OFR WORKING PAPER
OFFICE OF FINANCIAL RESEARCH

15-16 | August 19, 2015

Bounding Wrong-Way Risk in Measuring Counterparty Risk

Paul Glasserman

Office of Financial Research and Columbia University
paul.glasserman@treasury.gov

Linan Yang

Columbia University
ly2220@columbia.edu

Bounding Wrong-Way Risk in Measuring Counterparty Risk

Paul Glasserman* and Linan Yang†

Abstract

Counterparty risk measurement integrates two sources of risk: *market risk*, which determines the size of a firm's exposure to a counterparty, and *credit risk*, which reflects the likelihood that the counterparty will default on its obligations. Wrong-way risk refers to the possibility that a counterparty's default risk increases with the market value of the exposure. We investigate the potential impact of wrong-way risk in calculating a credit valuation adjustment (CVA) to a derivatives portfolio: CVA has become a standard tool for pricing counterparty risk and setting associated capital requirements. We present a method, introduced in our earlier work, for bounding the impact of wrong-way risk on CVA. The method holds fixed marginal models for market and credit risk while varying the dependence between them. Given simulated paths of the two models, we solve a linear program to find the worst-case CVA resulting from wrong-way risk. The worst case can be overly conservative, so we extend the procedure by penalizing deviations of the joint model from a baseline model. By varying the penalty for deviations, we can sweep out the full range of possible CVA values for different degrees of wrong-way risk. Our method addresses an important source of model risk in counterparty risk measurement.

Keywords: credit valuation adjustment, counterparty credit risk, wrong-way risk, iterative proportional fitting process (IPFP).

1 Introduction

Counterparty risk has taken on heightened importance since the failures of major derivatives dealers Bear Stearns, Lehman Brothers, and AIG Financial Products in 2008. Basel III includes a new capital charge for counterparty risk, which is among the largest changes to capital requirements for banks with major derivatives businesses (BCBS [1]). Accurate measurement of counterparty risk is essential to financial stability, yet it presents significant modeling and computational challenges for industry participants.

Counterparty risk combines market exposure and credit risk. Market factors determine the size of a firm's exposure to a counterparty, and credit risk determines the likelihood that the

*Office of Financial Research and Columbia Business School. This paper was produced while Paul Glasserman was a consultant to the Office of Financial Research. Email: paul.glasserman@treasury.gov.

†Industrial Engineering and Operations Research Department, Columbia University. Email: ly2220@columbia.edu.

counterparty will default, turning the exposure into a loss. The proper integration of these two sources of uncertainty is among the major challenges in counterparty risk measurement.

Wrong-way risk refers to the possibility that the two sources of risk move together, so that the market exposure increases just as the counterparty's risk of default increases. Wrong-way risk arises, for example, if one bank sells credit default swap protection on another bank with a similar profile. The value of the credit protection increases when the second bank faces financial difficulties; this is likely to be a scenario in which the bank that sold the protection is also at greater risk of default. In practice, the sources and nature of wrong-way risk are often less obvious.

The standard tool for quantifying counterparty risk in derivatives markets is the credit valuation adjustment (CVA). The CVA on a portfolio of derivatives is an adjustment to the market value of the portfolio to correct for the credit quality of the counterparty, relative to what the same trades would be worth with a hypothetical default-free counterparty. The CVA for a portfolio of derivatives will generally increase with wrong-way risk, but the correct degree of wrong-way risk is difficult to estimate in practice.

In Glasserman and Yang [8], we introduced a method for bounding wrong-way risk — that is, for finding the largest CVA that is consistent with fixed models for market exposure and credit risk, letting the dependence between market factors and credit risk vary. The present article provides a simplified exposition of the method in our earlier paper and additional examples. Our earlier paper also contains more extensive references to the literature.

Our approach builds on a standard simulation framework for CVA calculation: paths of underlying market factors are simulated over time; a portfolio is revalued (often using approximations) at fixed dates along each path of the market model; and the counterparty's time to default is either simulated from a credit risk model or extracted from a credit curve. Given a set of paths of market exposures and the distribution of the time to default, we find the worst-case CVA by solving a linear programming problem. The linear program finds the assignment of default times to paths resulting in the largest possible CVA, given the constraints on the default time distribution and the set of paths simulated from the market model. As a byproduct, the dual variables associated with constraints on the marginal default time distribution provide sensitivities of the worst-case CVA to the default probabilities.

A strength of this approach is that it yields the largest possible CVA value consistent with given models for market factors and credit risk. Because it finds the worst-case wrong-way risk, this approach can also be too conservative. In Glasserman and Yang [8], we extend the method by penalizing deviations from a baseline model and finding the resulting tempered CVA. With a large penalty for deviations, the tempered CVA will be close to the CVA under the baseline model;

with a small penalty, the tempered CVA will be close to the worst case found through the linear programming solution. A natural choice for the baseline model is to take market factors and credit risk to be independent of each other.[1] By varying the penalty for deviations, we can sweep out the full range of potential CVA values from the independent case to the worst possible wrong-way risk.

The penalized problem can no longer be solved through linear programming, but we formulate it as a tractable convex optimization problem. The special structure of the problem leads to a convenient solution through iterative rescaling of the rows and columns of a matrix.

