Terminal Risk Analysis

A forecasting and stress testing tool for assessing portfolio outcomes and their impact on loss reserves and capital

A White Paper from Jeff Judy & Associates

Table of Contents

Terminal Risk Analysis

What if you knew the risk rating of every dollar in your loan portfolio right before it left for good? With only two ways out, repayment and charge-off, couldn't you make a pretty good guess about how each dollar would leave? That's how "Terminal Risk Analysis" (TRA) works. After projecting the terminal risk rating of all the dollars in a loan portfolio, it estimates the split between those two final outcomes.

Why does it matter? As lenders, banks have frequently struggled to accurately discern this split. And that matters to a lot of people — depositors, regulators, directors, shareowners and employees — who are counting on banks to have a good enough understanding of the final outcomes to build an enduring business model around them. Given the low levels of capital that banks have traditionally operated with, there has never been a lot of room for error … and the tolerance for error seems to be shrinking.

Fear of charge-offs drives the two ongoing questions investors, directors and regulators ask about banks. Is the ALLL enough to protect against expected loan losses? Is there enough capital to protect against the losses no one is expecting? Without a good, evidence-based way to estimate either the expected or unexpected loss, you're left with back-of-the-envelope guesses. Wouldn't a face-to-face with directors and regulators go better if you had sound data to support your assessment of loss reserve and capital adequacy?

Community banks will be hard pressed to develop the data and requisite skills to analyze the complex dynamics of their lending behavior in a manner that would put them on a par with large international banks in terms of formal credit portfolio assessment. Most likely, it won't even be expected of them. However, simply piling on more loan loss reserves and capital "just to be sure" isn't much of an answer, either. It reinforces the notion that banks lack "self-knowledge" about their risks and leaves regulators little choice but to impose ever more burdensome capital constraints. Is there something in between that offers enough analytical sophistication for a bank to be seen differently from its peers?

The question becomes, what can you do to preserve a rational degree of freedom to pursue the best business model for your bank? Can you find a sweet spot that balances a clear understanding of credit risk against ALLL and capital adequacy demands? TRA offers a way to address this intransigent question using an objective, outcome based, forward looking methodology that relies primarily on data you most likely have already. It also offers inherent stress testing capabilities and naturally leads to more formal approaches to credit portfolio assessment.

A Starting Point

Consider this example. Assume you have just four risk rating categories and your current portfolio is distributed among them as shown in Figure 1. Bear in mind, you don't necessarily know when any of the loans originated, what their original balances were or what their original risk ratings were. All you know and care about is their current status.

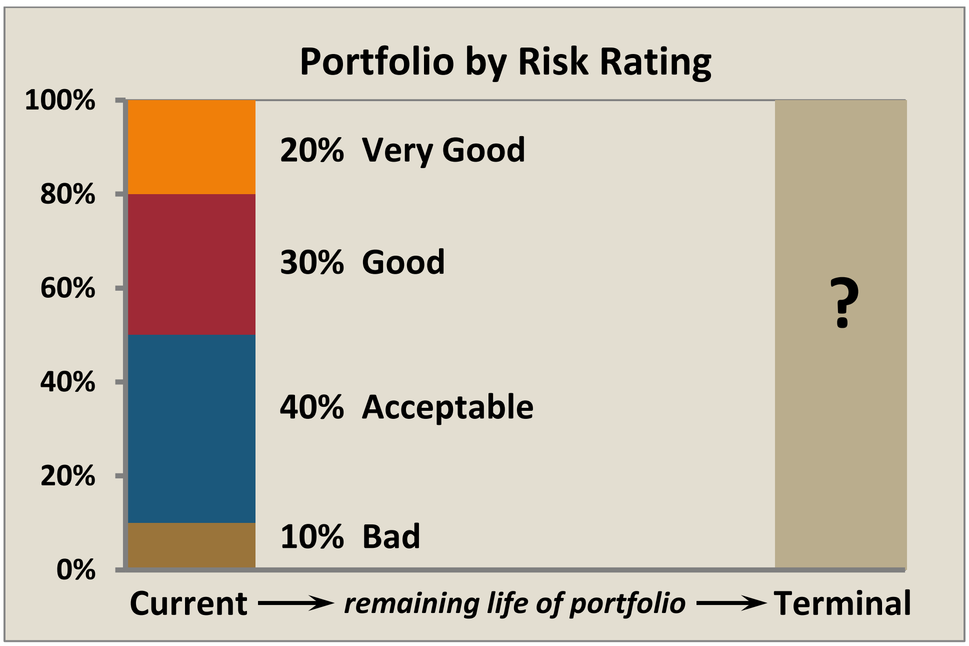

Figure 1

Although this gives you a picture of where things stand today, it doesn't tell you how they'll end up. Over the remaining life of the portfolio, you know risk ratings will shift and all the outstanding dollars will gradually take their leave, mostly by repayment but some by charge-off. How do you replace the question mark in Figure 1 short of waiting umpteen years for the portfolio to play out ... when it won't matter anymore?

