Forward-Looking Tail Risk Exposures at U.S. Bank Holding Companies

Martin Knaup · Wolf Wagner

Abstract This paper develops a simple method for quantifying banks' exposures
to large (negative) shocks in a forward-looking manner. The method is based on
estimating banks' share prices sensitivities to (market) put options and does not
require the actual observation of tail risk events. We find that estimated (excess)
tail risk exposures for U.S. Bank Holding Companies are negatively correlated
with their share price beta, suggesting that banks which appear safer in normal
periods are actually more crisis prone than their beta would suggest. We also study
the determinants of banks' tail risk exposures and find that their key drivers are
uninsured deposits and non-traditional activities that leave assets on banks' balance
sheets.

M. Knaup
CentER, European Banking Center, and Department of Economics, Tilburg University,
Tilburg, Netherlands
e-mail: m.knaup@uvt.nl

W. Wagner (✉)
CentER, European Banking Center, TILEC, and Department of Economics,
Tilburg University, Tilburg, Netherlands
e-mail: wagner@uvt.nl, wolf.wagner@dsf.nl

W. Wagner
Duisenberg School of Finance, Amsterdam, Netherlands

1 Introduction

A systemic banking crisis—a situation in which many banks are in distress at the same time—can induce large costs for the economy. The task of supervisors and regulators is to avoid and mitigate, as far as possible, such crises. For this they need advance information about how banks are exposed to shocks to the economy. This allows them to identify weak banks and put them under increased scrutiny but also to monitor general risks in the financial system. When evaluating the exposure of banks it is also of paramount importance to distinguish between exposures to normal market shocks, and exposures to large shocks. For example, a financial institution that follows a tail risk strategy (such as writing protection in the CDS market) may appear relatively safe in normal periods as it earns steady returns but may actually be very vulnerable to significant downturns in the economy.

Supervisors and regulators obtain their information to a large extent from information generated by the bank itself, such as its accounts. While these sources are a crucial ingredient of the evaluation process they are not free from drawbacks. For example, most of this information is under the discretion of banks and may be used strategically.[1] Moreover, this data is typically backward looking and available only at relatively low frequency. Accounting information also misses important aspects such as informal knowledge (e.g., CEO reputation) or information contained in analysts' reports.

In recent years there has been growing interest in using market-based measures of bank risk. This is on the back of evidence that market signals contain valuable information about banks' risks (see Flannery (1998, 2001) for surveys). Some of the measures explicitly take into account systemic and tail risk aspects (e.g., Acharya et al. 2009; Adrian and Brunnermeier 2009; De Jonghe 2010). They typically use information from historical tail risk events to compute *realized* tail risk exposures over a certain period.

In this paper we develop a forward-looking measure of bank tail risk. We define a bank's (systemic) tail risk as its exposure to a large negative market shock. We measure this exposure by estimating a bank's share price sensitivity to changes in far out-of-the-money put options on the market, correcting for market movements themselves. As these options only pay out in very adverse scenarios, changes in their prices reflect changes in the perceived likelihood and severity of market crashes. Banks that show a high sensitivity to such put options are hence perceived by the market as being severely affected should such a crash materialize. As this sensitivity reflects perceived exposures to a hypothetical crash, it is truly forward-looking in nature. This property is important to the extent that bank risks change quickly and hence historical tail risk exposures become less informative. Another advantage of this method is that it does not require the actual observation of any crashes, as the method relies on changes in their perceived likelihood.

---

[1] For evidence on such strategic use see, for example, Wall and Koch (2000) and Hasan and Wall (2004) for the reporting of loan losses and Laeven and Majnoni (2003) and Bushman and Williams (2009) for the provisioning of loan losses. Huizinga and Laeven (2009) also provide evidence that banks have used accounting discretion to overstate the value of their distressed assets in the current crisis.

We use our methodology to estimate tail risk exposures of U.S. bank holding companies. We find that the estimated exposures are inversely related to their CAPM beta. Since our methodology estimates tail risk over and above beta risk, this implies that low beta-banks have more tail risk than their beta would suggest. Thus, banks which appear safe in normal times are actually more exposed to a crash. Conversely, of course, high beta banks have lower tail exposure than their normal risk suggests. In other words, banks' risk exposures converge in the tails. This has interesting implications for financial regulation and we discuss various interpretations of this finding in the paper.