Models that explicitly describe dependence between market factors and credit risk in CVA calculation include Brigo, Capponi, and Pallavicini [2], Crépey [5], Ghamami and Goldberg [7], Hull and White [10], and Rosen and Saunders [13]; see Brigo, Morini, and Pallavicini [3] for an extensive overview of modeling approaches. See Canabarro and Duffie [4] and Gregory [9] for additional background on counterparty risk.

2 Problem Formulation and Worst-Case CVA

To help fix ideas, we start with an example. Consider a T-year foreign exchange forward contract between a U.S. bank, which receives U.S. dollar payments, and a foreign bank which receives its local currency. The contract has forward exchange rate K and notional size S. If the foreign currency weakens against the dollar, the foreign bank's credit quality is likely to deteriorate with its currency, just as the U.S. bank's exposure increases, so this transaction exhibits evident wrong-way risk.

Let U_t be the exchange rate, measured as the number of units of the foreign currency paid in exchange for one U.S. dollar at time t. Assume this exchange rate follows an Ornstein-Uhlenbeck process,

$$dU_t = \kappa(\bar{U} - U_t)dt + \sigma dW_t,$$

where $\bar{U}$ is the level toward which the exchange rate mean-reverts, and W_t is a standard Brownian motion.

CVA measures the discounted expected loss of a portfolio at the counterparty's default, so its calculation involves the default time of the counterparty, denoted by τ, and the discounted value of the portfolio exposure with this counterparty at the time of its default, denoted by $V(\tau)$. We assume that $V(\tau)$ accounts for all netting and collateral agreements, and let R denote the recovery rate of the counterparty. Only the positive part of exposure $V^+(\tau)$ represents a loss at default, so

[1] The Basel III standarized approach for CVA assumes independence and then multiplies the result by a factor of 1.4.

the CVA for a fixed time horizon T is the expectation

$$\mathrm{CVA} = \mathbb{E}[(1-R)V^{+}(\tau)\mathbf{1}\{\tau \leq T\}], \tag{2.1}$$

given a joint law for the default time τ and the exposure $V^{+}(\tau)$.

CVA is usually calculated over a fixed set of dates, so set $0 = t_0 < t_1 < \cdots < t_d = T$, and $t_{d+1} = \infty$. The dates can be monthly or quarterly, or the payment dates of the underlying contracts. We limit τ to values in $\{t_1, \ldots, t_d, t_{d+1}\}$ and let q_j, $j = 1, \ldots, d+1$, denote the probability that default occurs at t_k, or, more precisely that it occurs in the interval $(t_{k-1}, t_k]$. The default time distribution can be extracted from credit default swap spreads of the counterparty or from a more extensive credit risk model, such as a stochastic intensity model.

Simulation of market factors is used to generate exposure paths $(V^{+}(t_1), \ldots, V^{+}(t_d))$, and the calculation of each $V(t_j)$ needs to account for netting and collateral agreements and recovery rates if the counterparty were to default. In our foreign exchange example,

$$V(t_j) = e^{-\delta t_j} \cdot \mathbb{E}[e^{-\delta(T-t_j)} S(U_T - K)/U_T \,|\, U_{t_j}],$$

where δ is the discount rate and R is the recovery rate. The expectation gives the expected exposure of the contract at time t_j, and this value is discounted to t_0 and adjusted for partial recovery. The market factor model (in this example the exchange rate dynamics) implicitly determines the law of the positive exposure path $(V^{+}(t_1), \ldots, V^{+}(t_d))$, and we denote this law by a probability measure p on $\mathbb{R}^d$.

Let X denote this vector of positive exposures adjusted by recovery at the specified dates, and let Y be a vector of default indicators,

$$X = ((1-R)V^{+}(t_1), \ldots, (1-R)V^{+}(t_d)) \;\text{ and }\; Y = (\mathbf{1}\{\tau = t_1\}, \ldots, \mathbf{1}\{\tau = t_d\});$$

then CVA would reduce to the expectation of the inner product

$$< X, Y > = (1-R)\sum_{j=1}^{d} V^{+}(t_j)\mathbf{1}\{\tau = t_j\} = (1-R)V^{+}(\tau)\mathbf{1}\{\tau \leq T\},$$

if the joint law for X and Y were known.

However the joint law is in general unavailable and difficult to find because of limited data on the dependence between market factors and credit risk. With the marginals fixed, we need to assign a joint probability between X and Y to calculate CVA. As an upper bound, we seek to evaluate the *worst-case* CVA, defined by

$$\mathrm{CVA}_* := \sup_{\mu \in \Pi(p,q)} \int_{\mathbb{R}^d \times \mathbb{R}^d} < x, y > d\mu(x,y), \tag{2.2}$$

where $\Pi(p,q)$ denotes the set of probability measures on $\mathbb{R}^d \times \mathbb{R}^d$ with marginals p and q. Our procedure estimates (2.2) by solving a linear programming problem using the simulation results for X and Y. It builds on the following standard simulation of market exposures:

(1) Simulate path of all relevant market factors jointly. In our example, this requires simulating paths of the foreign exchange rate, $(U^i(t_1),\ldots,U^i(t_d))$ for $i=1,\ldots,N$.