Terminal Risk

Figure 2 shows a TRA projection of the ***terminal risk*** distribution for this portfolio. Without getting into the "how" just yet, consider what you might surmise from it. Regardless of its current risk rating, any loan dollar projected to leave the portfolio through one of the better risk ratings will almost certainly be repaid. Whatever the previous meanderings of this 84% of the portfolio, it's no longer a worry.

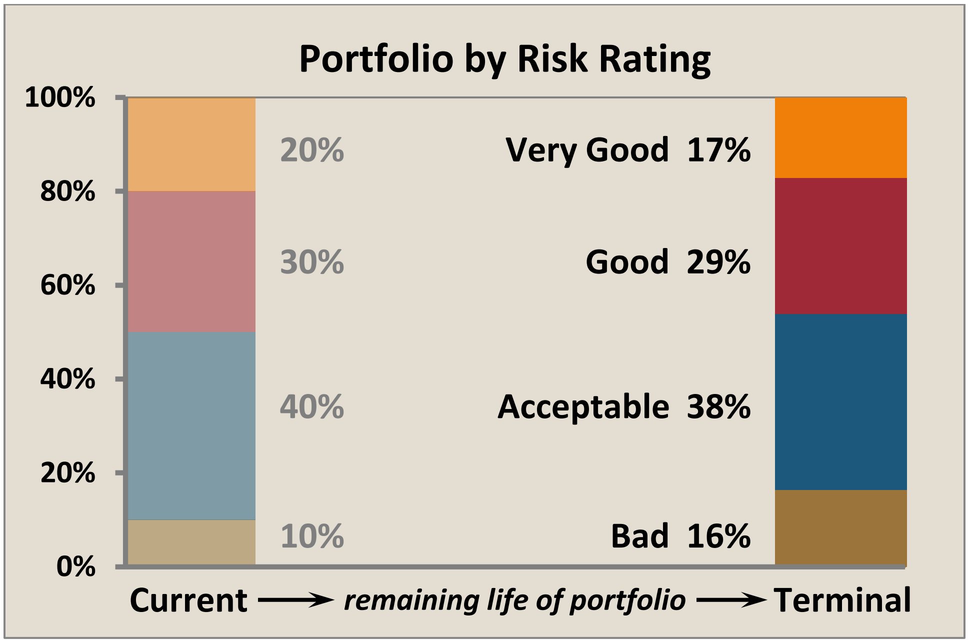

Figure 2

A Good Picture of "Bad"

In contrast, any loan dollar expected to leave via the "Bad" risk rating will bear some likelihood of charge-off. And it's safe to say virtually every charge-off will be taken against a loan that, by time of write down, is risk rated "Bad." So your real concern is focused on just 16% of the portfolio. And chances are your data show fairly clearly how you've fared with "Bad" loans, what the split between repayment and charge-off has been. You're now far beyond guessing about the nature of the inherent loss in your loan portfolio.

However, just knowing 16% of the dollars are likely to exit as "Bad" doesn't quite complete the picture. Remember, you're trying to assess the ***current*** portfolio. Only 10% of it is classified as "Bad." It's obvious some of the terminal 16% "Bad" are currently masquerading in other risk ratings. Where?

When Loans Leave

Again without explaining how, Figure 3 shows the complete detail of the terminal risk summarized in Figure 2. The "Risk Rating at Reduction" ***is*** the terminal risk rating for any loan dollar. As you can see, over the remaining life of the portfolio, even some dollars that are currently classified as "Very Good" are projected to leave as "Bad." And not all of the current "Bad" dollars are projected to leave that way. (Notice the percentages for each risk pool sum across to 100%. Also, understand that the percentages in the "Total" row are in terms of the total portfolio, which takes into account the relative size of each risk pool.)

Life of Portfolio		**Risk Rating at Reduction**				
		Very Good	Good	Acceptable	Bad	Total
Current Risk Rating	Very Good	66%	26%	7%	1%	100%
	Good	11%	68%	18%	3%	100%
	Acceptable	1%	8%	76%	15%	100%
	Bad	0%	1%	8%	91%	100%
	Total	*17%*	*38%*	*29%*	*16%*	*100%*

Figure 3

Got Data?

Having jumped ahead to some of TRA's conclusions, let's back up a bit to the "how." What information do you need? TRA requires loan-by-loan, period-by-period risk rating and loan balance data along with any selection criteria — loan type, organizational unit, etc. — you might use to refine risk pools. Clearly, risk rating reliability affects the quality of the results. And the more history you have, the richer the analytical possibilities.

You shouldn't be overly discouraged, though, by a short data set. You can still get some indication of what the information is trying to tell you. And by the time you begin to get comfortable with the methodology, you'll have even more data. In short, keep saving or start saving data … even if you don't yet know what it can do for you.

A Slice of Time

How do you do TRA? It helps to begin with the notion of a ***time slice***. In TRA, you examine a slice of time that has already passed, one with known outcomes, to determine the pattern of risk rating shifts and reductions in amount. You then simulate the continued play out of that behavior to its logical conclusion — the complete reduction of the portfolio to zero. With enough data, you can analyze ***multiple time slices*** over varying portfolio and economic conditions, resulting in even better estimates of the most likely outcome within a probable range of outcomes.

