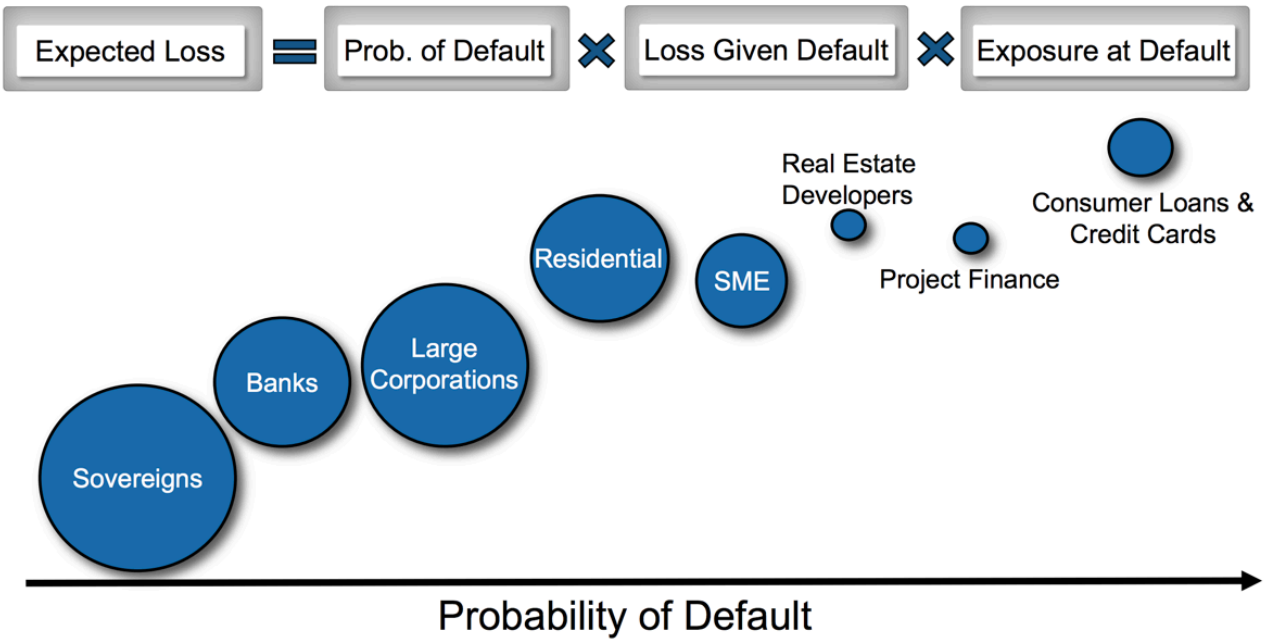

Figure 3. Quantifying loss with the probability of default model.

Our portfolio exposure to each entity isn't 100%, so we need to weight the amount by our actual exposure, which is represented by the size of the bubbles. We'd have more of our portfolio in sovereigns and less in consumer credit for this example.

This equation is going to be the basis for the CECL-compliant applications we build in our example. Furthermore, we're going to tease out some implicit assumptions, which we will relax as we generalize this approach.

So bear with me. We will answer the question about forecasting lifetime probabilities of default, but first we need to pick an example to illustrate how this can be done.

Q: What's the most difficult component of this model to estimate?

A: Some would say loss given default, which may be the most difficult to estimate with precision, but LGD is bounded and we can often be safe assuming a worst-case value here. With loss given default, you tend to know your recovery rates—what you will lose if somebody defaults—so there's generally some data you can use and apply to that. I would argue that probability of default tends to be more difficult because there are usually gaps in the data and it's bound to the counterparty's balance sheet and ability to generate cash flows for any immediate liabilities due. This change in the default probability can occur rapidly and be brought on by both firm-specific conditions (such as intense competition) and the global economy (limited credit available). A lot of people may not have enough insight into the specific firm situation for forecasting probability of default in the future.

For our purposes, we will focus on modeling the lifetime probability of default and leave modeling LGD and exposure at default (EAD) for another time.

Let's start with a retail (consumer) credit example to see how this all comes together.

Q: Can you describe the workflow process for a consumer risk model?

A: There are a few steps to the process (Figure 4):

1. **Gather the data**, and categorize it in a way that's going to make sense for the analysis you're trying to do.
2. Do credit scoring and **probability of default estimation** from the obligor data.
3. Integrate economic scenarios into the model to enable **stress testing forecasting.**
4. Project that out into the future and do the **lifetime loss estimation.**

Figure 4. Workflow for consumer credit modeling.

You're probably already doing the first three steps, for CCAR or stress-testing scenarios, which means these models and capabilities are already built into your workflow (Figure 5).