(2) At each date t_j, for each market scenario i, evaluate the expected exposure of the portfolio based on the market factors; discount and adjust by netting rules and take positive parts to get $(V^i(t_1)^+,\ldots,V^i(t_d)^+)$, $i=1,\ldots,N$.

The key point is that we assume the availability of independent copies of exposure paths and use them as input to our procedure.

Let $X_1,\ldots,X_N$ be N independent copies of X, and let $Y_1,\ldots,Y_N$ be N independent copies of Y. We can assume that p has no atoms, and we can then identify the empirical measure p_N on $\mathbb{R}^d$ with the uniform distribution on the set $\{X_1,\ldots,X_N\}$, where

$$p_N(\cdot) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\{X_i \in \cdot\}. \tag{2.3}$$

In particular, each path X_i gets equal weight $1/N$. On the other hand, Y is supported on the finite set $\{y_1,\ldots,y_{d+1}\}$, with $y_1=(1,0,\ldots,0),\ldots,y_d=(0,0,\ldots,1)$, and $y_{d+1}=(0,\ldots,0)$, and each y_j has probability $q(y_j)$. These probabilities may be known or estimated from simulation of N independent copies $Y_1,\ldots,Y_N$ of Y, in which case we denote the empirical frequency of each y_j by $q_N(y_j)$, defined as

$$q_N(\cdot) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\{Y_i \in \cdot\}. \tag{2.4}$$

With simulations of X_i and Y_i for $i=1,\ldots,N$, and empirical marginals p_N and q_N, we can find the worst-case joint mass function P^N_{ij} on the set of pairs $\{(X_i,y_j),\, i=1,\ldots,N,\, j=1,\ldots,d+1\}$. Let $\Pi(p_N,q_N)$ be the set of joint mass functions with marginals p_N and q_N. We estimate (2.2) using

$$\widehat{\mathrm{CVA}}_* = \max_{P^N\in\Pi(p_N,q_N)} \sum_{i=1}^{N}\sum_{j=1}^{d+1} P^N_{ij} < X_i, y_j > .$$

Letting $C_{ij} =< X_i, y_j >$ we can rewrite this estimate of the worst-case CVA as a linear programming problem:

$$\max_{\{P_{ij}\}} \quad \sum_{i=1}^{N}\sum_{j=1}^{d+1} C_{ij}P_{ij}, \tag{2.5}$$

$$\text{subject to} \quad \sum_{j=1}^{d+1} P_{ij} = 1/N,\ i=1,...,N, \tag{2.6}$$

$$\sum_{i=1}^{N} P_{ij} = q_N(y_j),\ j=1,...,d+1 \quad \text{and} \tag{2.7}$$

$$P_{ij} \geq 0,\ i=1,...,N,\ j=1,...,d+1, \tag{2.8}$$

Constraint (2.6) ensures that the paths $X_1,\dots,X_N$ of market factors get equal weight; constraint (2.7) ensures that the default-time distribution in the joint model has the correct marginal distribution. In our running example, we have

$$C_{ij} = (1-R)(V^i(t_j))^+ = (1-R)\cdot e^{-\delta t_j}\cdot \mathbb{E}^+[e^{-\delta(T-t_j)}S(U_T-K)/U_T \,|\, U^i_{t_j}].$$

A nice feature of this problem is that it has the structure of a transportation problem, so that efficient algorithms are available, including a strongly polynomial algorithm; see Kleinschmidt and Schannath [11].

In Glasserman and Yang [8], we have established the consistency of this estimator. In other words, we show there that the optimal value of the linear programming problem converges to CVA$_*$ in (2.2) as the sample size N increases.

The sensitivities of the worst-case CVA to the constraints on the marginal distributions are available from the solution of the dual to the linear program. Let a_i and b_j be dual variables associated with constraints (2.6) and (2.7), respectively. The dual problem is

$$\begin{aligned} \min_{a\in\mathbb{R}^N, b\in\mathbb{R}^{d+1}} \quad & \sum_{i=1}^{N} a_i/N + \sum_{j=1}^{d+1} b_j q_N(y_j) \\ \text{subject to} \quad & a_i + b_j \geq C_{ij}, i=1,...,N,\ j=1,...,d. \end{aligned}$$

Let (a^*, b^*) be the dual optimal solution. Consider a vector of perturbations $(\Delta q_1,\dots,\Delta q_{d+1})$ to the mass function q_N with components that sum to zero (so that the perturbed probabilities sum to one). Suppose these perturbations are sufficiently small to leave the dual solution unchanged. Then

$$\Delta\widehat{\mathrm{CVA}}_* = \sum_{j=1}^{d+1} b^*_j \Delta q_j.$$

In particular, we can calculate the sensitivity of the worst-case CVA to a parallel shift in the credit curve by setting $\Delta q_j = \Delta$, $j=1,\dots,d$, and $\Delta q_{d+1} = -d\Delta$, for sufficiently small Δ.

3 Tempering the Worst Case

The worst-case joint distribution is useful because it provides a bound, but it is not likely to be achieved in practice. To address this point, we want to control the degree of wrong-way risk in the CVA calculation. We accomplish this by introducing a penalty term in the objective function that penalizes deviations of the joint probability from a baseline model. Putting greater weight on the penalty limits the degree of potential wrong-way risk and tempers the worst-case CVA.