Consider another example. Say there's a car going down the road, and you're pretty sure you know where it's going and how many miles it takes to get there (Figure 4). For fun, you've decided to see how close you can come to estimating when the car will arrive at its destination. (A friend there will let you know how close you come.) You have a few buddies standing alongside the road with you. You're out in the countryside without any GPS or other fancy electronic tools. All you have are a tape measure and stop watch.

Ask one person to stand ready to signal a timer positioned further down the road. Measure the distance between the two. As the car moves along, the first person signals the other to time how long it takes the car to travel between them. With that single time slice of information, you can make a fair estimate of when the car will arrive at its destination.

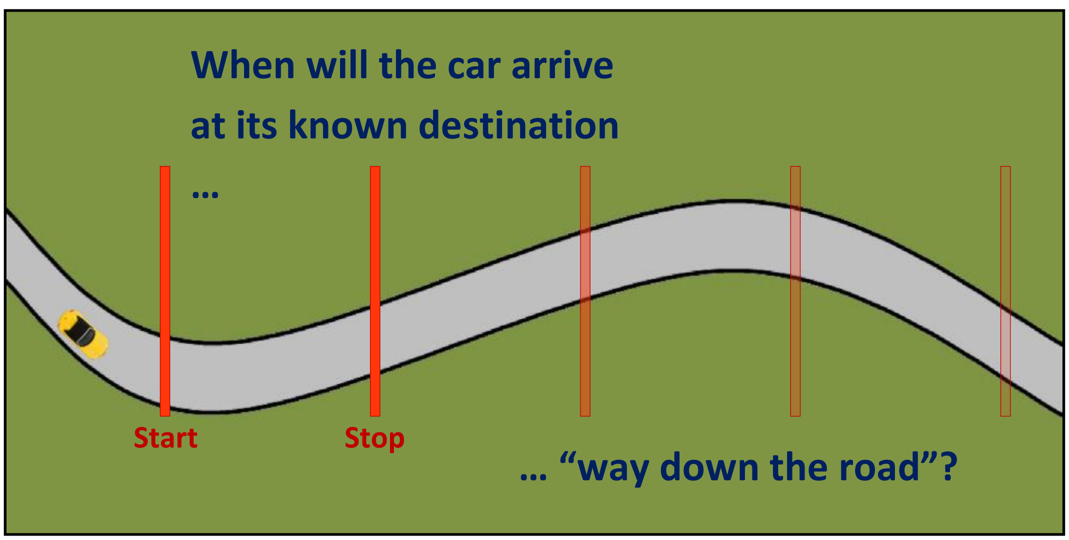

Figure 4

However, you can see there are some curves in the road along with some straight sections. It's almost certain the car will go different speeds over the different stretches of road. With a little more work, you can improve your estimate by taking as many measurements as practical — some on curves, some on straight sections — and averaging them.

The Road Unseen

There is something, though, that you're going to have to make your best guess about. You can only see that part of the road where you've been taking measurements. You're not really sure what it looks like over the entire remaining distance to the destination — which is a lot longer than the stretch you can see. You'll have to assume either the remaining distance is like where you've measured … or it's not. If you think it's not, then you'll have to decide whether it allows for a faster or slower average speed. Or you can say, "Here's the most likely arrival time, and here's the range I'm almost certain it will fall within." You simply can't make a perfect estimate … but you can sure make a good one.

What you do with TRA is similar to the car example. You're trying to measure something that's already in motion. You're trying to grab one or more time slices of representative information. And with that data, you're trying to predict a final outcome that's still some ways off in the future.

Setting Up

In TRA, time slices normally overlap (Figure 5), partly to increase the overall number of time slices and partly to capture the nuances of changing conditions. Note that each time slice is offset consistently from the previous one. Also, the duration of each time slice is the same. Most importantly, recognize that a given loan can be part of any number of time slices, but it begins each time slice at its then current risk rating and amount.