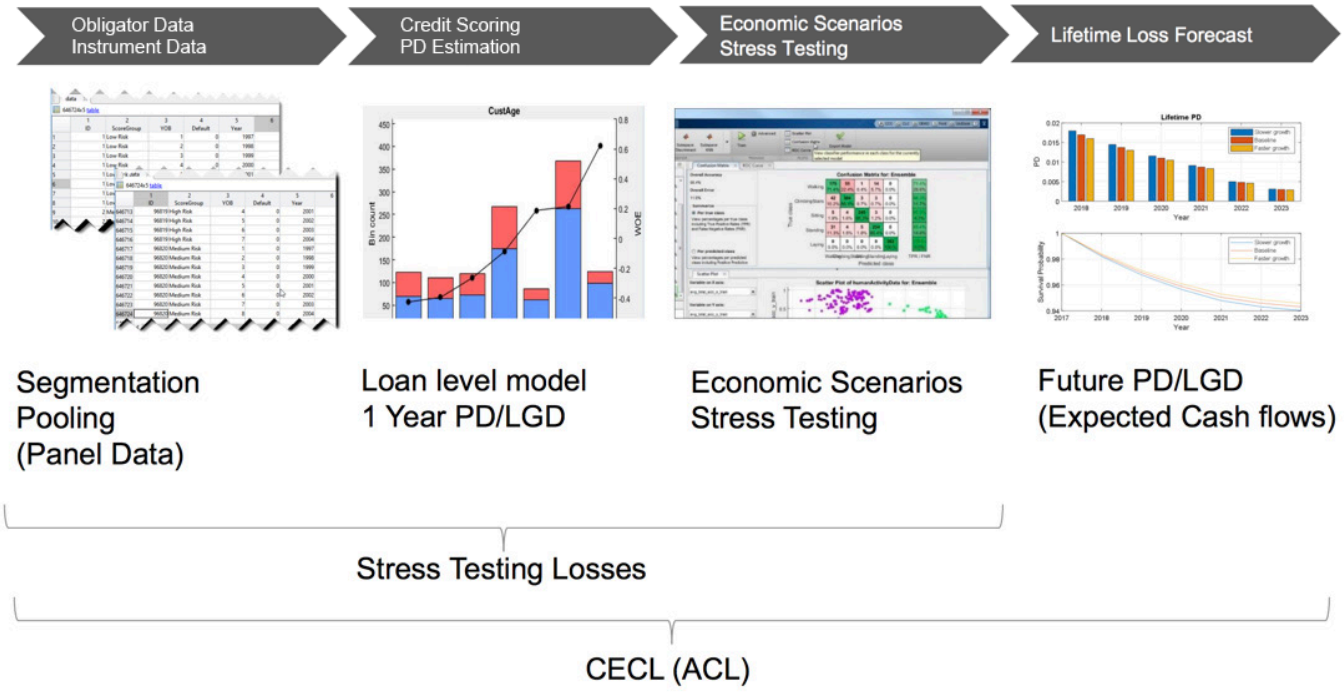

Figure 5. CECL workflow, with an additional step on top of your existing workflow.

Q: Where did the data for this example come from?

A: The data is from ***Lending Club***, where individual people can request a loan and can decide whether or not to invest in the loan portfolio Lending Club offers. This data set is rich with information about each loan requester, with information such as:

- Amount funded
- Employer
- Employment length
- Annual income
- Debt-to-income
- Two-year delinquencies
- Location

We've processed this data substantially and performed a scoring operation on it for the purpose of simplifying the explanation in the following sections.

The economic data is the U.S. gross domestic product (GDP), and the market is the total returns of the S&P 500 stock index.

Q: Are there any tricks to preparing the data set for model validation?

A: When you create a model, you first need to separate your data into two sets: a training set and a test set. You need to come up with the estimates of the parameters, and then use the test set to validate those parameters (Figure 6). If you went to a statisticians or even economists and didn't do the test set, they'd come back and say, "Well, you've over-fit the data. You've modeled everything based on history and haven't validated that your parameters are stable or have predictive power outside of your training data set." So it's good practice to separate the two, usually dividing the data set 80:20. There are other sophisticated techniques to segment it in different ways, but the standard approach is to split it with 20% reserved for testing and 80% for actual model calibration.