3.1 Penalty Formulation

For the penalty, we need a notion of distance from one distribution to another. Among many possible choices, relative entropy, also known as Kullback-Leibler divergence, is particularly convenient. For probability measures P and F on a common measurable space with $F >> P$, the relative entropy of P to F is defined as

$$D(P|F) = \int \ln\left(\frac{dP}{dF}\right) dP.$$

Relative entropy is nonnegative, and $D(P|F) = 0$ only if $P = F$. As a function of P, relative entropy $D(P|F)$ is convex in P and leads to a convex optimization problem. The problem can be solved through a simple algorithm, known as the iterative proportional fitting procedure. Relative entropy is not symmetric in its arguments, but we do not see this as a drawback because we think of the baseline model as a favored benchmark.

To be concrete, here we take the baseline model to assume independence between market factors and credit risk. This choice is convenient, but it is not critical. Let ν denote the independent joint law of X and Y, meaning that

$$\nu(A \times B) = p(A)q(B),$$

for all measurable $A, B \subseteq \mathbb{R}^d$, and note that $\nu \in \Pi(p,q)$. Under this baseline model, we have the independent case CVA given by

$$\mathrm{CVA}_\nu = \int_{\mathbb{R}^d\times\mathbb{R}^d} < x, y > d\nu(x,y) = \int_{\mathbb{R}^d\times\mathbb{R}^d} < x, y > dp(x)dq(y).$$

In the penalty formulation, we seek a joint distribution that maximizes CVA with marginals p and q and at the same time penalizes the objective function by the relative entropy divergence of this joint distribution from the baseline model ν. The penalty formulation with parameter θ is as follows:

$$\sup_{\mu\in\Pi(p,q)} \int_{\mathbb{R}^d\times\mathbb{R}^d} < x, y > d\mu(x,y) - \frac{1}{\theta}\int \ln(\frac{d\mu}{d\nu})d\mu. \tag{3.1}$$

With the solution μ^* of (3.1), we get the tempered CVA

$$\mathrm{CVA}_\theta = \int_{\mathbb{R}^d \times \mathbb{R}^d} < x, y > d\mu^*(x, y). \tag{3.2}$$

For $\theta = 0$, the penalty becomes infinite unless $\mu = \nu$, in which case $\mathrm{CVA}_0 = \mathrm{CVA}_\nu$. For $\theta = \infty$, the penalty term vanishes and we recover the worst-case linear program in Section 2, and $\mathrm{CVA}_\infty = \mathrm{CVA}^*$. For $\theta \in (0, \infty)$, we find a tempered CVA interpolated between the independent case and the worst-case wrong-way risk. We will see that negative values of θ correspond to "right-way" risk in which the likelihood of default of the counterparty decreases with the market value of exposure.

As in Section 2, we use simulation samples of X and Y to form an optimization problem to estimate CVA_θ. Let F_{ij} denote the baseline joint probabilities; in the independent case, these are given by $F_{ij} = q_N(y_j)/N$, $i = 1, \ldots, N$, $j = 1, \ldots, d + 1$. The finite penalty problem is

$$\max_{\{P_{ij}\}} \quad \sum_{i=1}^{N} \sum_{j=1}^{d+1} C_{ij} P_{ij} - \frac{1}{\theta} \sum_{i=1}^{N} \sum_{j=1}^{d+1} P_{ij} \ln\left(\frac{P_{ij}}{F_{ij}}\right) \quad \text{subject to (2.6)-(2.8).} \tag{3.3}$$

With P^θ denoting the optimal solution to this problem, we estimate CVA_θ by

$$\widehat{\mathrm{CVA}}_\theta := \sum_{i=1}^{N} \sum_{j=1}^{d+1} C_{ij} P^\theta_{ij}.$$

In our earlier paper [8], we established the consistency of this estimator of tempered CVA, showing that it converges to CVA_θ as the sample size $N \to \infty$

3.2 Algorithm

A general-purpose optimization method can be applied to solve the penalty problem in (3.3) because it is a convex program. However, a simple algorithm is available for solving this particular problem through the iterative proportional fitting procedure (IPFP) because we chose relative entropy in the penalty term. This method dates back to Deming and Stephan [6] and has applications in many areas; see Pulkelsheim [12] for an overview.

The IPFP method takes as input data a nonnegative matrix and row and column marginals. The output is a matrix with the specified row and column sums equal to the input marginals. The output is the matrix closest to the input matrix, in the sense of relative entropy, among all matrices with the target row sums and column sums.

In applying this method in our setting, we start with an input matrix that represents an initial joint probability matrix weighted by the exposure at each possible default time. Then we iteratively fit this initial joint probability matrix to the target marginals of market factors and credit risk.

The procedure converges to a proper joint probability that satisfies the marginal constraints and also accounts for exposure weighting.

