

Quantitative Finance in the Post Crisis Financial Environment

Kevin D. Oden

INTRODUCTION

Mathematics is a disciplined approach to solving problems. As a byproduct of that approach, beautiful truths are often discovered, problems sometimes solved, and these beautiful truths and sometimes solved problems usually lead to even more interesting problems.

Finance has historically been rich in problems to solve:

- portfolio optimization;
- investment strategy;
- performance analysis;
- estimating fair value;
- price prediction.

From attempting to solve these financial problems, many beautiful “truths” have been discovered:

The views expressed in this document are the author’s and do not reflect Wells Fargo’s opinion or recommendations.

K.D. Oden (✉)
Wells Fargo & Co., San Francisco, CA, USA

- mean-variance analysis;
- factor modeling and arbitrage pricing theory;
- performance ratios;
- risk neutral pricing;
- and more recently, market microstructure theory.

So far most studies in mathematical finance cover material to understand and analyze these mostly “classical” financial problems, with most putting a strong emphasis on the risk neutral pricing of derivatives. Though these problems remain an important aspect of quantitative risk management for banking organizations around the globe, the financial crisis brought to the forefront many old problems that now needed to be viewed in the light of a new financial environment, exposing flaws in how these problems were traditionally approached. Furthermore, the crisis raised new issues that required quantitative analysis and potentially a new set of tools to perform these analyses.

In the next sections we will briefly explore the quantitative problems associated with five risk areas: the fair value of credit, debt, funding and capital risk, collectively known as XVA risk; operational risk, fair lending risk, financial crimes risk, and finally model risk. The problems analyzed fall both in the category of “old with exposed flaws” as well as “new and in search of new tools”. However, each of these topics is worthy of a book in its own right, so by design we cannot delve deeply into any one of them, but provide relevant references for the reader to continue her research. Finally, we note that this is not intended as a comprehensive list of all of the quantitative problems facing the industry today. But in many respects these quantitative problems have emerged post-crisis and have found themselves on the top of many firm and regulatory agendas.

THE FAIR VALUE OF CREDIT, DEBT, FUNDING, AND CAPITAL RISK (XVA)

Credit has always been one of the largest exposures for commercial banks. And even prior to the financial crisis derivative traders at commercial banks realized that all derivative counterparties were not created equally from a credit perspective. The credit quality of the derivative counterparty should be taken into account either through collateral arrangements or through reserving a portion of expected profit on transactions with counterparties. These adjustments to the value of the contract were made to compensate for the possibility the counterparty defaulted before expiry of the transaction and the costs associated with replacing or offsetting the risk. The

notion of adjusting the fair value of a derivative to account for the credit quality of a counterparty became memorialized as an accounting standard in the USA in 2006 with the FASB 157 and in the European union in IAS 39.¹ These accounting standards require credit risk of both participants in the derivative transaction to be reflected in the fair value of the derivative. These adjustments are known as credit valuation adjustment (CVA) and debt valuation adjustment (DVA).

The crisis also revealed that a longstanding assumption that the cost of funding a collateralized trade was roughly the same as an uncollateralized trade could be dramatically wrong as the benchmark for funding most commercial banks (LIBOR) widened to historic levels versus the cost of carrying or the benefit of lending collateral which is typically benchmarked at the overnight index swap rate (OIS) in the USA, or the sterling overnight index average (SONIA) in the UK. This newly recognized risk led to a new adjustment to the fair value of a derivative known as a funding valuation adjustment (FVA). This adjustment could be a cost, if the derivative is an asset that needs to be funded or a benefit if the derivative is a liability.

Similarly, each derivative instrument the bank transacts attracts one or more capital charges. Typically, there will be a charge for the risk of loss associated with market movements (market risk capital charge) and there will be a charge associated with the potential for counterparty default (a counterparty credit capital charge). This capital must be held for the life of the transaction and will vary over the life of the transaction depending on the credit quality of the counterparty, the market, and the remaining maturity of the transaction. Clearly, the level of expected capital that must be held throughout the life of the transaction impacts the profitability of the trade and should be reflected as an “adjustment” to the fair value of the trade. This adjustment is known as a capital valuation adjustment or KVA.

In summary we have the following adjustments with some of their key properties:

- Credit valuation adjustment (CVA)
 - It is always a decrease in the fair value of a financial instrument or portfolio due to the risk of a counterparty defaulting before the expiration of the trade.

- An increase in credit risk of the counterparty results in a decrease in fair value.
- To hedge this risk requires a hedge of the market exposure as well as the credit exposure. Whereas a credit default swap (CDS) can hedge the credit quality of a counterparty for a fixed notional, the fair value of a derivative portfolio changes with the changes in the underlying market value (s) driving the derivative. Hence hedging CVA would require a contingent (upon the fair value of the derivative) CDS or a CCDS.
- Debt valuation adjustment (DVA)
 - The increase in fair value of a financial instrument or portfolio due to the commercial bank's (own entity's) risk of defaulting.
 - An increase in risk results in an increase in the fair value.
 - The increase in the likelihood of default of the commercial bank, implies an increase in the likelihood it will not have to pay all or some of its outstanding or potential liabilities. Though counterintuitive, from an accounting perspective this is a net benefit to the commercial bank.
 - Though a net benefit, DVA leads to profit and loss volatility as changes in market factors and the change in the credit quality of the commercial bank changes the value of DVA. Hedging this risk is difficult because most banks will or cannot buy credit protection on themselves. Therefore, typically hedging this exposure is done through buying or selling credit protection on a basket of names highly correlated with the credit risk of the bank.
- Funding valuation adjustment (FVA)
 - The cost (benefit) from borrowing (lending) the shortfall (excess) cash from daily derivatives operations.
 - FVA can be either a cost or benefit to the commercial bank.
 - A derivative asset or one that has a net positive fair value to the bank needs to be funded.² Similarly a derivative liability benefits from the bank's investment rate of return. But the funding and investing rates will differ from bank to bank. From a theoretical perspective, this brings up the question of should FVA be even considered a true adjustment to the price of a derivative since its inclusion breaks the rule of one price.
- Capital valuation adjustment (KVA)
 - The expected cost of capital associated with the derivative over the life of the trade.

