

Class 24: Fixed Income, Credit Risk

This Version: December 1, 2016

1 The Credit Market

Banks play an important role in creating credit. In our class on Risk Management, we talked about a simple model of a bank, which holds only a fraction, say 10%, of its total deposits in its cash reserves and lends out the rest in loans. In doing so (often called fractional reserve banking), credit is created. Another important factor in credit creation is the Federal Reserve. In our class on Monetary Policy and Yield Curve, we talked about how the Fed injects or withdraws the total amount of bank reserves from the entire banking system via open market operations (i.e., purchasing or selling securities). In doing so, the Fed influences the amount of credit creation by banks. When the economy needs help, the Fed inject reserves into the system. With more reserves, banks are more willing to extend credit to their customers. When the economy is at risk of over-heating, the Fed withdraw reserves from the system and the banks are less willing to extend credit.

For anyone in Finance, understanding how the credit market works is essential. At the macro level, it informs you on the overall condition of the economy, given that the economic cycles are often linked to the credit market conditions. At the micro level, many of the financial instruments and much of the financial transactions contain credit components. So knowing how to model and price credit risk is an important skill to have.

Figure 1 plots the debt securities and loans outstanding in the US. By the second quarter of 2015, the total credit in the system is \$62.30 trillion, among which \$14.49 trillion goes to the federal government, \$15.20 trillion goes to the financial sectors, \$12.48 trillion goes to non-financial business and \$14.05 trillion goes to households and non-profit organizations.

Figure 2 tracks the relative importance of the six sectors in Figure 1 by measuring their debt and loan outstanding as a fraction of the total credit. As you can, the fraction of borrowing from the financial sectors increased steadily from 12% in 1980 to a plateau of around 31% in the mid-2000 and decreased quite significantly to 24% in 2015. By contrast, the fraction of borrowing from the Federal Government increased quite dramatically from 11%




Figure 1: Debt Securities and Loans Outstanding.




Figure 2: The Fraction of Debt Securities and Loans Outstanding by Sector.

before the 2008 crisis to 23% in 2015. Interestingly, the borrowing from non-financial business decreased from 30% in 1980 to about 19% before the 2008 crisis, dipped to an all-time low of 18.26% in the first quarter of 2011 and has recovered to 20% in 2015. For households, a large amount of borrowing is in mortgage. Compared with the federal government, non-financial business and financial sectors, the fraction of borrowing from the households remains relatively stable over time. This makes more obvious the significant increase and then the subsequent reduction in household borrowing surrounding the 2008 crisis. The fraction of households borrowing peaked at 27.9% in 2006 to dropped quite steadily to 22% in 2015.

2 Corporate Bonds

- **Default Probability:** In extending credit to a counterparty or an issuer, there are two considerations. First, the credit worthiness of the creditor. For this, we use the concept of *default intensity* to model the likelihood of the creditor's default. Second, in the event of a default, how much can we expect to recover or what is the expected loss? For this, we use the concept of *loss given default*. For the rest of the class, we will focus mostly on the corporate bond market, which by far is the most important component of the credit market.

Figure 3 plots the annual issuer-weighted corporate default rates reported by Moody's in its "Annual Default Study: Corporate Default and Recovery Rates." This is a very useful document updated annually by Moody's and I would encourage you to take a look if you are interested in the credit market. To link the credit market condition to the business cycles, I've also plotted the NBER-dated recession periods in the same plot. The average annual default rate for investment-grade corporates is about 14 basis points. Excluding the great depression era, the average default rate is about 5.84 basis points. In general, the likelihood of default for an investment grade bond (the so-called fallen angle) is low. Since the great depression, the two largest investment-grade defaults were: the default of WorldCom in 2002 with \$33B and the default of Lehman in 2008 with \$120B.

The annual default rates for speculative grades are significantly higher, with an average of 2.83% including the great depression era and 2.77% excluding the great depression. The connection of corporate defaults and business cycle is also stronger for this sector. For the three most recent recessions, the default rates all peaked above 10%, with the 2008 number reaching 13.3%.




Figure 3: Annual Issuer-Weighted Corporate Default Rates for Investment and Speculative Grades.