Define an $N \times (d+1)$ matrix M^θ with entries

$$M_{ij}^\theta = \frac{e^{\theta \cdot C_{ij}} \cdot F_{ij}}{\sum_{i=1}^N \sum_{j=1}^{d+1} e^{\theta \cdot C_{ij}} \cdot F_{ij}}.$$

The row index $i = 1, \ldots, N$, is the market exposure path index, and the column index $j = 1, \ldots, d+1$, is the default time index for t_j. As before, F_{ij} is the baseline model. Each $C_{ij} =< X_i, y_j >$ is the loss on exposure path i if the counterparty defaults at time t_j. The parameter θ is the penalty parameter in the penalty problem. The matrix M^θ emphasizes wrong-way risk because for $\theta > 0$ it puts greater weight on the combinations of exposure path and default time that produce larger losses.

The denominator of M_{ij}^θ normalizes all entries to sum to one, but this matrix does not satisfy the target marginals p_N and q_N. The IPFP algorithm iteratively repeats the following steps:

(a) For $i = 1, \ldots, N$, set $M_{ij} \leftarrow M_{ij} p_N(i) / \sum_{k=1}^{d+1} M_{ik}$ for $j = 1, \ldots, d+1$.

(b) For $j = 1, \ldots, d+1$, set $M_{ij} \leftarrow M_{ij} q_N(j) / \sum_{n=1}^{N} M_{nj}$ for $i = 1, \ldots, N$.

After step (a), the new matrix has row sums equal to the market marginal p_N, regardless of the marginals of the initial matix. After step (b), the new matrix has column sums equal to the credit risk marginal q_N, however the row sums will no longer necessarily match p_N. Let $\Phi(M)$ denote the operation of applying both steps (a) and (b) to M, and write $\Phi^{(n)}$ for the n-fold composition of Φ. We have shown in Glasserman and Yang [8] that as $n \to \infty$, $\Phi^{(n)}(M)$ converges to a limit P^θ with marginals p_N and q_N that solves the penalty problem (3.3). With $\theta < 0$, the limit solves the penalty problem with max replaced by min, corresponding to right-way risk in the sense that it minimizes the CVA subject to the marginal constraints and the penalty on deviations from the baseline model.

3.3 Sensitivity Through Dual Variables

Consider the dual of the convex optimization problem in (3.3),

$$\min_{a \in \mathbb{R}^N, b \in \mathbb{R}^{d+1}} \sum_{i=1}^N a_i/N + \sum_{j=1}^{d+1} b_j q_N(y_j) + \frac{1}{\theta} \sum_{i=1}^N \sum_{j=1}^{d+1} F_{ij} e^{\theta(C_{ij} - a_i - b_j)}. \tag{3.4}$$

Let (a^*, b^*) denote the optimal dual solution, and consider a vector of small perturbations $(\Delta q_1, \ldots, \Delta q_{d+1})$ to the marginal distribution q_N with components that sum to zero. For perturbations small enough to keep the dual solution unchanged, we can estimate the change in CVA, without resolving problem (3.3), using

$$\Delta\widehat{\mathrm{CVA}}_\theta = \sum_{j=1}^{d+1} b_j \Delta q_j.$$

We can calculate the sensitivity to a parallel shift in the credit curve by setting $\Delta q_j = \Delta$, $j = 1, \ldots, d$, and $\Delta q_{d+1} = -d\Delta$, for sufficiently small Δ.

The dual solution can be obtained as a byproduct of the IPFP algorithm. The optimal primal solution takes the form $P^\theta_{ij} = F_{ij} e^{\theta(C_{ij} - a^*_i - b^*_j)}$, where a^* and b^* are optimal dual variables, so we can define scalars u_i and v_j such that

$$P^\theta_{ij} = \frac{M_{ij} w}{u_i v_j},$$

where $w = F_{ij} \sum_{i=1}^N \sum_{j=1}^{d+1} e^{\theta \cdot C_{ij}}$ is the normalization term in M_{ij}.

Let $r_i(n)$ be the i-th row sum of $\Phi^{(n)}(M)$ and let $c_j(n)$ be the j-th column sum of $\Phi^{(n)}(M)$ after step (a) in the $(n+1)$-th iteration. By Pukelsheim [12],

$$u_i = \lim_{n\to\infty} \prod_{t=0}^{n} \Big(\frac{r_i(t)}{p^N(X_i)} \Big) \quad \text{and} \quad v_j = \lim_{n\to\infty} \prod_{t=0}^{n} \Big(\frac{c_j(t)}{q^N(y_j)} \Big).$$

The optimal dual variables are then given by $a^*_i = \frac{1}{\theta}\ln(u_i) + \frac{1}{2}\ln w$ and $b^*_j = \frac{1}{\theta}\ln(v_j) + \frac{1}{2}\ln w$.

4 Examples

In this section we illustrate the method with examples, starting with the example in Section 2. Assume that the counterparty's default time has an exponential distribution with hazard rate $\lambda = 0.04$. The other model parameters are $(U_{t_0}, \bar{U}, K, \kappa, \sigma, \delta, S) = (1000, 1000, 1000, 0.3, 50, 0.03, 10^6)$. Assume the time horizon T is 10 years. We simulate $N = 1000$ paths, each divided into 20 time steps. After applying the algorithm in Section 3.2 for various values of θ, we get CVA values corresponding to different degrees of wrong-way risk ($\theta > 0$) and right-way risk ($\theta < 0$).