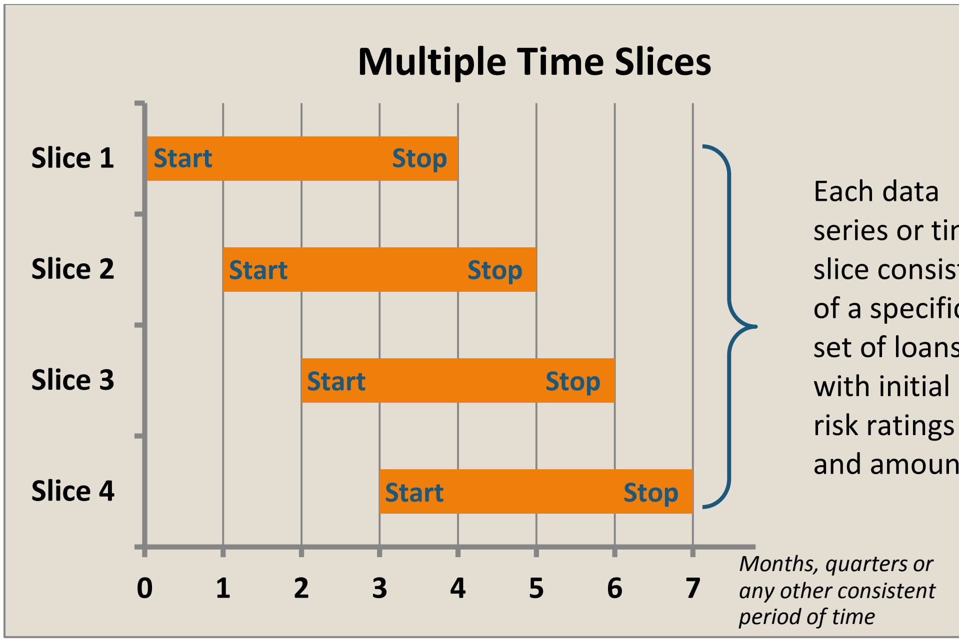

Figure 5

Tracking Loans

Assume you have a ridiculously small and simple portfolio. There are only two loans, each risk rated "Good" and each with $100 outstanding. You track each loan individually throughout the duration of a time slice, which you've set at four quarters. As mentioned above, you're primarily looking for changes in risk ratings and reductions in amount. Figures 6 and 7 show what you might find.

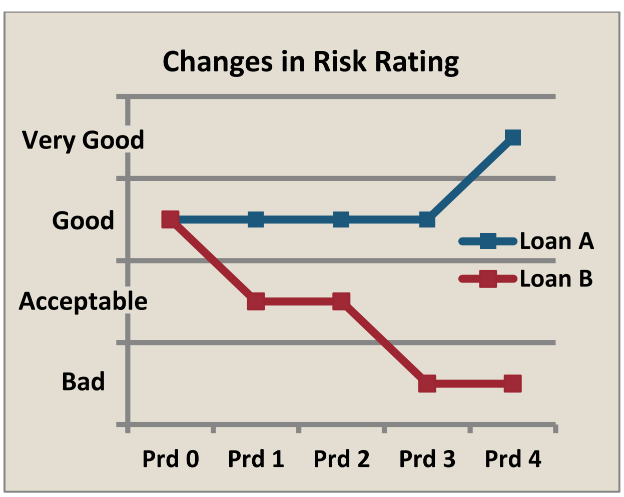

Figure 6

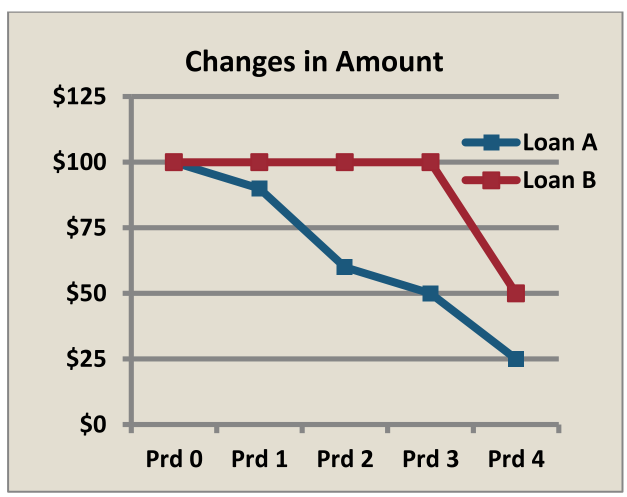

Figure 7

From time slice start to time slice stop, Loan A improved from "Good" to "Very Good" and its dollar amount decreased from $100 to $25, which you can safely assume, occurred by repayment. Loan B, on the other hand, deteriorated from "Good" to "Bad" while its dollar amount decreased from $100 to $50, which is assumed here to have resulted from a charge-off. Figure 8 summarizes these outcomes.

	Dollars			Percentage		
Rating	Outstdg	Repaid	C/O	Outstdg	Repaid	C/O
Very Good	$ 25	$ 75	$ 0	12.5%	37.5%	0.0%
Good	0	0	0	0.0	0.0	0.0
Acceptable	0	0	0	0.0	0.0	0.0
Bad	50	0	50	25.0	0.0	25.0
Total	$ 75	$ 75	$ 50	37.5%	37.5%	25.0%

Figure 8

By looking at the "Total" row, you can see that all $200 present at the start of the time slice have been accounted for. In percentages, you can say that of the $200 ***pool*** of "Good" rated loans you started with 25% shifted to "Bad," 12½% shifted to "Very Good," 37½% repaid and 25% was charged off … which accounts for all 100%.

These are the most basic TRA time slice mechanics — tracking each loan found in a given risk pool to identify risk rating changes and reductions in amount, then aggregating the loan-by-loan information for the pool as a whole. If there are tens of thousands of loans, the task requires some computing power, but the basic calculations are not complicated.

Risk Pools

So far, you've been working with just a two-loan example. Now consider a more complete view of risk pool percentages, one with a realistic number of loans distributed among all risk ratings. Figure 9 shows how every loan dollar begins right where it's found at the start of a time slice … because there haven't been any ratings shifts or reductions yet.

Time Slice Start		Reduction Outcomes					
		Very Good	Good	Acceptable	Bad	Paid	C/O
Beginning Risk Rating	Very Good	100%					
	Good		100%				
	Acceptable			100%			
	Bad				100%		

Figure 9

Reduction Outcomes

By the time you get to the end of a time slice (say, four quarters), the dollars in each risk pool are distributed across various ***reduction outcomes*** (Figure 10). The perceived quality of some loans improved while some worsened. A good portion of each pool was repaid. And, as you would have guessed, some of the dollars in the "Bad" pool were charged off. Still, if you add across any row, you get 100%. The beginning amount of each pool is fully accounted for. (Notice the repayment rates and the tendency of loans to stay put.)