- Loss Given Default:** In the event of a default, it matters for the bond holder how much the bond is worth. For example, Fannie and Freddie had recovery rates close to 100 percent after receiving government guarantees on their senior unsecured bonds. Interestingly, the terms of the government rescues meant that “credit events” were deemed to have occurred, triggering CDS auctions. For most bonds, however, the recovery rates are much lower than 100%. For the purpose of credit pricing, the loss given default (LGD) is often set to a constant level around 50%.

Figure 4 summarizes the average recovery rates reported by Moody’s. For senior unsecured bonds (which is the category for most corporate bonds), the average recovery rate from 1982-2014 is about 37.4%, making the LGD around $1-37.4\%=62.5\%$. In the most recent years (2013 and 2014), the recovery rates are higher compared with their long-term averages. This in part is due to the positive relation between higher recovery and economic conditions. For example, the recovery rate was 33.8% in 2008. It should be noted that measures such as recovery rates are sensitive to the specific default events. With defaults being rare in general, what we can learn from these default events is difficult to generalize. For example, in 2014, the recovery rate for senior sub-

EXHIBIT 7

Average corporate debt recovery rates measured by post-default trading prices

Lien Position	Issuer-weighted			Volume-weighted		
	2014	2013	1982-2014	2014	2013	1982-2014
1st Lien Bank Loan	78.4%	75.1%	66.6%	80.6%	67.7%	62.5%
2nd Lien Bank Loan*	10.5%	78.7%	31.8%	10.5%	69.2%	28.5%
Sr. Unsecured Bank Loan	n.a.	n.a.	47.1%	n.a.	n.a.	40.2%
Sr. Secured Bond	59.5%	59.8%	52.8%	76.5%	59.5%	52.4%
Sr. Unsecured Bond	43.3%	43.8%	37.4%	34.3%	29.2%	33.6%
Sr. Subordinated Bond*	46.9%	20.7%	31.1%	28.3%	26.6%	26.0%
Subordinated Bond**	38.8%	26.4%	31.4%	38.0%	33.7%	26.3%
Jr. Subordinated Bond	n.a.	n.a.	24.7%	n.a.	n.a.	17.1%

* The average recovery rates for 2014's and 2013's second lien bank loans and senior subordinated bonds were each based on fewer than five defaults.

** The average recovery rates for 2014's subordinated bonds were based on fewer than five defaults.

Figure 4: Average Corporate Default Recovery Rates Measured by Post-Default Trading Prices.

ordinated bonds was 46.9%, slightly higher than the senior unsecured bonds (43.3%). This is due to the fact that the 46.9% number was based on only four defaults.

EXHIBIT 8

Average corporate debt recovery rates measured by ultimate recoveries, 1987-2014

Lien Position	Emergence Year			Default Year		
	2014	2013	1987-2014	2014	2013	1987-2014
Loans*	81.0%	76.7%	80.2%	68.4%	76.6%	80.2%
Senior Secured Bonds**	57.1%	84.2%	63.0%	59.4%	56.9%	63.0%
Senior Unsecured Bonds***	44.6%	61.3%	48.8%	0.0%	34.4%	48.8%
Subordinated Bonds	0.0%	21.0%	28.2%	0.0%	21.0%	28.2%

Figure 5: Average Corporate Default Recovery Rates Measured by Ultimate Recoveries.

The above recovery data are based on trading prices averaged over 30 days after a default event. An alternative recovery measure is based on ultimate recoveries, or the value creditors realize at the resolution of a default event. Going through the bankruptcy court is usually a lengthy process of 1-2 years following the initial default date. As shown in Figure 5, for senior unsecured bonds, the average ultimate recovery rate is 48.8% compared with the average recovery rate of 37.4%. In other words, funds specializing in distressed debt can purchase defaulted bonds and hope to recover the extra 10% by going through the bankruptcy court.

- **Estimating Default Probability:** We will now focus on modeling and estimating the default probability of a bond issuer, while keeping the recovery rate at a constant level.