- KVA is a random variable depending on the anticipated future fair value of the derivative and the present value of the capital associated with the derivative.

The ideas behind these concepts are relatively straightforward. However, they are surprisingly difficult to implement in practice. We describe their calculation a little further in the following and outline some of the practical difficulties.

If we define V_R to be the default free and capital free price of an asset and let V denote the default risky price of the asset adjusted for the cost of capital, then one can write

$$V = V_R - CVA + DVA \pm FVA - KVA$$

In particular, if the buyer and the seller of the derivative agree on their respective risk of default (or credit spread) than there is an arbitrage free “price” agreed on by both parties, \tilde{V} :

$$\tilde{V} = V_R - CVA + DVA$$

However, FVA can be a cost or benefit and more importantly depends on the funding costs of the holder of the derivative and therefore apparently breaks the single price paradigm of arbitrage-free pricing. Furthermore KVA also depends on the capital requirements of a particular bank which, among other things, could depend on the bank’s resident jurisdiction, as well as its size and level of sophistication.

To give some idea of the complexity of accurately calculating each of these adjustments, we observe that CVA and DVA will require the estimation of the counterparty’s and the bank’s credit quality throughout the life of transaction, the correlation between these two and their correlation with underlying market risk factors. In order to estimate FVA and KVA we will need to know the cost of funding the derivative and the capital throughout the life of the transaction. See [16] for a good treatment of XVA risk generally and plenty of references.

Except for CVA and DVA, how these adjustments should impact the fair value of the derivative is an open debate in the academic world as well as the accounting world. However, in the industry, whether accounted for in the fair value of the derivative or not, there is a growing realization that there are risks associated with each of these that must be managed as they certainly impact the economic value of the derivatives portfolio.

In summary, we can say the fair (or economic) value of an instrument is impacted by the non-performance risk of the counterparty (CVA), the legal entity's own credit risk (DVA), as well as the cost or benefit of funding the instrument (FVA) and the opportunity cost of capital associated with holding or selling the instrument (KVA). Each of the adjustment concepts is relatively easy to grasp. Yet, even in isolation each can be a difficult quantity to compute, depending on forward implied credit spreads, underlying market risk factors and their implied correlations, to list just a few of the driving factors. These complications alone will provide a fruitful source of research for quantitative financial researchers for years to come (see [17]).

OPERATIONAL RISK

The original Basel Accord set aside capital requirements for credit and market risk. Losses associated with operational failures or the legal fallout that followed were mostly associated with failures in credit processes or market risk management lapses. But over time it became increasingly clear that many losses were not breakdowns in credit or market risk management; but failures in processes were clear and distinct from these two disciplines and could ultimately result in significant credit or market risk losses or even more punitive legal claims. Therefore, the second Basel Accord (Basel II) [2] clearly defined operational risk as “the risk of loss resulting from inadequate or failed internal processes, people and systems, or external events” and prescribed three approaches to calculate capital for this risk.

The financial crisis highlighted how pervasive and impactful the poor management of operational risk could be to financial institutions, in particular, commercial banks with lapses in sound mortgage origination practices to errors in home foreclosures. Some of the costlier operational risk losses before and after the financial crisis are listed below [13, 15]:

- \$25 billion—Ally Financial Inc., Bank of America Corp., J.P. Morgan Chase & Co., Citigroup Inc., Wells Fargo & Co., 2012: The five banks agreed to pay \$25 billion in penalties and borrower relief over alleged foreclosure processing abuses.

- \$13 billion—J.P. Morgan Chase & Co.—2013: J.P. Morgan and the Justice Department agreed to a landmark \$13 billion settlement that resolved a number of legal headaches. Of the \$13 billion settlement, \$9 billion was set aside to pay federal and state civil lawsuit claims over residential-backed mortgage securities. Of that \$9 billion, \$2 billion was a civil penalty to the Justice Department, \$1.4 billion was to settle federal and state claims by the National Credit Union Administration, \$515 million to settle Federal Deposit Insurance Corp. claims, \$4 billion to settle Federal Housing Finance Agency claims, nearly \$300 million to settle claims by California state officials, nearly \$20 million to settle claims by Delaware, \$100 million to settle claims from Illinois, \$34 million to settle claims by Massachusetts, and nearly \$614 million to settle claims by New York State.
- €4.9 billion—Société Générale (SocGen)—2008: A rogue trader, Jerome Kerviel, systematically deceives systems, taking unauthorized positions worth up to €4.9 billion in stock index futures. The bank has enough capital to absorb the loss but its reputation is damaged.