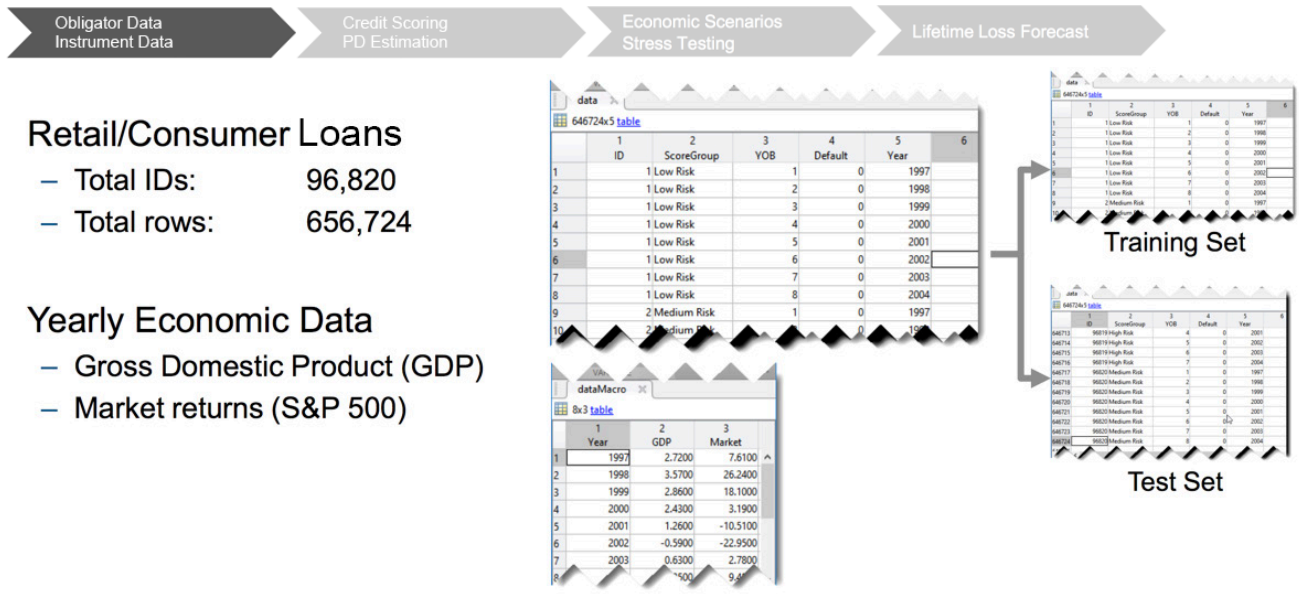

Figure 6. Data set for retail consumer loan model.

Q: How do you determine the probability of default in retail loans?

A: We're going to use the loan data and the economic data to determine the probability of default. The data is pre-scored, with a score grouping that looks at low, medium, and high in terms of a credit score. The columns show the number of years on the book, whether it's defaulted or not, and the year. In addition to that, we have two variables for our economic model, which are really the gross domestic product and the market. We'll take all this information into account to create a lifetime default model. So the process is to create a standard probability of default model tied to the economic data and then use that model to forecast out for the expected lifetime of the loan portfolio (Figure 7).

- Create Credit Scorecards
- Determine Probability of Default
- Calculate Expected Loss

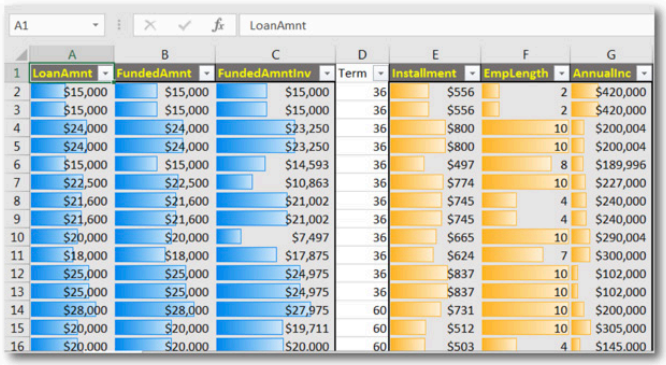

	A	B	C	D	E	F	G
1	LoanAmnt	FundedAmnt	FundedAmntInv	Term	Installment	EmpLength	AnnualInc
2	$15,000	$15,000	$15,000	36	$556	2	$420,000
3	$15,000	$15,000	$15,000	36	$556	2	$420,000
4	$24,000	$24,000	$23,250	36	$800	10	$200,004
5	$24,000	$24,000	$23,250	36	$800	10	$200,004
6	$15,000	$15,000	$14,593	36	$497	8	$189,996
7	$22,500	$22,500	$10,863	36	$774	10	$227,000
8	$21,600	$21,600	$21,002	36	$745	4	$240,000
9	$21,600	$21,600	$21,002	36	$745	4	$240,000
10	$20,000	$20,000	$7,497	36	$665	10	$290,004
11	$18,000	$18,000	$17,875	36	$624	7	$300,000
12	$25,000	$25,000	$24,975	36	$837	10	$102,000
13	$25,000	$25,000	$24,975	36	$837	10	$102,000
14	$28,000	$28,000	$27,975	60	$731	10	$200,000
15	$20,000	$20,000	$19,711	60	$512	10	$305,000
16	$20,000	$20,000	$20,000	60	$503	4	$145,000