The CVA values are shown in Figure 1, reported as a percentage of the independent case, for which $\theta = 0$. The left panel shows how the CVA bound increases with θ, approaching the worst-case bound as θ becomes sufficiently large. At the other extreme, as θ becomes a large negative number, the CVA approaches zero. These results should be contrasted with the right panel, which shows the result of linking market factors and credit risk through a Gaussian copula, following Rosen and Saunders [13]. Dependence in the Gaussian copula is indexed by the correlation parameter ρ, which ranges from -1 to 1. Restricting the dependence between market factors and credit risk through a Gaussian copula does not achieve the full range of potential wrong-way risk.

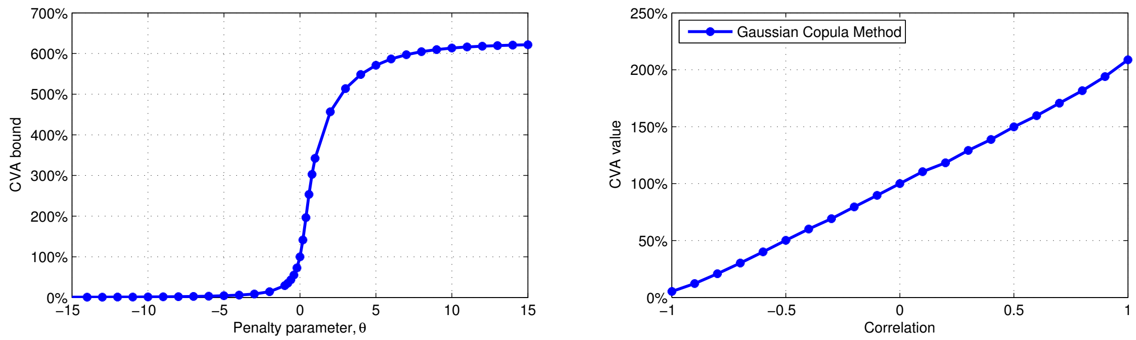

Figure 1: CVA Stress Test. The left panel shows the full range of right-way and wrong-way risk in the CVA valuation for the foreign exchange forward example as we vary the penalty parameter θ. The right panel shows results using a Gaussian copula to link market factors and credit risk. The Gaussian copula does not achieve the full range of potential wrong-way risk.

In Figure 2 we show that varying the volatility of the market factor and the counterparty's default hazard rate have significant impact on CVA. Increasing either of these parameters shifts the curve up for $\theta > 0$. Increasing the volatility of the market exposure or the level of the credit risk in this example increases the potential impact of wrong-way risk, relative to the benchmark of independent market and credit models.

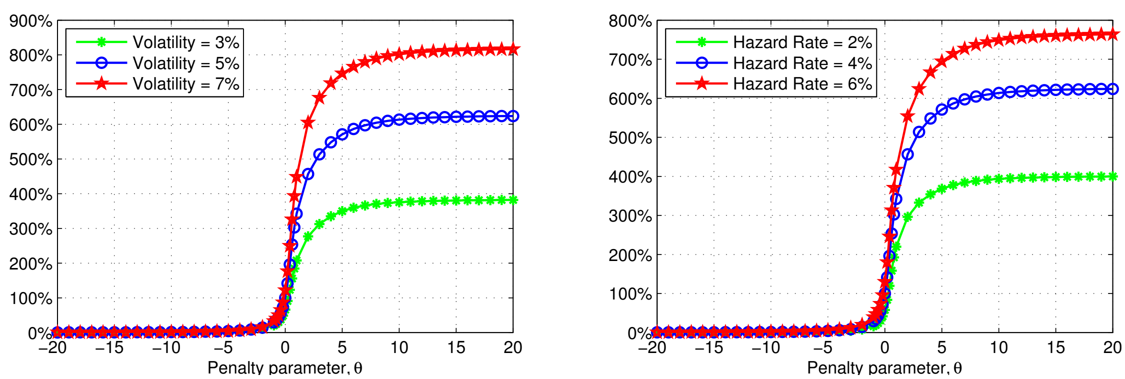

Figure 2: CVA with Different Volatilities and Hazard Rates.

We next consider a 10-year fixed-for-floating cross-currency swap, in which a U.S. bank receives a fixed rate in dollars and pays a floating rate in foreign currency, with a notional size of \$5 million. At the same time, this U.S. bank enters a three-year and a six-year foreign exchange forward contract with the same counterparty in the same currency, each with a notional of \$0.5 million. We use the Vasicek model for the U.S. interest rate,

$$dr_t = \kappa_r(\bar{r} - r_t)dt + \sigma_r d\tilde{W}_t,$$

with parameter values $(r_0, \bar{r}, \sigma_r, \kappa_r) = (0.05, 0.05, 0.0005, 0.8)$.

We consider three different portfolios for the U.S. bank with the same counterparty. The first two portfolios contain multiple transactions of different maturities. The third one contains a single transaction with multiple cash flows.

Portfolio 1 : \$5 million 10-year cross-currency swap, \$0.5 million three-year and six-year foreign exchange forwards. The U.S. bank is the U.S. dollar receiver in all transactions.

Portfolio 2 : \$5 million 10-year cross-currency swap, \$0.5 million three-year and six-year foreign exchange forwards. The U.S. bank is the U.S. dollar receiver in the cross-currency swap and the U.S. dollar payer in the three-year and six-year forward contracts.