The increased losses leading to and immediately after the financial crisis increased pressure to improve the models assessing operational risk capital and more broadly enhance and standardize the practices related to operational risk management. On the capital front, Basel III [5] provided three approaches for calculating operational risk capital:

- the basic indicator approach;
- the standardized approach;
- the advanced measurement approach (AMA).

We focus here on the last approach, because it gives the industry the most latitude to produce modeling techniques that address the idiosyncratic nature of operational risk at the particular commercial bank. This latitude also means a wide range of practice has developed around the calculation under the approach and the realization by regulators and practitioners alike that the problem of quantitatively assessing capital for operational risk is a difficult problem still in its infancy.

The regulatory requirements for calculating operational risk capital under the AMA framework (Basel III [3], [4]) are essentially the following:

- one-year holding period and loss (gross of expected loss) at the 99.9th percentile
- the use of business environment and internal control factors (BEICFs)
- the use of internal data
- the use of external data
- the use of scenario analysis

One common approach to addressing this problem is to assume that over the year-long period the number of loss events has a prescribed distribution and the severity of each loss, given an event has a stationary conditional distribution. Denoting the collective loss over the holding period by L , then collecting these facts the formal model for losses can be described as follows:

- Let n be a random number of loss events during the holding period;
- Let X_i be a random variable representing the magnitude of loss for event i ;
- Then $L = \sum_{i=1}^n X_i$ is the total loss in the holding period.

The challenge then becomes estimating the CDF for L in order to find quantiles,

$$F_L(y) = \text{Prob}(L \leq y).$$

Once the distribution is determined we can assess capital, K :

$$K = y^* \text{ given } \text{Prob}(L \leq y^*) = .001.$$

This is very similar to the market risk Value-At-Risk (VaR) framework. So in theory, the exercise is very tractable, but in practice, there are many difficulties with this approach. First, VaR in the market risk setting is typically measured at the 99th percentile for capital purposes or at the 97.5th percentile for day-to-day risk management purposes and typically on a one- or ten-day basis. The operational risk 99.9th percentile over a year period requires the measurement of one in 1000 year events. There simply is not enough data at any commercial bank to accurately measure the tail of the probability distribution. Even combining data from various

institutions and using methods from the theory of extreme value theory (EVT) still make the task practically difficult.

As we have already noted, most institutions do not have enough operational loss data (internal data) to estimate the extreme tail of the loss distribution reliably. Even if an institution does have enough data for a particular type of loss, loss data is inherently non-homogenous and the tails for some types of losses (e.g. employee fraud) may have different distributional characteristics than the tails for other types of losses (e.g. sales practice failures). These groupings of losses are typically known as operational risk categories (ORC). So in practice banks must estimate multiple losses, L_i where i ranges over all ORCs. In this case, data becomes even sparser.

If external data is brought in to augment internal data (which it must by the rule), how do we scale the external data to fit the risk characteristics of the firm being modeled? For example, using external loss data for credit card fraud from an organization that has multiples of exposure to the modeled company without some kind of scaling of the data would lead to misleading and oversized capital related to credit card fraud.

Besides the challenge of estimating the distribution for each L_i there is the task of modeling the co-dependence of the frequency and severity of each loss in each ORC along with modeling the co-dependence structure between the ORC groupings.

Given the complications both theoretical and practical outlined above, many market practitioners and recent regulatory agencies have questioned the feasibility of a modeled capital charge for operational risk and whether a more standard and possibly punitive charge should be levied. This approach too has its disadvantages, as simply adding capital without any relationship to the risk it is meant to cover is not only bad for the industry but in fact poor risk management. So regardless of the direction the Basel Committee ultimately takes, the industry will need to tackle this difficult problem to better manage this risk.

FAIR LENDING RISK

Fair lending risk is the risk that a financial institution's lending operations treat applicants and borrowers differently on a prohibited basis, treat applicants in an unfair or deceptive manner, or subject applicants to predatory or abusive lending practices.

Fair lending risk analysis aims to monitor compliance with the fair lending laws and statutes, in particular the Equal Credit Opportunity Act (ECOA) and the Fair Housing Act (FaHA). But the origins of fair lending analysis go back to at least 1991 when data collected under the Home Mortgage Disclosure Act (HMDA) was first released [24].

In 1975 Congress required, through HMDA, loan originators to maintain data on mortgage originations, mainly to monitor the geography of these originations. In 1989 after further amendments to HMDA, requirements were included to retain the race and ethnicity of loan applicants along with denial rates. When this information was released in 1991, the results not only fueled outrage in some circles because of the disparate loan denial rates between blacks, Hispanics and whites, but instigated the Federal Reserve Board of Boston to perform a detailed statistical analysis of the data in order to draw conclusions about discriminatory practices in mortgage lending. This now famous, heavily scrutinized, and often criticized study is popularly known as the “Boston Fed Study” (see [24] for references) and in many ways laid the foundation for all fair lending analysis that followed.