Time Slice Stop		**Reduction Outcomes**					
		Very Good	Good	Acceptable	Bad	Paid	C/O
Beginning Risk Rating	Very Good	35% →	25%	0%	0%	40.0%	0.0%
	Good	10% ←	40% →	13%	0%	37.0%	0.0%
	Acceptable	0%	6% ←	50% →	8%	36.0%	0.0%
	Bad	0%	0%	7% ←	35%	48.0%	10.0%

Figure 10

The reduction outcomes shown in Figure 10 are at the core of TRA methodology. What if you assumed these patterns continued repeating until all the dollars were gone from the portfolio? That would mean the 25% of "Very Good" loans that fell to "Good" would begin behaving like all the other "Good" loans. Some would actually return to "Very Good" (at least for a time) and some would decline further to "Acceptable." And all the while, repayments would continue to shrink all the pools.

Mapping out the entire array of movement is a bit daunting … and doing period by period calculations, even on a spreadsheet, is difficult. With some knowhow, there are faster and better ways to do the calculations. Explaining the math, though, is not the purpose here; rather, it's to show how to use the resulting information.

Life of Portfolio

Figure 11 shows the projected ***life of portfolio*** reduction outcomes. By then, there would be no more balances outstanding. Every dollar would have to be accounted for in one of only two ways, repayment or charge-off. And, as you can see, 100% of the beginning amount of each risk pool is indeed accounted for by repayment or charge-off. Also, as you would expect, the worse the risk rating at the start of a time slice, the worse the projected charge-off over the remaining life of the portfolio.

Life of Portfolio		**Reduction Outcomes**					
		Very Good	Good	Acceptable	Bad	Paid	C/O
Beginning Risk Rating	Very Good	0%	0%	0%	0%	99.8%	0.2%
	Good	0%	0%	0%	0%	99.4%	0.6%
	Acceptable	0%	0%	0%	0%	97.4%	2.6%
	Bad	0%	0%	0%	0%	84.3%	15.7%

Figure 11

You can also display reduction outcomes graphically (Figure 12). This example begins (left side) with the portfolio risk distribution shown in Figure 1. The very next data points are the known outcomes of the single time slice (Figure 10). Then, from left to right, you see the full sequence of projections that culminated in the table in Figure 11.

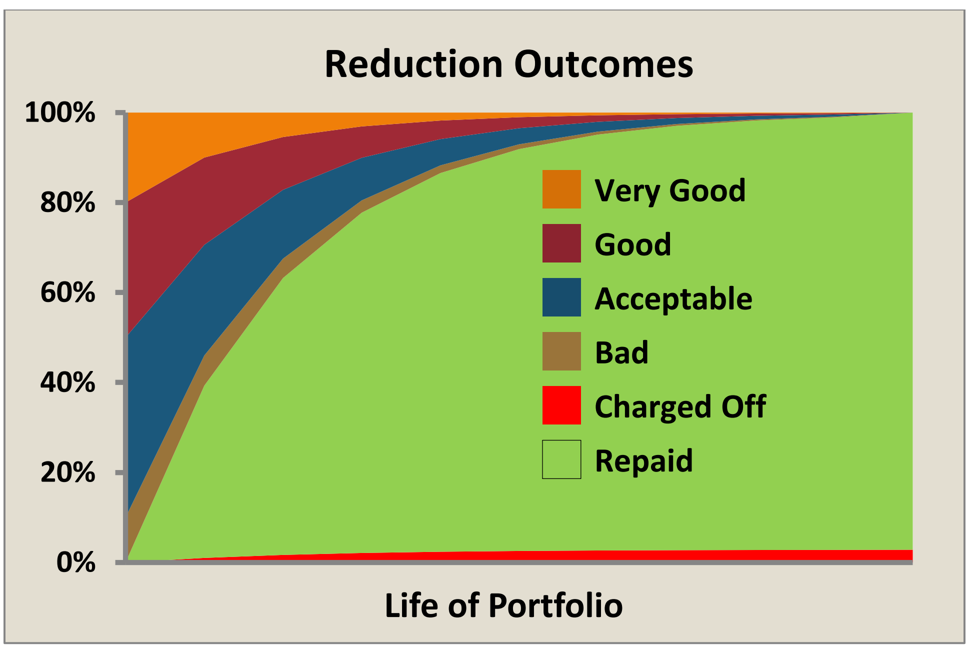

Figure 12

By the end of the remaining life of the portfolio (right side), all the beginning dollars are projected to have been repaid or charged off. Clearly, the most important thing on the chart is the far right side of the red charge-off line. The thickness of that red line suggests the eventual impact of the losses inherent in the current portfolio. It needs to be a very, very thin line.