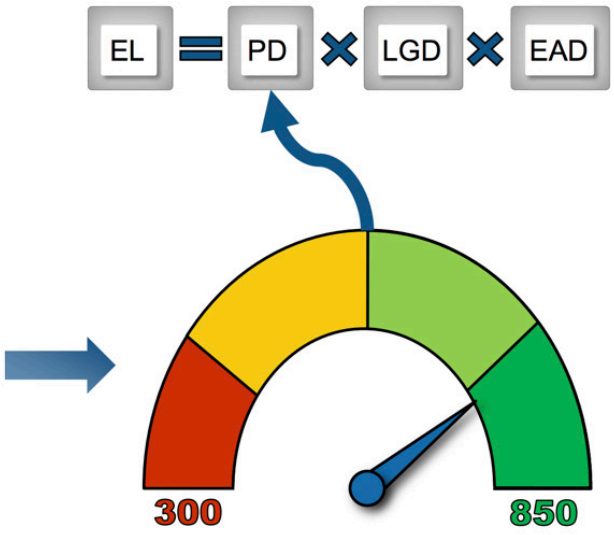

Figure 7. Modeling consumer credit risk.

Q: Can you show how you model the probability of default?

A: On the test set, we come up with an empirical cumulative distribution function, shown in Figure 8. This is for all the data collectively. From the cumulative distribution, you can get the probability of default. You can also do the same analysis on the data that's segmented by the credit scoring group, and you'll get a plot that looks like the one on the right in Figure 8. So this is the actual data segmented by the pools if you want to look at the different buckets that you have.

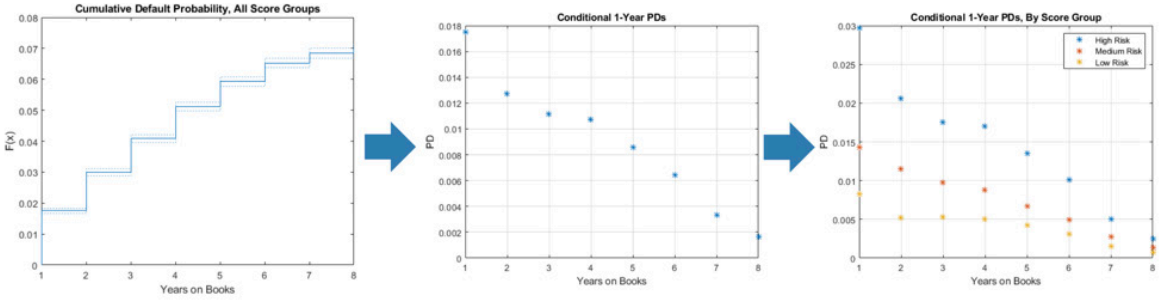

Figure 8. Probability of default estimated from the data.

Next, we're going to use the Cox proportional hazard or Cox proportional regression model, because it has some nice characteristics in terms of being able to project into the future. You can just chain the models together and project into the future using this Cox proportional regression. Other models would work just as well but may have a few more steps involved. So, in the top of Figure 9, you see the validation on the training set for the overall portfolio. The bottom shows the model by score group. Visual inspection shows that the model has some regions where it doesn't capture the data well. Let's expand the model to include the economic data and see if that fits the data better.

- PD Modeling
  - Cox-proportional regression
  - 1 Yr PD (total)
  - 1 Yr PD by ScoreGroup

$$h(X_i,t) = h_0(t)\exp\left(\sum_{j=1}^{p} x_{ij}b_j\right)$$

where:

- $X_i = (x_{i1},\ldots,x_{ip})$ are the predictor variables for the ith subject.
- b_j is the coefficient of the jth predictor variable.
- $h(X_i,t)$ is the hazard rate at time t for X_i.
- $h_0(t)$ is the baseline hazard rate function.

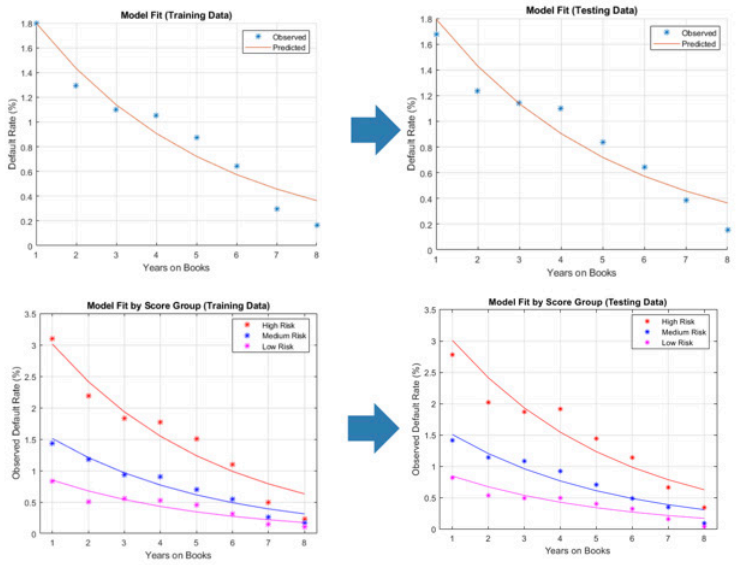

Figure 9. Modeling the probability of default using Cox proportional regression.