However, fair lending analysis now extends to all forms of credit, ranging from auto loans to credit cards; from home improvement loans to home equity lines. Beyond the origination of credit, there are requirements to identify abusive practices, like predatory lending (e.g. NINJA loans³) and unfair foreclosures. And, in the wake of the financial crisis, the Dodd-Frank Wall Street Reform and Consumer Protection Act of 2010 created the Consumer Financial Protection Bureau (CFPB), whose primary task is to protect consumers by carrying out federal consumer financial laws. In particular, as it relates to fair lending, the CFPB is the primary regulator that attempts to detect and enforce remediation related to unfair lending practices. Typically, these policies look to detect discriminatory treatment of persons in protected classes. Though the definition of protected class varies by jurisdiction and regulation, most laws provide protection based on race, color, religion, gender, national origin, and sexual orientation.

Fair lending risk can be broken down into two broad types:

- Disparate impact risk
 - This occurs when the policies, practices, or procedures have a statistically different impact on individuals in a protected class compared to similarly situated credit qualities of non-members of the protected class.

- Disparate treatment risk
 - This occurs when individuals in a protected class are intentionally treated differently than non-members of the protected class.

One problem facing commercial banks is then is to determine if their credit approval processes have disparate impact or treatment. To frame this problem mathematically, we follow closely the presentation of Ross and Yinger [24]. Suppose that π is the expected profitability (or performance) of a credit product (e.g. mortgage, auto loan, credit card, etc.). The profitability of the loan is dependent on factors related to the characteristics of the loan, the applicant’s credit quality, and the underlying asset (e.g. location and size of property, type of car, etc.), which we denote by L , C , and A , respectively. Each of these characteristics may have a number of variables which describe their quality. Denoting these variables generically by X_i , $i = 1, \dots, n$. We write

$$\{L, C, A\} = \{X_1, X_2, X_3, \dots, X_n\},$$

and the lending profitability function becomes

$$\pi(L, C, A) = \pi(X_1, X_2, X_3, \dots, X_n).$$

So, the lending problem, absent overt discrimination by race, ethnicity, gender, sexual orientation or another protected class, reduces to making a loan when π is above some threshold π^* and denying the application otherwise. That is,

$$\begin{cases} \text{approve } \pi > \pi^* , \\ \text{deny } \pi \leq \pi^* . \end{cases} \tag{1}$$

The set-up in Eq. (1) lends itself nicely to the credit scoring analysis typically performed using logit, probit, or even ordinary least squares (OLS) analysis. However, one of the drawbacks of the problem as stated is due to historic overt discriminatory practices (redlining, for example, in the USA). Therefore, any historical calibration of the performance model would immediately suffer from an omission-in-variables (OIV) bias. To account for this we include the protected class variables, P_1, P_2, \dots, P_M and

modify our profit formula to include a discriminatory factor $D = D(P_1, P_2, \dots, P_M)$. This leads us to modify the approval–denial criteria (1) to the following:

$$\begin{cases} \text{approve } \pi + D > \pi^*, \\ \text{deny } \pi + D \leq \pi^*. \end{cases}$$

In regression form this reduces to estimating the combined coefficients of the modified performance equation

$$\pi = -D + \sum_{k=1}^n \beta_k X_k + \varepsilon = \alpha - \sum_{k=1}^m \lambda_k P_k + \sum_{k=1}^n \beta_k X_k + \varepsilon. \quad (2)$$

We note first that Eq. (2) implies the profit equation takes into account characteristics of the protected classes, for example, race, gender, and so on. From a profit perspective this may be true and in fact necessary due to historical discriminatory practices leading to inequities in education, compensation, or even job retention. In fact, current discriminatory practices may exist which will impact the ability of a protected class to repay a loan. However, under ECOA and FaHA, banks are not allowed to use protected class information in their decision processes related to loan origination. This would be disparate treatment. Therefore, in assessing the approval processes for adherence to fair-lending practices, the regression Eq. (2) is used to assess whether the coefficients of the protected class characteristics are significantly different from zero.

The approach just outlined is now typically used by commercial banks and regulatory agencies to identify disparate impact or disparate treatment in lending practices. But there are a number of practical and theoretical difficulties with the approach. As noted earlier, there may be any number of relevant variables that determine the credit quality of the borrower. If those variables are omitted in the regression equation, then their impact may bias one or more of the protected class coefficients. This is one type of OIV problem. There are more subtle types of OIV problems, such as unobservable variables that influence the outcome of the lending process that are difficult to assess, whose omission could lead to correlation between the error term and the outcome variable (approval or denial), leading to coefficient bias.

At a more fundamental level, one can question the appropriateness of regression to analyze the problem, as regression analyses are meant to adjust for small imbalances of the covariates between control and treatment groups in randomized designs. However, this problem is not dealing with a randomized design, as race, gender, and other protected classes cannot be randomized. Recently the work of Berkane [6] has attempted to address some of these theoretical problems using a different type of classification analysis with some success.

Analysis of lending practices for disparate impact or disparate treatment is a difficult and important problem facing all commercial banks as well as the agencies that regulate them. The industry practice is evolving rapidly as the consequences of unfair lending practices become more severe.

FINANCIAL CRIMES RISK

There are many slightly nuanced definitions of financial crimes. However, for our purposes we shall define financial crimes as crimes against customers, the commercial bank or leveraging the financial system to facilitate a crime. To go along with the many definitions of financial crimes there are a number of types of financial crime. These can be broadly classified into at least the following three categories:

- money laundering;
- fraud;
- tax avoidance.