More Than One Slice

So far, you've been working with just a single time slice. Remember the example of the car, where you wanted to improve your estimate by taking additional measurements along different stretches of road? You do something similar with TRA. You gather data from as many available time slices as appropriate, perform the same sequence of steps … and average the results. You can call the ***average*** projected charge-off for a given risk pool the probable or ***expected loss*** for that pool.

Figure 13 shows a chart of the projected charge-off percentage for the "Acceptable" pool of loans covering ***twelve*** time slices. As you can see, the average projected charge-off for the series is 2.6%. In other words, you can say the expected loss for loans in the current portfolio that are risk rated "Acceptable" is 2.6%.

Should you notice, the dim green line is a linear trend put there as one example of several ways a pattern can be "interpreted." In this case, the line slopes downward showing a favorable trend. Among other possibilities, you could also plot a moving average.

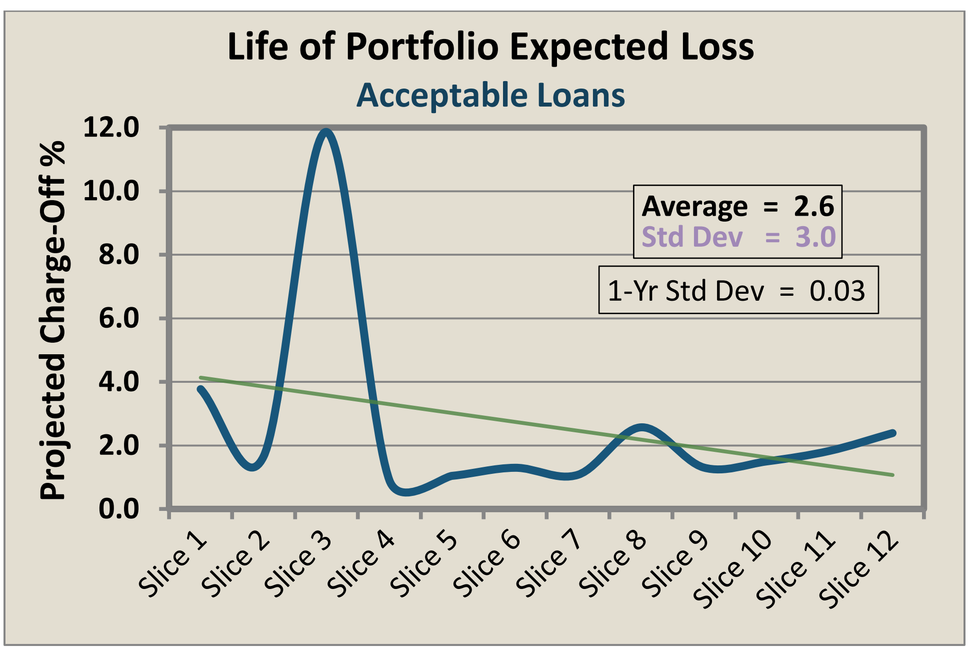

Figure 13

Directly below the average in Figure 13, you'll see a "grayed out" standard deviation. (You'll recall that the standard deviation is a common measure of volatility — the higher the value, the more volatile the measure.) While the life of portfolio standard deviation is quite meaningful, it's been grayed out nonetheless to focus on a different value. For now, just make note of the "other" standard deviation, the one pertaining to a ***one-year time horizon***. That value will be important later during a discussion of ***unexpected*** loss.

The table in Figure 14 summarizes projected charge-offs for all risk pools from the twelve time-slice TRA. It shows averages and standard deviations for both the life of the portfolio and for just a one-year horizon. (Yes, the life of portfolio averages exactly match the projected charge-offs for the single time slice above (Figure 11). This was intentional to minimize confusion caused by too many different numbers.)

	Life of Portfolio		One-Year Horizon	
Risk Pool	Average	Standard Deviation	Average	Standard Deviation
Very Good	0.2%	0.3%	0.01%	0.01%
Good	0.6%	0.9%	0.02%	0.01%
Acceptable	2.6%	3.0%	0.06%	0.03%
Bad	15.7%	11.1%	9.53%	5.63%

Figure 14

Selecting Time Slices

Before leaving the topic of multiple time slices, consider the issue of time slice selection. There are two areas of concern, one internal and one external. Regarding the internal, if you knew there was a major change in credit behavior at some point in your bank, perhaps due to a policy change or revision to the approval process, you'd be aware that time slices occurring before the change might produce different outcomes than time slices occurring after. You'd benefit by doing several analyses. For instance, by selecting distinct sets of time slices reflecting before, during and after change scenarios, you could arguably make a better case about which scenario best reflects the inherent loss in the ***current*** portfolio.

Figure 13 also provides a cautionary tale about time slice selection and potential data manipulation. Notice the spike in the projected charge-off percentage for Time Slice 3. Clearly, some atypical losses occurred, and that spike affected the average and the standard deviation. There may be a tendency by some to argue, "Well, this is an anomaly. These things don't really happen to us. We should drop that from the data." That would be a mistake. So-called "strange things" happen with surprising regularity ... meaning they're really not that strange. They're a normal part of data. Don't try to excise them. Instead, try to interpret them properly.