Information about default probability can be collected from multiple sources. First, we can use credit ratings posted by the three rating agencies (Standard and Poor, Moody's, and Fitch). Using the historical default rates for each rating category, we can get a sense of the likelihood of default knowing the rating of the issuer. For example, as mentioned earlier, the average default rate of an investment-grade issuer is about 14 basis points including the great depression, and 5.84 basis points excluding the great depression. The corresponding numbers for a speculative-grade issuer are higher: 2.83% including the great depression era and 2.77% excluding.

In addition to measuring the default probability using the actual default experiences, there are two other markets from which we can collect information about default. First, a borrower with access to the corporate bond market is in general a large and mature company. In addition to debt financing, it also finance itself through the equity market. As such, information such as the firm fundamentals (e.g., financial statements) and equity market pricing and volatility becomes valuable information for us to assess the credit quality of the bond issuer. For this to work, however, we need to have a model that takes into account of the firm's total value (or cashflow) and prices the equity and bond securities simultaneously. These models are often called structural models of default, pioneered by the work of Merton (1974).

In addition to the equity market, the credit market itself provides a direct measure of default probability. We can use the bond yield spreads in the corporate bond market and the pricing of credit-default swaps to gauge the likelihood of default of an issuer.

The fact that the default probability of an issuer can be collected from these multiple sources implies that these markets are inter-related and any mis-pricing across the multiple markets could be an "arbitrage" opportunity. In fact, understanding the credit market requires a broad knowledge base that includes fixed-income, accounting, equity, and macro-economics. Once, a former MBA student visited my office with his boss and we spent a very nice hour discussing the CDS-bond basis. Toward the end, the boss said, with a full range of emotion, "When I met credits, that's when my love for Finance really started." Not that I wanted to stereotype middle-aged Wall-Street guys, but it was such an usual experience for me. And I fully understand what he meant.

- **Structural Model of Default:** This class of models are pioneered by the work of Merton (1974), and followed up and refined by Black and Cox (1976) and Leland (1994).

In Merton (1974), the total asset value V of a firm is modeled as a geometric Brownian motion with growth rate μ and volatility σ_A :

$$dV_t = \mu V_t dt + \sigma_A V_t dB_t.$$

There are two classes of claimants for V : equity and bond holders. Let K be the book value of the bond, which is a zero-coupon bond that matures in some future date T . Let's consider the fixed time horizon, say t , which is prior to the maturity of the bond. In this model, as long as the time- t firm value V_t is above K , the firm is solvent. And default happens when V_t falls below K .

Using what we learned from the Black-Scholes model, the distance to default can be measured by

$$DD = \frac{\ln(V/K) + (\mu - \sigma_A^2/2) t}{\sigma_A \sqrt{t}}$$

and the probability of default is $N(-DD)$. Although this is a very simple and somewhat unrealistic model, it captures the essence of what drives the probability of default for a name issuer. When the debt-to-asset ratio is high, the firm is closer to the default boundary K . As a result, its distance to default is small and the firm is more likely to default. When the firm's growth potential is good, then the growth rate pulls the firm value away from the default boundary, making it less likely to default. A firm with more volatile asset value is more likely to touch the default boundary and therefore more likely to default.

In addition to probability of default, the Merton model can also price equity and bond simultaneously. The equity of a firm is essentially a call option of the firm's asset value V with strike price K . A low leverage ratio K/V implies the option is deep in the money. For example, the leverage ratio of an Aaa firm is around 13% while the leverage ratio of a Baa firm is around 43%. A typical single A-rated issuer has a leverage ratio around 30%. Even for a single B-rated issuer, the ratio of K/V is around 65%, implying a call option that is deep-in-the-money. In the Merton model, the maturity of the call option is the maturity of the zero-coupon bond. But in practice, firms return to the capital market periodically to manage the maturity structure of their debt. This is where the Merton model (and many of the structural models of default) becomes inadequate.

On the bond side, buying a defaultable bond is the same as holding a default-free bond and selling a deep out-of-the-money put option on the firm's asset value V with strike price

K .