This list is neither mutually exclusive nor intended to be exhaustive, as one type of financial crime may necessarily involve many elements. For example, money laundering will typically involve some type of fraud. The more seasoned reader may believe we have omitted customer due diligence (CDD), know your customer (KYC), terrorist financing, cyber security, and watch/sanctions list violations. However, the omission was somewhat intentional as the sorts of activities that go into monitoring these types of potential problems are typically covered by the techniques to address the three categories outlined above. Moreover, a failure to address one of these omitted categories is typically coupled with one of the categories we have listed. For example, transacting improperly with a politically exposed person (PEP) is typically part of a money laundering, fraud or even a terrorist financing investigation. Therefore, this list is

essentially representative of the types of financial crimes risk facing most financial institutions today.

The Bank Secrecy Act of 1970 (or BSA) requires financial institutions in the USA to assist US government agencies to detect and prevent financial crimes. Specifically, the act requires financial institutions to keep records of cash purchases of negotiable instruments, and file reports of cash purchases of these negotiable instruments of more than \$10,000 (daily aggregate amount), and to report suspicious activity that might signify money laundering, tax evasion, or other criminal activities.⁴ These reports are commonly known as suspicious activity reports or SARs and have become the cornerstone of investigations into criminal activity. In 2013 alone more than 1.6 million SARs were filed according to the Treasury Department's Financial Crimes Enforcement Network.

There are many potential paths leading to a SARs filing. It could be a suspicious deposit at a teller or a recurrent wire transfer from an overseas account. However, given the number of filings, it should be apparent that the amount of data that must be analyzed to produce a single filing is enormous. However, the cost of lax oversight in the area of anti-money laundering, fraud, or detecting tax avoidance can be severe as demonstrated by several recent high profile settlements below [14]:

- 2.6 billion—Credit Suisse AG—May 2014: Credit Suisse Group became the first financial institution in more than a decade to plead guilty to a crime when the Swiss bank admitted last month that it conspired to aid tax evasion and agreed to pay \$2.6 billion to settle a long-running probe by the US Justice Department.
- \$1.9 billion—HSBC Holdings—2012: HSBC agreed to pay \$1.9 billion to US authorities over deficiencies in its anti-money laundering controls. US officials hailed the settlement as the largest penalty ever under the Bank Secrecy Act. The agreement between the USA and HSBC also represented the third time since 2003 the bank agreed to US orders to cease lax conduct and correct failed policies.

As we have already noted, the detection of money laundering, fraud, and tax evasion typically involve the analysis of massive data sets. For instance, looking through hundreds of thousands if not millions of transactions to detect financial crime candidates that will then require additional analysis. Broadly speaking the techniques to perform these analyses

fall into two broad categories, supervised methods and unsupervised methods, respectively.

Outlier detection is a common form of unsupervised technique, while classification analyses like discriminant analysis, logistic regression, Bayesian belief networks, and decision trees would fall under the supervised learning methods. [17] provides a good overview of various classification schemes of both financial crimes as well as the techniques to analyze them.

To give an idea of the complexity of detecting financial crimes and the techniques used we focus on one particular type of fraud, credit card fraud, and briefly delve into some of the complexities. Credit card fraud cost banks billions of dollars annually [9, 10], and this is above the costs associated with the reputational damage once credit card fraud is identified.

Statistical learning approaches have become common in recent years to approach credit card fraud detection. These approaches fall under the supervised learning methods and have progressed greatly since their early use in the 1990s with neural networks. The statistical learning approach we review here is the support vector machines (SVMs) algorithm and the presentation follows [7] closely.

The SVMs method is a binary classification method that essentially embeds the classification features into a high-dimensional space and finds the hyper-plane which separates the two classes, fraudulent transactions and legitimate transactions, respectively. Due to the embedding in a high-dimensional space, the optimization process is linear. Moreover, the risk of overfitting, which exists for most neural network-like schemes, is minimized by finding the hyper-plane with maximal margin of separation between the two classes. Mathematically the problem can be described as the following quadratic programming problem:

Maximize

$$w(\alpha) = \sum_{k=1}^m \alpha_k - \sum_{j,k=1}^m \alpha_k \alpha_j \gamma_k \gamma_j k(x_k, x_j) \tag{3}$$

subject to

$$0 \leq \alpha_k \leq \frac{C}{m}, (k = 1, \dots, m), \tag{4}$$

$$\sum_{k=1}^m \alpha_k \gamma_k = 0, \quad (5)$$

where x_k , $k = 1, 2, \dots, m$, are the training data describing the credit card transactions⁵ which we collectively denote by X , k is a kernel function mapping $X \times X$ into an m dimensional space H . C is the cost parameter and represents a penalty for misclassifying the data while γ_k are the classification labels for the data points (i.e. one or zero, depending on whether x_k is a fraudulent or legitimate transaction).