Q: How do you include economic factors in your probability of default model?

A: If you look at the math, you just add some terms for your macroeconomic variables into the Cox proportional regression model; then you can capture the gross domestic product as well as the returns of the market into this model. Figure 10 shows the results. We get a better fit to the model, and intuitively this makes sense. When the market is up, defaults are down. When the market is down, defaults are up. When gross domestic product is increasing, you're not in a recession. When gross domestic product is decreasing or not growing as fast, you're closer to a recession. So, intuitively, it fits the model as well, so you can see how well it fits for both the training set and the test set.

- Macro modeling
  - GDP
  - Market returns

$$h(X_i,t) = h_0(t)\exp\left(\sum_{j=1}^{p1} x_{ij}b_j + \sum_{k=1}^{p2} x_{ik}(t)c_k\right)$$

where:

- x_{ij} is the predictor variable value for the ith subject and the jth time-independent predictor.
- $x_{ik}(t)$ is the predictor variable value for the ith subject and the kth time-dependent predictor at time t.
- b_j is the coefficient of the jth time-independent predictor variable.
- c_k is the coefficient of the kth time-dependent predictor variable.
- $h(X_i(t),t)$ is the hazard rate at time t for $X_i(t)$.
- $h_0(t)$ is the baseline hazard rate function.

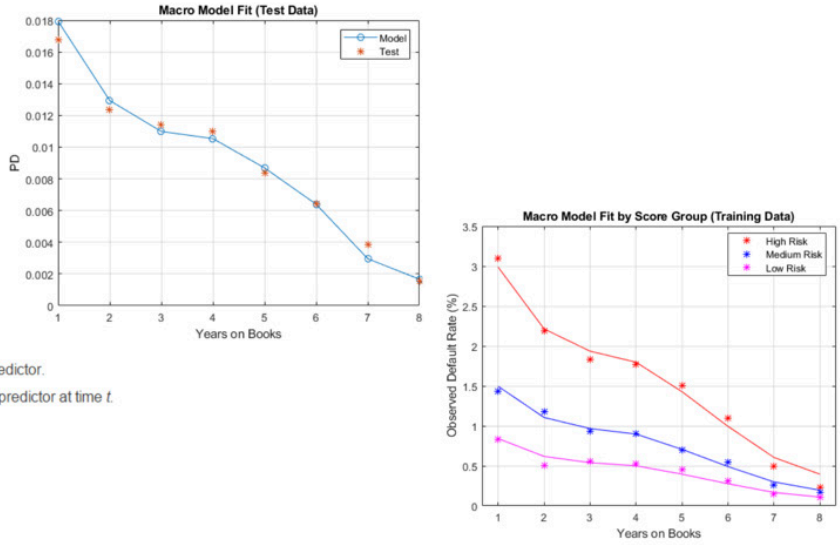

Figure 10. Getting a better fit to the model using macroeconomic modeling.

Q: How do you integrate this with stress tests?

A: Most likely today you're already stress testing a model. You probably do at least three scenarios: a baseline scenario, an extreme scenario, and a less extreme but adverse scenario. In this example, you define these scenarios around your macroeconomic variables, GDP, and market returns. Then you perform a stress test for the probability of default forecast for each case. This approach lets you test for extreme movements in future conditions. However, CECL does not require worst-case forecasts, just reasonably expected future outcomes.

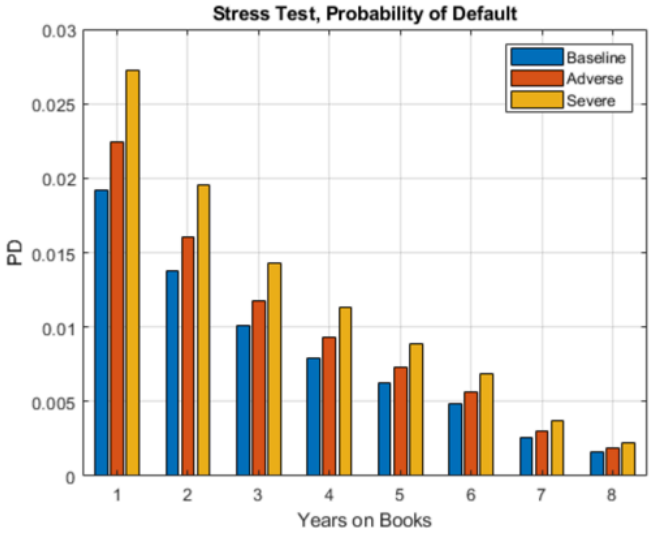

Figure 11. Stress testing for changing economic factors.