Portfolio 3 : A simple interest rate swap with notional size \$5 million. The U.S. bank receives the floating rate and pays a fixed rate $r_{\text{fix}} = 5$ percent.

The sample average positive exposures for these three portfolios are shown in Figure 3. In the top two panels, the average positive exposure increases with time because the largest payments are exchanged at maturity. For Portfolio 1, the drop in exposure at year three and year six results from the expiration of the foreign exchange forward contracts. For Portfolio 2, since the portfolio is more balanced, the exposure path is smoother. In the bottom panel, for Portfolio 3, the average positive exposure decreases to 0 at maturity because the total exposure decreases with time in an interest rate swap.

Figure 4 shows CVA values as θ varies. We report CVA as a percentage of the Portfolio 2 CVA in the independent case. Portfolio 1 has the greatest sensitivity to wrong-way risk because all its transactions run in the wrong-way direction. Portfolio 2 is better diversified, with both wrong-way and right-way transactions. Because the average positive exposure for Portfolio 1 is higher than that of Portfolio 2, it has a higher CVA for all θ, and with θ increasing, Portfolio 1 attains much higher CVA values near the worst-case bound than does the more diversified Portfolio 2. For Portfolio 3, because of its lower and less concentrated positive exposure, the CVA bound is much lower compared with the other two portfolios.

Figure 5 shows the sensitivity of the CVA estimates to a change in the counterparty's default hazard rate. We increase the hazard rate by 1 basis point from $\lambda = 0.04$ to $\lambda' = 0.0401$ and show the estimated change in CVA using dual variables and the actual difference based on resolving the optimization problem at each θ. To put these sensitivities in perspective, the CVA estimate

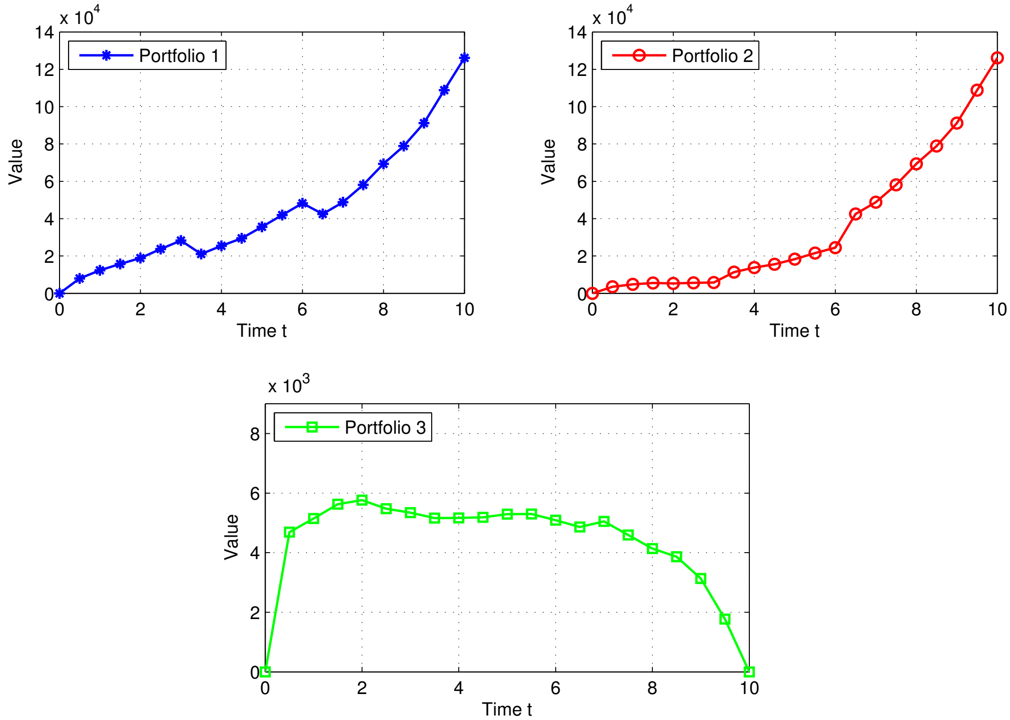

Figure 3: Sample Average Positive Exposures for Three Portfolios.

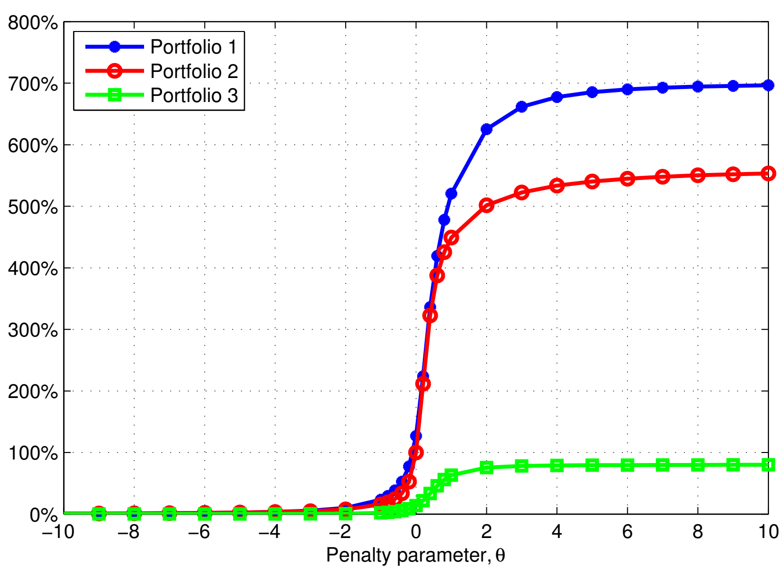

Figure 4: CVA Bounds for Three Portfolios.