The solution to (3), (4), and (5) provides the (dual) classification function:

$$\sum_{k=1}^m \alpha_k \gamma_k k(x_k, x) + b = 0. \quad (6)$$

There are several aspects of this problem which are practically and theoretically challenging. First, due to the high dimensionality the solution of the programming problem is computationally difficult, though there are iterative approaches, see [19] for example, that can scale large problems for SVM implementations. Second, the choice of the kernel function and the cost parameter can greatly influence the outcome of the classification function and its effectiveness. The cost parameter is often difficult to estimate and only experimenting with choices of k and reviewing results is currently available. Last, and probably most pressing, there is no clear-cut best measure of model performance. The industry has used the receiver operating characteristic (ROC) and the area under the ROC curve (AUC) as well as functions of AUC, like the Gini coefficient (see [7] for a fuller discussion), but each has weaknesses when encountering imbalanced data; that is, data where the occurrence of one class, for example fraud, has a very low probability of occurring. A frequently used example (see [8] for instance) to describe this difficulty as it applies to accuracy as performance measure is the following: Suppose in our credit card example, the probability of correctly detecting legitimate activity as legitimate is $\frac{99}{100}$ and the probability correctly detecting fraudulent activity as fraudulent is $\frac{99}{100}$. This would appear to be a very accurate detection system. However, now suppose we have a data imbalance. For example, we know that one in 1000 records are fraudulent. Then on an average in a sample of 100 records flagged

as fraudulent we would expect only nine to really be fraudulent. But this would require the commercial bank to review 100 records to possibly zero in on the nine true offenders. Imagine the cost if this were 1000 flagged records or hundreds of thousands like a typical commercial bank would have with SARs records. Data imbalance requires thoughtful choices of the various parameters in the modeling effort as well as careful choices of the model's performance measurement. As of this writing, these and other topics related to the quantitative analysis of financial crimes remain fertile ground for research.

MODEL RISK

Models are used pervasively throughout all commercial banks. In fact, this chapter has discussed just a small subset of the types of models used daily in most banks throughout the world. Furthermore, with the ability to store and manipulate ever larger data sets, more computing power and the increased packaging of models into easy to use software, the upward trend in model use in the banking industry is likely to continue unabated. But with model use come model risks. This risk was highlighted with the notable model risk management failures prior to and during the financial crisis. The pre-crisis pricing of CDOs using Gaussian copula models (see [18] for an in-depth discussion) or the models used by rating agencies to rate structured products are just two of many examples.

Though not the driving factor behind the financial crisis, the regulatory agencies realized that poor model risk management was likely a contributing factor to the crisis and that guiding principles for the proper management of model risk were needed. This framework was provided in the form of a joint agency bulletin [20], known typically as "SR-11-7" or "2011-12" in the banking industry.⁶ We shall simply refer to it as the agency guidance.

The agency guidance defined a model as a "quantitative method, system, or approach that applies statistical, economic, financial, or mathematical theories, techniques, and assumptions to process input data into quantitative estimates". Furthermore, it stated "that a model consists of three components: an information input component, which delivers assumptions and data to the model; a processing component, which transforms inputs into estimates; and a reporting component, which translates the estimates into useful business information". The document goes on

to define model risk as “the potential for adverse consequences from decisions based on incorrect or misused model outputs or reports”.

The regulatory definition of model, for all practical purposes, expanded the scope of model risk. Early attempts at defining and measuring model risk primarily focused on the “transformation component” of the model (the “quant stuff”) and largely ignored the input and output components. Moreover, most of the model risk work pre-crisis focused on risks associated with derivative pricing models ([11, 12, 21]), though the largest risk in most commercial banks comes from credit and its approval and ongoing monitoring processes, which are increasingly model driven.

Fundamentally, model risk can be broken down into three categories—*inherent, residual, and aggregate risks*. These risks can be described as follows:

- **Inherent Risk**
 - All models are simplifications of real-world phenomena.
 - This simplification process leads to risk of omitting relevant features of the process one wishes to model.
 - Some inherent risks can be mitigated or reduced while others cannot or may not even be known at the time of model development.
- **Residual Risk**
 - The risk that remains after mitigating all known inherent risks that can be managed or are deemed cost effective to manage.
 - Accepted risk for using a particular model.
- **Aggregate Risk**
 - The risk to the firm from all model residual risks.
 - Not simply an additive concept as there will likely be complex dependencies between models either directly or through their residual risks.

Within this framework, most model risk work has focused on analyzing inherent risk and has attempted to measure model misspecification within a well-defined class of models in order to make the problem tractable. Bayesian model averaging is one such approach that has been explored extensively ([15, 22]). Cont [11] refers to this type of model misspecification risk as “model uncertainty” and asks the fundamental questions relate to it:

- How sensitive is the value of a given derivative to the choice of pricing model?
- Are some instruments more model-sensitive than others?
- How large is the model uncertainty of a portfolio compared with its market risk?
- Can one capitalize for “model risk” in the same way as one capitalizes for market and credit risk?

Cont approaches the problem by looking at the payoff V of a derivative or a portfolio of derivatives which all have well defined values for pricing models \mathcal{Q} , contained in a class of pricing models \mathcal{Q} . He then defines model uncertainty (within the class of models) as

$$\mu_{\mathcal{Q}} = \sup_{\mathcal{Q} \in \mathcal{Q}} E^{\mathcal{Q}}[V] - \inf_{\mathcal{Q} \in \mathcal{Q}} E^{\mathcal{Q}}[V]$$

where expectation is with respect to the risk-neutral measure. Cont goes on to show that $\mu_{\mathcal{Q}}$ is a coherent measure of model uncertainty⁷ and for a fixed model \mathcal{Q} defines the model risk ratio

$$MR(V) = \frac{\mu_{\mathcal{Q}}(V)}{E^{\mathcal{Q}}[V]}.$$

This is essentially the ratio of the range of potential values of V within the class of admissible functions to the value of V under the proposed model.