Q: How do you forecast the survivability of a portfolio for CECL?

A: Now, the trick is to take this model and then forecast it into the future to account for all the different scenarios. You chain the yearly Cox proportional regression models together, and have a forecast into the future, or the lifetime probability of default. The Cox model is really the survivability of the portfolio, and you subtract one from it to get the probability of default. Figure 12 illustrates how to do this.

$$
\begin{aligned}
&S(0) = 1,\\
&S(1) = (1 - h(1)),\\
&\dots\\
&S(t) = S(t-1)(1 - h(t)) = (1 - h(1)) \cdots (1 - h(t))
\end{aligned}
$$

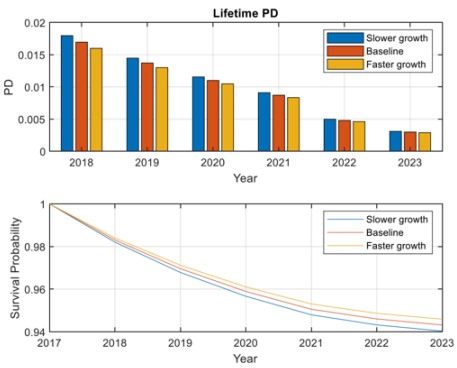

Figure 12. Creating the lifetime probability of default for the Cox proportional hazard model.

In this example, we've changed our scenarios not to be a stress test, but to forecast out into the future for a better-than-expected and worse-than-expected economic situation (+/- 10–20% from the baseline). So you can see how the probability of default changes over time as well as the survivability of your loan portfolio. At the beginning, it's 1, and then it decays down to about 0.94.

We've created a justifiable lifetime probability of default model tied to economic variables that can be used for CECL compliance. The only remaining part is to calculate the expected loss of your portfolio, which is relatively straightforward. From your portfolio, you multiply the value of the loans outstanding at default (the LGD) and your exposure to get the expected loss, which is what is recorded as your CECL impairment.

Q: From a technology perspective, what are the top challenges to implementing CECL?

A: What I see among our customers is that once they figure out that they've had data challenges or controls issues—and they actually try to implement models and systems—a big challenge for them is trying to do it in the timeframe they need to or want to. It comes down to automating all of these processes. Finance organizations have a lot of disparate legacy data systems, even data modeling systems, that they've got to bring together to do all these parallel runs and get reporting out through automated report generation. We spend the bulk of our time with customers helping them plug all those pieces together and building an automation platform to help them move through the process quickly. Data collection and automation are the most time-consuming and error-prone parts of the process. The modeling is relatively easy once the data is available.

Q: What should banks be looking for in a solution?

A: We see a lot of banks and firms considering developing a model execution platform that fits within their IT infrastructure and enables rapid evolution of their models while allowing model governance.

We have gone through DFAST and CCAR, but through all those years very few banks have implemented or built really good, controlled environments where they can run their credit loss models and update them in an agile fashion. Most banks are still running by duct tape and baling wire. And when you think about how long it takes to run CCAR now, CECL isn't as complicated because it doesn't have the balance projections, but it will introduce a different complication because the current systems are built for lifetime estimation.

But still, you're going to be running a somewhat complex system on a monthly basis, and you're going to report on it on a quarterly basis. So the time crunch and the compliance are enormous, and there are very few banks that could do it at the moment. So there's a lot of work on accelerating data collection through a model execution platform that needs to be implemented.

What banks need is a technology platform that can be managed using standard IT processes but is flexible enough to support agile model development, while also supporting the disruptive move toward cloud computing architectures. This calls for a model development environment that caters to the risk manger's way of thinking (the math) and the ability to take the risk manager's model into production without having to recode or redevelop that model from scratch, which can introduce translation errors or integration problems. Also, IT wants to follow an agile development process. MATLAB has been designed from the ground up to bring mathematically focused solutions into IT architectures and processes.

Q: Do you have any project management tips for teams implementing CECL?

A: With these various complex issues that we've discussed from data, decision-making, controls, managing the timeline, getting automation in place, getting model execution set up, what does it mean for companies as they try to implement this from the actual organization program and project management perspective? Once you've decided you're going to do this, how do you actually run the organization to do it effectively, and educate and change the DNA of how your firm works?