- **Moody's KMV:** The Merton model was used by KMV to calculate expected default frequency (EDF). One of the key innovations of KMV was to recognize that, in the Merton model, the mapping between the distance to default to probability of default relies on the assumption of normal distribution: $N(-DD)$. Instead of using the mapping prescribed by the model, they use the actual default rates for companies in similar ranges to determine a mapping from DD to EDF. Effectively, they are using an empirical mapping. The EDF service provided by KMV has been quite successful and was acquired in 2002 by Moody's in a \$210 million cash transaction.
- **Model Default using Reduced-Form Approach:** Let \tilde{T} be the random default time of a credit issuer. Let t be the horizon over which we care about the survival of this issuer. If $\tilde{T} \geq t$, then issuer is able to survive over the time horizon of our interest and the probability of survival can be summarized by $\text{Prob}(\tilde{T} \geq t)$. Conversely, the probability of default before time t is $1 - \text{Prob}(\tilde{T} \geq t)$.

Let's model this random default time \tilde{T} by exponential distribution:

$$\text{Prob}(\tilde{T} \geq t) = e^{-\lambda t}$$

, where the constant parameter λ captures the default intensity. An issuer with large λ defaults faster. To see this, let's consider the one-year default rate under this model. Setting t to one year, the one-year survival probability is $e^{-\lambda}$ and the one-year default probability is $1 - e^{-\lambda} \approx \lambda$, where I used the linear approximation for e^x for small x . As you can see from this exercise, λ is directly linked to default probability.

We can now use this model to price defaultable bonds. I am going to side step the question about risk-neutral pricing for now. Let's consider a one-year zero-coupon bond with a face value of \$1 and assume that the loss given default is 100%. Let r be the riskfree interest rate (continuously compounded), the present value of the defaultable bond is:

$$P = e^{-r} \text{Prob}(\tilde{T} > 1) = e^{-r} \times e^{-\lambda} = e^{-(r+\lambda)}$$

So the yield to maturity of this defaultable bond is $r + \lambda$ and the credit spread is λ .

If the loss given default is not 100%, then, assuming the loss given default is L , we

have

$$P = e^{-r} \text{Prob}(\tilde{T} > 1) + e^{-r} \times \text{Prob}(\tilde{T} \leq 1) \times (1 - L) = e^{-(r+\lambda)} + e^{-r} (1 - e^{-\lambda}) \times (1 - L),$$

where the second term comes from the recovery rate $(1-L)$. For small λ , the credit spread can be approximated by $\lambda \times L$. Being able to recover part of the bond value makes the credit spread smaller. For two issuers with the same default intensity λ , the one with higher loss rate is priced with a higher credit spread. For example, an issuer might issue two bonds with different seniority, say senior vs. subordinated. Then the difference in the pricing of these two bonds bolts down to the recovery rate.

- **Historical Default Rates and Credit Spreads:** The reduced-form approach gives us a useful tool to connect the historical default rates to credit spreads. In Figure 6, I plot the Moody's one-year default rates together with the credit spreads of Barclays' investment-grade and speculative-grade bond indices. In calculating the credit spreads, I use Barclay's Treasury bond index. For teaching purpose, this is Okay, but the maturity match between the credit indices and the Treasury index are not very well done.

In the first panel of Figure 6, the blue line plots the credit spreads for investment grades and the average is around 150 basis points. The red line plots the one-year default rates and the average over the same sample period is about 9 basis points. Recall that the one-year default probability is $1 - e^{-\lambda} \approx \lambda$ for small λ . Let the red line is essentially λ for investment grade issuer and the average default intensity is 9 basis points. Also recall that the credit spread can be approximated by $\lambda \times L$, for small λ . So the blue line is essentially $\lambda \times L$. If we start with the red line, and multiply it by L , how can we get the blue line?

Of course, in doing this calculation, we assume a constant default intensity and we make no distinction between risk-neutral and actual pricing. In other words, there is no role of default risk premium in our very simple pricing framework. In the academic research, this disconnect between the actual default experiences and the credit spreads has been studied quite extensively. It is called the "credit spread puzzle." The main puzzle is that the credit spreads are too high (or corporate bonds are too cheap) compared with the actual default experiences. Possible explanations include: credit risk premium and liquidity premium.