Glasserman and Xu [13] take a similar approach. Denoting by X the stochastic elements of the model, the risk neutral value of the derivative under the presumed distribution of X , given by f , is once again $E[V(X)]$, where the dependence on f is implicit. They then allow alternative distributions of X (alternative models) denoted by \tilde{f} and proceed to solve the constrained maximum and minimum problems to find the range of model risk:

Solve

$$\inf_m E[m(X)V(X)] \text{ and } \sup_m E[m(X)V(X)]$$

subject to

$$D(m) = E[m \log m] \leq \eta, \text{ where } m(X) = \frac{\tilde{f}}{f}$$

Essentially, the solution distributions to this max/min problem are restricted to a relative entropy distance η from the presumed base distribution f . At this point one can create model risk ratios like Cont. One of the drawbacks of the Glasserman and Xu approach is that the class of models under consideration is not necessarily calibrated to a base set of instruments (e.g. European options, or swaptions), which is a desirable if not required feature for many derivatives models.

Abasto and Kust [1] take a novel approach and define a Model “01”, in the spirit of risk sensitivities used by market risk practitioners, like DV01,⁸ by employing weighted Monte Carlo (WMC) techniques. Their technique allows the calculation of a Model 01, while ensuring that the target model is calibrated to a set of calibration instruments $\{C_j\}$, $j = 1, \dots, M$. Mathematically, if X_i , $i = 1, \dots, N$ are the realizations of the stochastic parameters driving the value of the derivative, V , and p_i , $i = 1, \dots, N$ are probabilities of the i th stochastic event being realized then an estimate of the value of the derivative is given by

$$E^p[V(X)] = \sum_{i=1}^N p_i V(X_i) = \sum_{i=1}^N p_i V_i.$$

Abasto and Kust then solve the constrained minimization problem⁹:

$$\min_p D(p||p_0)$$

subject to

$$\begin{aligned} \sum_{i=1}^N p_i V(X_i) &= V(1 + \alpha), \\ \sum_{i=1}^N p_i g_{ij} &= C_j, \quad j = 1, \dots, M, \\ \sum_{i=1}^N p_i &= 1. \end{aligned}$$

Here $D(p | p_0)$ is the Hellinger distance between p and the target model p_0 , g_{ij} is the payoff of the j th calibration instrument C_j under the i th scenario X_i , and α is, initially, some fixed small increment.

Finally, they use the fact that the square root vector, $(\sqrt{p_1}, \sqrt{p_2}, \dots, \sqrt{p_N}) = P$ for our probabilities resides on a unit hyper-sphere so they fix a small angle φ^* (say = .01) and find two models p^- and p^+ corresponding to small increments $\alpha < 0$ and $\alpha > 0$. These two models lie in an “01” normalized distance of the target model in the following sense:

$$\text{Model 01} = E^{p^+} [V] - E^{p^-} [V],$$

subject to

$$P^+, P^- = \cos(\varphi^*).$$

As noted, all of these techniques are designed to assess inherent model risk, not residual or aggregate model risk; however, they all assess inherent risk within a well-defined class of admissible models. Therefore the measure of risk depends greatly on the family of models chosen. In fact, in some of the approaches, ensuring that all models in the admissible class are calibrated to a set of base instruments is construed as eliminating inherent risk and only leaving residual risk. This is not a view shared by the author.

A more serious weakness of most of the techniques is their heavy reliance on risk-neutral pricing apparatus. They are, therefore, very well suited for analyzing derivative model risk but are not readily amenable to assessing the risk of the broad array of models that are widespread throughout banks, like credit scoring models, in particular. This is not a weakness of the Bayesian averaging approaches.

Finally we note that methods for addressing aggregate model risk are still in their early stages. At its core, the problem of measuring aggregate model risk requires understanding and quantifying complex dependencies across a myriad of models. This is a complex problem, with the most basic attempts trying to assess sensitivities to common variables or parameters (like “01”s) across similar models.

CONCLUSION

We have given a flavor of the types of pressing quantitative problems facing the commercial banking industry in the post-crisis financial environment. This list is far from exhaustive and in the limited space available we could only scratch the surface of these nuanced and complex issues. There are many other quantitative problems facing the industry which are equally rich in their complexity and importance and this fact leads the author to believe that the golden age of quantitative finance is not in its twilight but stretches before us on the horizon.

NOTES

1. Financial Accounting Standards (FASB) 157 in the USA <http://www.fasb.org/summary/stsum157.shtml> . International Accounting Standards (IAS) 39 http://ec.europa.eu/internal_market/accounting/docs/consolidated/ias39_en.pdf.
2. To be exact, the uncollateralized derivative asset which is hedged with a collateralized derivative instrument requires funding as the hedging liability will require collateral funding. Conversely, an uncollateralized derivative liability will benefit from collateral inflows from the hedging asset.
3. NINJA loans are lightly documented loans which have been viewed as predatory. The NINJA acronym comes from No Income, No Jobs no Assets.
4. https://www.fincen.gov/statutes_regs/bsa/.
5. More accurately, x_i are the derived attributes of the training data. For each credit card transaction for instance, a set of attributes will be aggregated like the number of transactions at a particular location or the average size of transactions over the last three months.
6. The Federal Reserve board-issued document is known as SR-11-7 while the OCC document is known as 2011-12.
7. (1) Coherent in the sense that model uncertainty reduces to uncertainty in market value (bid-ask spread), (2) a derivative that can be replicated in a model-free way has no uncertainty, (3) diversification and hedging with traded options decrease uncertainty.
8. DV01 = Dollar Value of a basis point decrease in “interest rates”.
9. Abasto and Kust actually perform the minimization relative to the equal weight probability measure $p_i = 1/N$ for all i in the Hellinger distance and demonstrate that the end results are identical.