Obviously, external conditions also affect credit outcomes. Most loan portfolios will exhibit different reduction patterns at different points in the economic cycle. Knowing these cycle-pattern differences and knowing where you are in the current cycle can lead to more refined — and more realistic — estimates of loan loss. Figure 15 shows a four-quarter moving average of real GDP over the past 15 years.

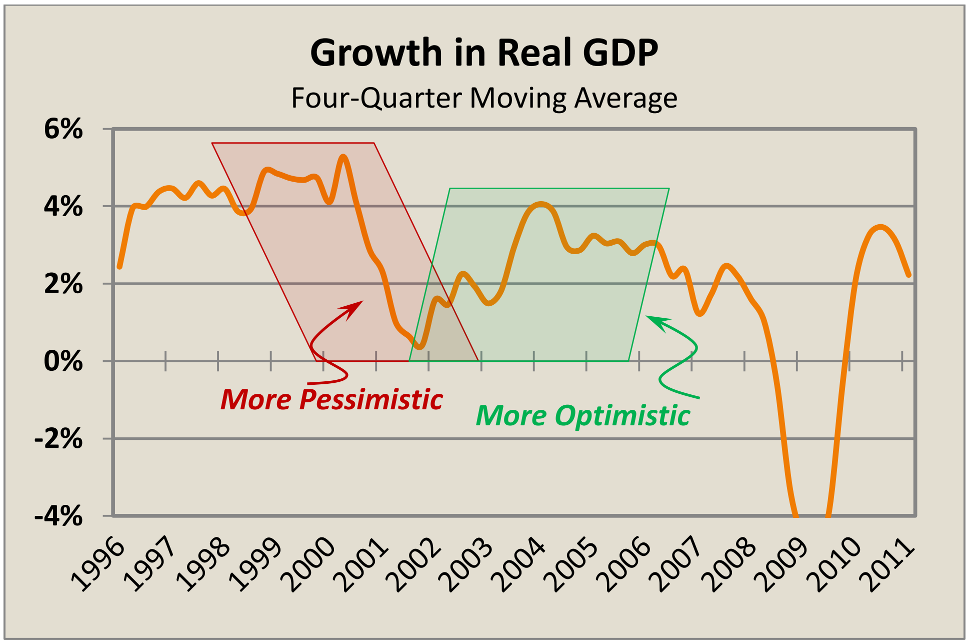

Figure 15

Time slices spanning roughly 1999 to 2003 would probably produce more pessimistic outcomes than time slices spanning, say, 2002 to 2006. The key consideration in selecting an appropriate time slice range is how well it matches up with where you ***think*** the economy is headed today. Investigating several different time slice ranges might also be helpful. Although, you would normally want to use time slices from recent experience, there might be justification for using older data when the underlying economic cycle more closely matches today's circumstances.

As an aside, there's good reason to question the statistical value of the 2008–2009 near collapse of the global economy. It was the proverbial "100-year flood" (more accurately, the 70-year flood). You obviously can't pretend it didn't happen, but it was so extreme that incorporating it into a model — other than for identifying worse case limits — isn't going to be especially helpful.

FAS 5

Having generated all this information, it would now seem appropriate to put it to good use. Regardless of the type of business and most assuredly for banks, FAS 5 says that any loss that is "probable and estimable" is to be reserved for. The life of portfolio average projected charge-off or ***expected loss*** shown in Figure 14 for each risk pool perfectly meets this classic FAS 5 test. TRA is precisely the kind of analysis and produces precisely the type of loss factors FAS 5 calls for.

If you were to ignore FAS 114 for the time being and further assume it's acceptable to lump all your loans together in one general category — a gross oversimplification that's only appropriate for illustration — then the calculation of required loss reserves is very straightforward (Figure 16). The required reserve for each risk pool is obtained by multiplying its expected loss factor by its respective distribution percentage.

Current Portfolio	Risk Pool Distribution	Expected Loss	Required Reserve
Very Good	20%	0.20%	0.04%
Good	30%	0.60%	0.18%
Acceptable	40%	2.60%	1.04%
Bad	10%	15.70%	1.57%
Total	100%		2.83%

Figure 16

FAS 114

In real life where you're dealing with actual dollars, to incorporate FAS 114, you simply subtract the balances of any FAS 114 loans from their respective risk pools. You multiply the remaining ***true pool*** balances by their respective expected loss factors to obtain the FAS 5 portion of the required reserve. To that you add the loan-by-loan reserve amounts determined in accordance with FAS 114. You now have your total required reserve.

It would be worthwhile, though, to compare the percentage reserve you get on just FAS 114 loans to the appropriate TRA expected loss factor(s). You might find your bank reserving ***less*** on FAS 114 loans than what TRA analysis suggests based on actual experience. One of the objectives of FAS 114 is to bring more rigor and intimate analysis to reserving for the "biggest and baddest" loans. In practice, however, if familiarity doesn't breed outright contempt, it can still lead to a loss of respect for just how risky FAS 114 loans can be.

Stress Testing

In the past year or so, business news media have devoted a measure of attention to how regulators have pressed banks to conduct "stress tests." Unfortunately, neither media nor regulators have made it clear just what stress testing is or how it should be done. You can be sure, though, that it usually involves credit risk and often seems to be tied to loss reserves. However, what isn't so apparent is that even though stress measurements may be focused on credit risk and loss reserves, what's really being tested is capital adequacy. The threat comes from credit exposure. The first line of defense is the ALLL … and the backstop is capital.