In the second panel of Figure 6, the same exercise is done for speculative grades. There, the disconnect is not as severe. The average credit spread is 570 basis points while




Figure 6: Corporate Default Rates and Credit Spreads for Investment and Speculative Grades.

the average one-year default rate for the same sample period is 464 basis points. If we start with the red line, and multiply it by $L = 1 - 37.4\% = 62.6\%$ to get the credit spread, we will still have a gap. In other words, even for the speculative grade, the credit spreads are too high. But the gap is not as dramatic as that for the investment grade.

3 Credit Default Swaps

- **Introduction:** The US corporate bond market is among the most illiquid markets. For a market of \$8T in 2015, the average daily trading volume is only \$25B. By comparison, the average daily volume is \$499B for US Treasury and \$321B for US Equity.

In buying a corporate bond, investors take on both duration and credit exposures. To have a pure positive exposure to credit risk, investors have to hedge out the duration risk. To have a pure negative exposure to credit risk, investors have to locate, borrow, and then sell the bonds and buy back the duration exposure. The emergence of credit derivatives was in part a response to the limitations of corporate bonds as a vehicle for credit risk.

Figure 7 plots the size of the CDS market along with the market for interest rate swaps. The market for CDS started out in the late 1990s. It really took off in the mid-2000 and peaked to a notional amount of \$58T in 2007. The notional amount of CDS has declined quite rapidly in recent years and is at \$16T in 2014.

- **CDS:** An investor enters into a CDS contract to either buy or sell credit protection on a named issuer, who could be a corporate issuer (e.g., Ford, GM, etc) or a Sovereign issuer (e.g., Russia, Mexico, etc). In 2009, the CDS market has gone through a pretty large change, which is called the “CDS big bang.”

Let me first describe the contract specification in the “old-fashioned” way. The swap has two legs: the fixed leg consists of quarterly fixed payments indexed to the CDS spread/price and the floating leg pays nothing as long as the named issuer is not in default. In the event of a default (before the maturity of the CDS contract), the payment of the fixed leg stops and the floating leg pays the full face value of the defaulted bond minus the recovery of the bond. In other words, in the event of a default, the seller of the protection makes the bond whole for the buyer.

When the CDS was first introduced, the settlement involves the buyer locating and delivering the physical bond to the seller in exchange for the face value of the bond.