at $\theta = 0$ is \$11,241, and at $\theta = 20$, it is \$62,659. The sensitivities in Figure 5 are in dollars. Overall, the dual variables provide good estimates of the change in CVA under a small change in the default probability. Compared with resolving the optimization problem at the perturbed λ, the dual variables slightly underestimate the change in wrong-way scenarios ($\theta > 0$) and overestimate the change in the right-way scenarios ($\theta < 0$).

5 Concluding Remarks

We have presented and applied the method we developed in Glasserman and Yang [8] for bounding the effect of wrong-way risk in measuring counterparty risk through CVA. Our approach provides a practical way to find the largest possible CVA value consistent with separate models for market exposure and credit risk. In this sense, it addresses the model risk that arises from the uncertain dependence between market factors and credit risk. Our penalty formulation tempers the worst-case CVA by allowing the user to control the degree of wrong-way risk.

An important practical problem is the choice of the penalty parameter θ, which reflects the user's confidence in the baseline model. A large θ leads to more conservative values; a small θ produces values very close to the baseline. Inevitably, the choice of θ involves some judgment. However, this judgment is best anchored in observable data. In our examples, each value of θ implies some level of correlation between the exchange rate and the credit spread for the counterparty. This correlation is a limited measure that cannot determine the full dependence between the market and credit models, but it can help pin down an appropriate value for θ.

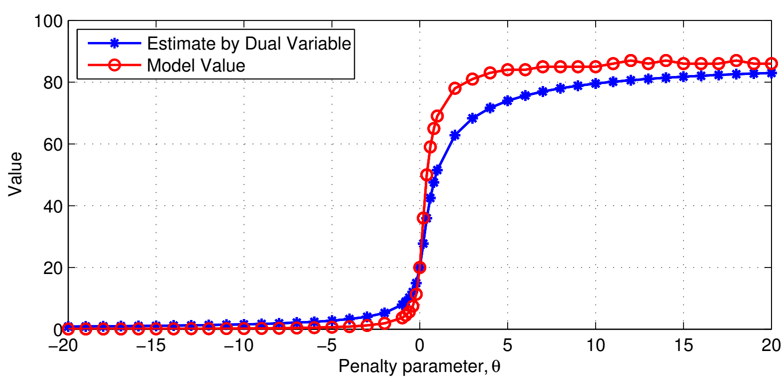

Figure 5: Change in CVA for a 1 Basis Point Change in Hazard Rate at Various Levels of the Penalty Parameter θ.

References

[1] Basel Committee on Banking Supervision (2010) Results of the comprehensive quantitative impact study. Bank for International Settlements, Basel, Switzerland.

[2] Brigo, D., Capponi, A., and Pallavicini, A. (2014). Arbitrage-free bilateral counterparty risk valuation under collateralization and application to credit default swaps. *Mathematical Finance* 24(1), 125–146.

[3] Brigo, D., Morini, M., and Pallavicini, A. (2013). *Counterparty Credit Risk, Collateral and Funding*, Wiley, Chichester, U.K.

[4] Canabarro, E., and Duffie, D. (2003). Measuring and marking counterparty risk, in *Asset/Liability Management for Financial Institutions*, Institutional Investor Books.

[5] Crépey, S. (2012) Bilateral counterparty risk under funding constraints — part II: CVA. *Mathematical Finance.*

[6] Deming, W.E., and Stephan, F.F. (1940) On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *Annals of Mathematical Statistics* 11, 427–444.

[7] Ghamami, S. and Goldberg, L. (2014) Stochastic Intensity Models of Wrong Way Risk: Wrong Way CVA Need Not Exceed Independent CVA. *Journal of Derivatives* 21(3), 24-35.

[8] Glasserman, P., and Yang, L. (2015). Bounding Wrong-Way Risk in CVA Calculation. Available at SSRN: http://ssrn.com/abstract=2607649.

[9] Gregory, J. (2012). *Counterparty Credit Risk and Credit Value Adjustment: A Continuing Challenge for Global Financial Markets*, Wiley, Chichester, U.K.

[10] Hull, J.C., and White, A. (2012) CVA and wrong way risk. Working paper, Rotman School of Management, University of Toronto, Toronto, Canada.

[11] Kleinschmidt, P., and Schannath, H. (1995). A strongly polynomial algorithm for the transportation problem, *Mathematical Programming* 68, 1-13.

[12] Pukelsheim, F. (2014). Biproportional scaling of matrices and the iterative proportional fitting procedure. *Annals of Operations Research* 215(1), 269–283.

[13] Rosen, D., and Saunders, D. (2012). CVA the wrong way. *Journal of Risk Management in Financial Institutions* 5(3), 252–272.