REFERENCES

1. Abasto, Damian and Kust, Mark P., “Model01: Quantifying the Risk of Incremental Model Changes”, September 5, 2014. Available at SSRN: <http://ssrn.com/abstract=2492256> or [10.2139/ssrn.2492256](http://dx.doi.org/10.2139/ssrn.2492256).
2. Basel Committee on Banking Supervision (BCBS), “International Convergence of Capital Measurement and Capital Standards A Revised Framework Comprehensive Version”, June 2006 <http://www.bis.org/publ/bcbs128.pdf>.
3. Basel Committee on Banking Supervision (BCBS), “Observed range of practice in key elements of Advanced Measurement Approaches (AMA)”, July 2009 <http://www.bis.org/publ/bcbs160b.pdf>.
4. Basel Committee on Banking Supervision (BCBS), “Operational Risk – Supervisory Guidelines for the Advanced Measurement Approaches”, June 2011 <http://www.bis.org/publ/bcbs196.pdf>.
5. Basel Committee on Banking Supervision (BCBS), “Basel III: A Global Regulatory Framework for more Resilient Banks and Banking Systems (Revised)”, June 2011 <http://www.bis.org/publ/bcbs189.pdf>.
6. Berkane, Maia. Wells Fargo & Co., Private Communication, March 2015.
7. Bhattacharyya, Siddhartha., Jah, Sanjeev., Tharakunnel, Kurian., Westland, J. Christopher., “Data Mining for Credit Card Fraud: A Comparative Study”, *Decision Support Systems*, 50(3), February 2011, 602–613.
8. Bolton, Richard J., and David J. Hand, “Statistical Fraud Detection: A Review”, *Statistical Science*, 17(3), 2002, 235–249.
9. Chan, P.K., Fan, W., Prodromidis, A.L., Stolfo, S.J., “Distributed Data Mining in Credit Card Fraud Detection”, *Data Mining*, (November/December), 1999, 67–74.
10. Chen, R.C., Chen, T.S., Lin, C.C., “A New Binary Support Vector System for Increasing Detection Rate of Credit Card Fraud”, *International Journal of Pattern Recognition*, 20(2), (2006), 227–239.
11. Cont, Rama. “Model Uncertainty and Its Impact on the Pricing of Derivative Instruments”, *Mathematical Finance*, 16, July 2006.
12. Derman, E, “Model Risk”, *Risk*, 9(5), 139–145, 1996
13. Glasserman, P., and Xu, X. “Robust Risk Measurement and Model Risk”, *Journal of Quantitative Finance*, 2013.
14. Grocer, Stephen, “A List of the Biggest Bank Settlements.” *The Wall Street Journal*, 23 June 2014.
15. Hoeting, J., Madigan, A.D, Raftery, A. E, Volinsky, C. T., “Bayesian Model Averaging: A Tutorial”, *Statistical Science* 14(4), 382–417.
16. Jorion, Philippe. GARP (Global Association of Risk Professionals) (2009-06-08). *Financial Risk Manager Handbook* (Wiley Finance) Wiley. Kindle Edition.

17. Kenyon, Chris., Stamm, Roland., “*Discounting, Libor, CVA and Funding: Interest Rate and Credit Pricing*”. Houndmills, Basingstoke: Palgrave Macmillan, 2012. Print.
18. Morini, Massimo., “*Understanding and Managing Model Risk: A Practical Guide for Quants, Traders and Validators*”, Hoboken: Wiley 2011
19. Ngai, W.T., Hu, Yong., Wong, Y. H., Chen, Yijun., Sun, Xin. “The Application of Data Mining Techniques in Financial Fraud Detection: A Classification Framework and an Academic Review of Literature.” *Decision Support Systems* 50, 3 (February 2011), 559–569.
20. OCC Bulletin 2011-12/Federal Reserve Bulletin SR 11-7, “Supervisory Guidance on Model Risk Management”, April 4, 2011. <http://www.occ.gov/news-issuances/bulletins/2011/bulletin-2011-12a.pdf>.
21. Platt, J.C., “Fast Training of Support Vector Machines Using Sequential Minimal Optimization”, in: B. Scholkopf, C.J.C. Burges, A.J. Smola (Eds.), *Advances in Kernel Methods—Support Vector Learning*, MIT Press, Cambridge, MA, 1998, 185–208.
22. Raftery, A. E., “Bayesian Model Selection in Structural Equation Models”, in *Testing Structural Equation Models*, K. Bollen and J. Long, eds. Newbury Park, CA: Sage, 1993 163–180.
23. Rebonato, R., “Theory and Practice of Model Risk Management”, in *Modern Risk Management: A History*; *RiskWaters Group*, 2003. 223–248.
24. Ross, S. L., Yinger, J., “*The Color of Credit: Mortgage Discrimination, Research Methodology, and Fair-Lending Enforcement*”. Cambridge, Mass: MIT Press, 2002.