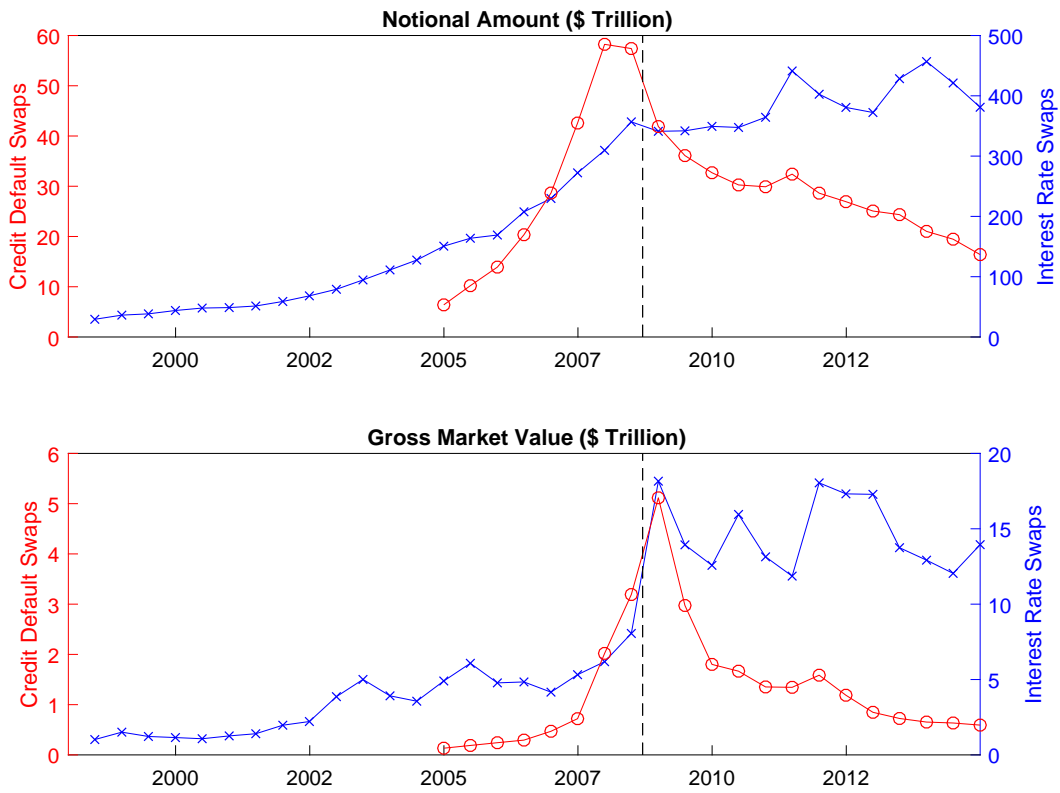


Figure 7: Interest Rate Swaps and Credit Default Swaps, Notional Amount and Gross Market Value.

In some default cases, however, the amount outstanding of CDS contracts became so large that it had the potential to drive up the price of the defaulted bonds as investors scrambled to acquire bonds to deliver. Because of this concern, in some cases, auctions can take place to determine the final recovery rate of a defaulted entity.

As with any swap, the present value of the fixed leg must equal the present value of the floating leg at the start of the contract and this is how the CDS spread/price is determined. Let's think of a very simple example to get some intuition. Suppose there is a one-year CDS on a named issuer. The fixed leg pays only annually. So the present value of the fixed leg (annuity):

$$\text{CDS} \times P(\tilde{T} > 1) \times e^{-r}$$

And the present value of the floating leg (insurance protection):

$$\text{Loss} \times P(\tilde{T} \leq 1) \times e^{-r}$$

We set CDS so that the two legs have the same present value:

$$\text{CDS} = \frac{P(\tilde{T} \leq 1) \times \text{Loss}}{1 - P(\tilde{T} \leq 1)}$$

As you can see, for small one-year probability of default $P(\tilde{T} \leq 1)$, the CDS can be approximated by

$$\text{CDS} \approx P(\tilde{T} \leq 1) \times \text{Loss} = \text{1yr Default Rate} \times \text{Loss}$$

Now let's apply the constant default intensity model to the above calculation:

$$\text{one-year default probability} = 1 - e^{-\lambda}$$

So the one-year CDS price is

$$\text{CDS} = \frac{(1 - e^{-\lambda}) \times \text{Loss}}{e^{-\lambda}} \approx \lambda \times \text{Loss},$$

where the approximation works for small λ .

- **A Few Examples:** By now, it should be clear that the CDS price/spread is simply

the credit spread of a name issuer. For a corporate bond, we can invert its yield from its price. To determine the credit quality, we have to locate a Treasury bond of similar maturity and subtract the Treasury yield from the corporate yield to get the credit spread. For example, during the time period leading up to the GM's default in 2009, the price of the GM bonds were actually going up because interest rates were going down. It is only after taking out the duration exposure and focusing on the credit spread, can we know the true credit quality of GM. By comparison, a CDS on GM gives us direct information about GM's credit quality. Figure 8 plots the time-series of CDS for a few corporate issuers. And Figure 9 plots the time-series of CDS for a few sovereign issuers.

- **The CDS Big Bang:** After the 2008 crisis, the CDS market has gone through some very important change. Let me focus on just one change in contract specification that I think you should know. As I mentioned earlier, a swap involves two counterparties. At the start of the swap, the present value is zero for both counterparties. One important change in the CDS market is such that this is no longer true. As of now, the contracts are still quoted in the “old-fashioned” way, but the actual contract specification is different. For a high-quality (investment grade) named issuer, the fixed leg of the CDS contract is indexed to 100 bps. If the actual credit spread is 150 bps for this named issuer, then the buyer of the protection has to pay upfront fee equaling the present value of the difference between 150 bps and 100 bps. For a low-quality (high yield/speculative grade) named issuer, the fixed leg is index to a fixed rate of 500 bps. Again, an upfront fee (or rebate) is made to adjust for the difference between the deal spread (i.e., 500 bps) and the actual credit spread. To be honest, I am surprised at this change, but it seems to be taking place in the market.
- **CDS-Bond Basis:** For the same issuer, we now have two credit spreads: one from the corporate bond market and the other from the CDS market. The difference between the two is called CDS-bond basis: CDS spread minus the bond yield spread. In theory, these two spreads should be close or the basis should be small.