Because of how cross-functional CECL is, from a project management perspective, you need to bring together a range of in-house experts—including accounting people, risk people, business owners, and quantitative modelers. In my experience dealing with large financial institutions, very often those departments may not even talk to each other prior to CECL forcing them together.

In terms of best practice, either finance or accounting and credit risk and modeling need to co-lead the project. It definitely needs a dedicated point person, certainly for midsize and large institutions to run with this. It's nearly their full-time job. Many institutions try to wing it to a certain extent, with people supplementing their existing responsibilities, and that seems to be a recipe for failure. A major area of weakness is the lack of good, robust project management documenting who owns or who is responsible for what pieces of the puzzle.

When you consider how much optionality there is within CECL, someone actually needs to sit down and make decisions and own the process. Very often nobody really knows who owns that decision-making process, or it's not clear whether it's accounting's responsibility because it's an accounting standard, or whether it's the modeler's responsibility because they're building the model, or whether it's credit risk's responsibility because they're ultimately responsible for governance of that whole process. Clear ownership of the different work streams and who's actually responsible for making decisions is a key best practice. Not defining that has been an area of weakness, where others have not been very good at allocating responsibility, essentially, for the different strands within CECL.

Q: What should the role of executive sponsorship for CECL projects be?

A: There's no question you also need executive sponsorship and leadership for a successful CECL implementation. That leader needs to be engaged, involved, and willing to be an arbitrator from time to time, because you're going to have different groups, using different technologies, all coming together to try to solve something. Somebody needs to be that mediator and also be the person who says, "Well, if we need more resources, we'll go get them." So it really has to come from the top down that they're behind this, they're committed to it, and they're going to do it. They also need to show that they're involved. Once you get executive sponsorship involved, and the other teams are all onboard, then they all come together behind that project and go toward the common goal.

Q: Are there any other implementation pitfalls you've seen?

A: One of the problems is when teams get way too big. Amazon has this principle about two-pizza teams. Two-pizza teams are purposely small and can get things done and get them done quickly. When teams get too big, you get a lot of fiefdoms developing, with a lot of different competing initiatives. So that's the other thing to watch out for—not letting the teams that own it get collectively too big. Maybe overall they're big, but each smaller team owns a small piece with clear deliverables and interfaces to the other teams. Working meetings should never be larger than this. Bigger teams can handle information sharing and reporting but never decision making.

Q: How does this extend to smaller institutions, in the less-than-a-trillion, less-than-10-billion range?

A: I think that's where you get some of the most varied reactions to this, and some are really on top of that. But they also are leveraging less infrastructure, so they probably have less historical data. They probably have fewer people who have been thinking about these things and have been responsible for these things in the past, and they're not leveraging a CCAR framework, for instance. So they may be building their first real enterprise-strength models, not leveraging one, adapting it, and looking for a challenger to it. But this is also the group that was most hopeful that this would go away or that for the smaller banks it would be deferred. So I think there's a lot more catch-up that's being played among smaller institutions.

We're seeing something interesting with some of the smaller firms: they are focusing on probabilities of default as their preferred approach. They may have smaller C&I portfolios, so pooling is more questionable, whether they've got a wide enough array of assets in a pool to make it in and of itself robust. So they are doing case-by-case probabilities of default and currently relying on a one-year PD as opposed to extending it out for the life of the loan. I think that will come back and bite some of them, or they're going to have to, prior to filing, actually adapt and extend those PDs out. But from a modeling perspective, that's a hole that we're seeing some of them falling into.

What I've seen from a community banking perspective is the view that CECL is scary. Community banks frequently look at CECL and say, "Well, let's just hire a vendor or a consultant and let them do all of the work," which from a consultant and a vendor perspective should sound great, but that's a recipe for disaster. Because there's so much choice and optionality within CECL—and ultimately the institutions that are applying CECL need to be able to own those decisions and defend them and be able to understand them—outsourcing all that decision-making responsibility is essentially relinquishing all ownership of that process. I think there are probably very few consultants or vendors that would be comfortable making a lot of those core decisions on behalf of the community banks. It's just a very uncomfortable place to be, especially when the management and the steering committees within these banks need to be owning that process and the decision-making, even if the ultimate legwork and heavy lifting is done by an external third party. So I think that's a pain point in dealing with the smaller institutions: Smaller community banks expect that a consultant will solve all their problems, and my experience is that that's an expectation gap that needs to go away sooner rather than later.

Q: Are banks the only financial organizations affected by CECL?