With your TRA results, you can engage in some fairly straightforward stress testing. For instance, based on the range of time slices and your assessment of where the economic cycle is, you think your portfolio might be on the verge of, say, a 20% across the board, downward shift in risk ratings. Figure 17 examines the impact of such an occurrence.

20% Shift	Out	In	Resulting Distribution	Expected Loss	Required Reserve
Very Good	−4%		16%	0.20%	0.03%
Good	−6%	+4%	28%	0.60%	0.17%
Acceptable	−8%	+6%	38%	2.60%	0.99%
Bad		+8%	18%	15.70%	2.83%
Total			100%		4.02%

Figure 17

The math is straightforward. 20% of each risk pool falls into the next worse pool … except "Bad," which can't get any worse. Take the "Good" pool as an example. It gains 4% from "Very Good" (20% X 20%) while losing 6% to "Acceptable" (20% X 30%) and declining overall to 28%. What's remarkable is that a 20% shift in risk ratings can cause a 42% increase in the required reserve (from 2.83% to 4.02%). The question here is not about the reserves withstanding the increase. The reserves are to be whatever they need to be. The question is about earnings and the adequacy of capital.

Here's another variation on stress testing (Figure 18). In addition to a 20% downward shift in risk ratings, assume collateral values suffer a serious decline causing expected losses to increase by, say, 25%. Keeping the portfolio distribution from Figure 17 and increasing each loss factor by 25% results in a required reserve of 5.02%. Meeting that reserve level would have a major impact on capital. Proper analysis of the issue will focus on probabilities. What's the likelihood of this scenario playing out? The higher the likelihood, the more necessary it is to respond with plans of action.

20% Shift *25% Worse Loss*	Out	In	Resulting Distribution	Expected Loss	Required Reserve
Very Good	−4%		16%	0.25%	0.04%
Good	−6%	+4%	28%	0.75%	0.21%
Acceptable	−8%	+6%	38%	3.25%	1.24%
Bad		+8%	18%	19.63%	3.53%
Total			100%		5.02%

Figure 18

Capital Adequacy

Finally, there's the matter of unexpected loss and capital adequacy. (The previous stress testing examples get at some of these same concerns.) Everyone talks about unexpected loss and its impact on economic or risk capital ... but almost no one offers any meaningful way to calculate that impact. Basel II offers some explicit instructions for measuring unexpected loss. This isn't the place to delve into the Basel Accords, but there is value in reviewing its recommended treatment of unexpected loan losses.

What's the worst possible — or nearly worst possible — loss that could occur over, say, a one-year time horizon in your bank? (There's no point considering repetition of this loss for another year. If that happens, there simply won't be another year.) It's difficult to put a number to. But through the twelve time-slice TRA example (Figures 13 and 14), you have some real measure of the volatility of projected losses based on extensive data. Basel II suggests that a loss rate that's 99% likely ***not*** to be exceeded over a one-year time horizon is an acceptable measure of unexpected loss that should be explicitly covered by capital.

The chart in Figure 13 reintroduced you to standard deviations, specifically, the standard deviation for projected loss over a one-year horizon. If you accept that multiplying one standard deviation times 2.3265 gets to a 99% level of probability (see Wikipedia), then you can calculate the capital needed to support the loss volatility in your loan portfolio.

Figure 19 shows a way to do this. The column labeled "One-Year Standard Deviation" lists the corresponding values from Figure 14. Multiplying those amounts times 2.3265 results in the "Unexpected Loss" values. Multiplying the unexpected loss for each risk pool times its respective portfolio percentage results in the "Additional Capital" needed. In other words, to support the volatility of expected losses in the loan portfolio, the bank needs a dedicated amount of capital equal to 1.35% of total loans.

Current Portfolio	Risk Pool Distribution	1-Year Std Dev	Unexpected Loss	Additional Capital
Very Good	20%	0.01%	0.02%	0.00%
Good	30%	0.01%	0.02%	0.01%
Acceptable	40%	0.03%	0.07%	0.03%
Bad	10%	5.63%	13.10%	1.31%
Total	100%			1.35%

Figure 19

Don't be overly alarmed by any suggested need for additional capital. Most community banks with more than 10% total capital (i.e., "well-capitalized" banks) probably already have sufficient capital to cover credit loss volatility. They just haven't had a way to isolate the information to show and explain it.

Is TRA Worth the Effort?

You've seen a lot of tables, charts and numbers. Terminal Risk Analysis does take some effort. Because of the dynamics of loans shifting in either direction concurrently with reductions in dollar amounts, mathematical computations are relatively complex. However, the complexity of the mathematics is far less challenging than building and maintaining a data warehouse of accurate and reliable information. Still, it's all very doable, even for community banks.

Is Terminal Risk Analysis worth the effort? If you're looking for a rigorous, evidence-based methodology for supporting your assessment of ALLL and capital adequacy, the answer is yes. Next to the analytical effort expended to make loan approval decisions, few other exercises can produce better insight into expected loan portfolio outcomes.