Figure 10 plots the basis during the 2008 financial crisis. During the depth of the crisis, the CDS-bond basis became very negative, to a level close to -300 bps on average. For some named issuers, the basis were as negative as -500 bps. A negative CDS-bond basis implies that, for the same named issuer, the bond yield spread is larger than the CDS spread. In other words, the cash bond is cheaper than the CDS “bond.” An arbitrage trade on negative CDS-bond basis would involve buying the corporate bond,

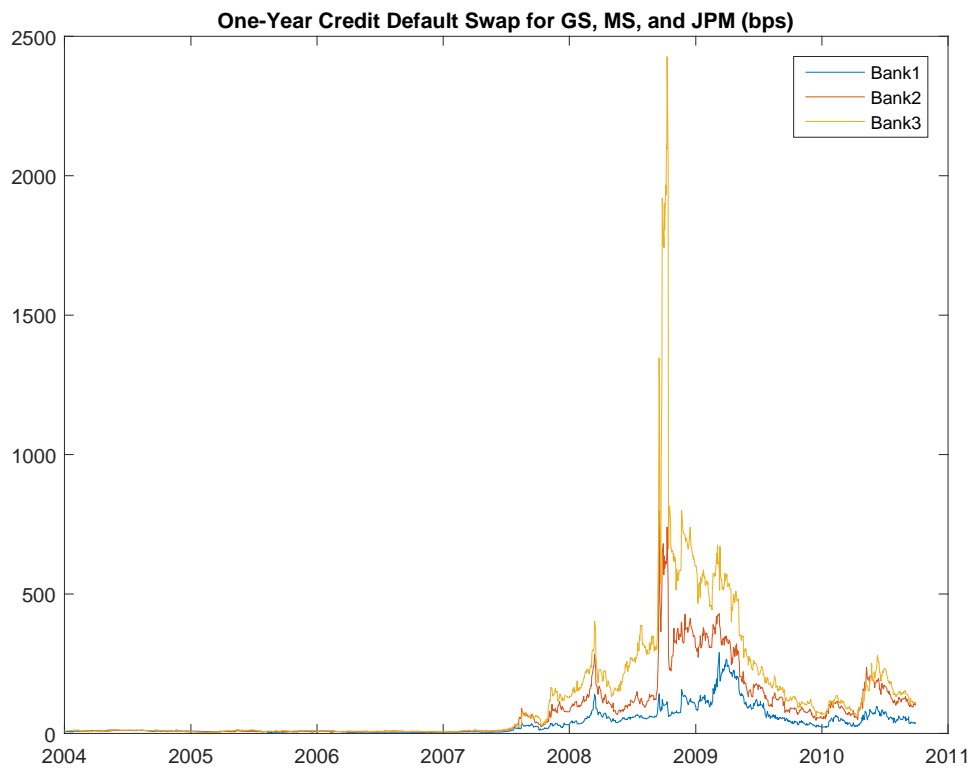




Figure 8: One-Year Credit Default Swaps on a few Corporate Issuers.



Figure 9: One-Year Credit Default Swaps on a few Sovereign Issuers.



Source: J.P. Morgan.

Figure 10: The CDS-Bond Basis in 2008-09.

hedging out the duration exposure and buying protection in the CDS market.

The evolution of the negative CDS-bond basis in 2008-09 was a story of limits to arbitrage. As shown in Figure 10, the basis turned negative after the Lehman default. According to the press, arbitrage trades were put on to bet that the negative basis would converge to zero. But instead of converging, the basis turned more negative. For example, Boaz Weinstein, a trader and co-head of credit trading at Deutsche Bank was down \$1bn, Ken Griffin of Citadel was down 50% and John Thain of Merrill was said to be down by more than \$10bn. The big part of these losses was due to the “negative basis trade.” As they unwound their negative basis trades due to losses, the basis further widened because of their unwinding. Eventually, the basis converged in the second half of 2009 as the financial markets recovered from the crisis.