A: We've spoken a lot about banks because I think they're, obviously, the ones most critically hit by the new standard, but CECL is not a banking standard. It affects many different organizations: any organization that applies U.S. GAAP or that has anything on the balance sheet at amortized cost, any company that simply has trade receivables, etc. All of that is within the scope of CECL. So, if we think about our banks, at least larger or small, they have some metrics or some element or some ability to be able to measure credit risk. But non-banking organizations don't have that, and will need to figure out how to fill the gap before auditors hit them with noncompliance.

Q: Are there lessons learned from IFRS 9 that teams can apply here?

A: Some large corporate banks engaged with IFRS 9 extremely late. And some of them are learning—to their detriment at this point, given that we're now live with IFRS 9—that they got it wrong, that they have significant exposure to IFRS 9 and don't have the expertise in house or the infrastructure to be able to deal with it. I think we'll see

probably in the next year or two, once IFRS 9 and financial statements start coming out, that a lot of those large corporate banks will really be scrambling to fix that hole on their books. But the pain IFRS 9 introduced will help elevate the urgency to address CECL, and these teams will not make the same mistakes with CECL.

Q: Do you have any advice about outsourcing CECL compliance?

A: It's important to remember that while you can outsource elements of your risk management process, you can't outsource risk management. And I think that's a pretty good adage to think of in this context as well. You can use outside consultants. You can use outside vendors, data providers, and technology providers for aspects of what you do, but you still have to own it yourselves. You can't outsource it completely.

Q: Does CECL create unintended consequences or any perverse incentives or anything like that that would cause businesses to exit certain products or categories? What are you seeing in terms of fallout from changes?

A: Well, I haven't seen any fallout yet because I think people are still scrambling to figure out how they're going to put it into their books. But any time you add a line item that's taking up capital, it's reducing the funds that you can use to fund their business, or it's increasing the cost of getting capital. Just like anything else, you're going to see companies shift their product strategy to the more profitable lines of business—maybe away from credit and loans, or some other kinds of debt instruments. I imagine you're going to see some change in how these firms do it. Some might divest some of their more capital-intensive credit operations and spin them off, and put them under different divisions, or maybe not do them anymore.

We saw this with the fallout from the recession with proprietary trading, because it became unprofitable. Certain types of credit risk operations will no longer meet the profitability criteria for some types of firms, and you will see them drop some products. The other thing that might be interesting, is any time you add a capital change that's going to involve auditors viewing it, you might get auditor scrutiny of your model that you never had before, and that could actually have more implications. The auditors will be getting into your risk management processes where they weren't typically going before. So now you're going to have oversight from somebody else, and I don't think we've figured out what that's going to mean yet to these corporations and these institutions that have to report to their auditors.

Q: The life insurance industry has embraced this concept or been associated with this concept for about 25 years, and they establish reserves for credit-related losses. How does this compare with that? How well are life insurers prepared for this?

A: From an insurance perspective, they're subject to a different regulatory environment than the banks that are obviously wrestling with CECL right now. They don't have the benefit of learning from the IFRS 9 experience because the international banks got a deferment, or the insurers got a deferment for IFRS 9 while the new insurance standard under IFRS was being finalized.

There are core differences, though, and maybe the result will actually be counterintuitive. The banks would logically expect that their moving to a lifetime model will mean higher reserves than what they have today. The life insurance

business may take a more conservative approach, where, if you compare it with CECL, CECL is not meant to be a conservative model. It's meant to be your best estimate, so you may have a potential unexpected result for insurers where they have to unwind reserves in order to be able to maintain compliance with the accounting standards, which would be an uncomfortable place for insurers to be because they've built up those reserves around those products with a view that that's the correct approach, the prudent approach, and that's the way they absorb potential losses on those products. As insurance companies start to wrestle with CECL, there will be an interesting conversation around whether insurers today are over-reserved compared with what CECL will ultimately require.

Q: What is the one thing organizations should know about CECL today?

A: You should be preparing for CECL now, and not wait for the final ruling. CECL requires more data to develop PD and LGD models from historical data, and requires the development of economic models taking into account the credit portfolios locality. If you don't have this work under way, it will be a scramble to get it in place when needed.

Meet Stuart Kozola

Stuart Kozola leads product management for computational finance and fintech at MathWorks. He has over 15 years of experience in data analytics, quantitative finance and risk management, simulation, and designing and implementing large-scale computational systems. Stu is a sustaining member of PRMIA and holds the FRM designation from GARP.

Learn More

CECL with MATLAB - Overview

IFRS 9 with MATLAB - Overview

Modeling Lifetime Probabilities of Default with MATLAB - Example

MATLAB for Quantitative Finance and Risk Management - Overview

Federal Reserve FAQ on CECL - FAQ

© 2018 The MathWorks, Inc. MATLAB and Simulink are registered trademarks of The MathWorks, Inc. See mathworks.com/trademarks for a list of additional trademarks.
Other product or brand names may be trademarks or registered trademarks of their respective holders.