We also use our methodology to understand the main drivers of bank tail risk. Understanding these drivers is important for regulators as it gives them information about which activities should be encouraged and which not. There is so far very little research on this question (a notable exception is De Jonghe 2010). Our main findings are that variables which proxy for traditional banking activities (such as lending) are associated with lower perceived tail risk. Several non-traditional activities, on the other hand, are perceived to contribute to tail risk. In particular, we find that securities held for-sale, trading assets and derivatives used for trading purposes are associated with higher tail risk. These findings are consistent with the experience of the crisis of 2008 and 2009. Interestingly, securitization, asset sales and derivatives used for hedging are not associated with an increase in tail risk exposure. This suggests that a transfer of risk itself is not detrimental for tail risk, but that non-traditional activities that leave risk on the balance sheet are. On the liability side we find that leverage itself is not related to tail risk but that large time deposits (which are typically uninsured) are. We also find that perceived tail risk falls with size, which is indicative of bail-out expectations due to too-big-to-fail policies.

The remainder of this paper is structured as follows. In Section 2 we briefly review existing measures of tail risk. Section 3 develops the methodology for measuring tail risk exposure using put option sensitivities. Section 4 contains the estimation of tail risks. Section 5 studies the determinants of tail risk. Section 6 concludes.

2 Existing tail risk measures

The Value-at-Risk (VaR) has for many years been the standard measure used for risk management. VaR is defined as the worst loss over a given holding period within a fixed confidence level.[2] A shortcoming of the VaR is that it disregards any loss beyond the VaR level. The expected shortfall (ES) is an alternative risk measure that addresses this issue. The ES is defined as the expected loss conditional on the losses exceeding the VaR level. Another frequently used measure is Moody's KMV. Essentially, Moody's KMV is a distance to default measure that is turned into an expected default probability with the help of a large historical dataset on defaults.[3]

---

[2]See Standard & Poors (2005) for a general overview and Jorion (2006) for a more comprehensive treatment.

[3](Subordinated) debt and CDS spreads are an alternative and attractive measure of a bank's default risk. A shortcoming of these measures is that these spreads are not available for many banks (in the case of CDS spreads) and often not very liquid (in the case of bonds).

While these measures focus on individual bank risk, there has been a growing interest in recent years in systemic measures of bank risk. One strand of the literature focuses on tail-betas (e.g., De Jonghe 2010). This concept applies extreme value theory to derive predictions about an individual bank's value in the event of a very large (negative) systematic shock. Loosely speaking, this method uses information from days where stock market prices have fallen heavily and considers the covariation with a bank's share price on the same day. It thus focuses on realized covariances conditional on large share price drops. A difficulty encountered when applying this method is that tail risk observations are rarely observed, and hence a large number of observations are needed to get accurate estimates (De Jonghe (2010) suggests at least six years of daily data).

Acharya et al. (2009) develop a measure similar to the concept of market dependence, which is based on expected shortfalls instead of betas. They propose measuring the Marginal Expected Shortfall (MES), which is defined as the average loss by an institution when the market reaches a certain quantile of its left tail. Huang et al. (2010) propose a related measure focusing instead on a threshold loss for a portfolio of large banks as the tail risk event. Adrian and Brunnermeier (2009) consider a different aspect of systemic risk. They estimate the contribution of each institution to the overall system risk. A bank's CoVaR is defined as the VaR of the whole financial sector conditional on the bank being at its own VaR level. The bank's marginal contribution to the overall systemic risk is then measured as the difference between the bank's CoVaR and the unconditional financial system VaR. An advantage of the CoVaR is that it is relatively simple to estimate, as it is based on quantile regressions. In terms of its informational properties it is similar to the tail risk beta in that it focuses on realized tail risk.

Our measure is most similar to the tail risk betas as we also measure bank exposures to large market swings. A difference that is important for the interpretation of the estimates, however, is that while the tail risk beta relates to large daily market drops, we estimate exposures to a large prolonged downturn in the market (e.g., several months).

There is literature on hedge funds performance which uses a methodology similar to ours. This literature estimates tail exposures for various styles of hedge funds with a non-linear market factor that takes the shape of an out-of-the-money put option (see, for example, Agarwal and Naik 2004). However, the focus of this literature is different. While we are interested in estimated tail risk exposures per se (i.e., regulators want to know which banks are more exposed to tail risk), the hedge fund literature looks at whether tail risk exposures can be used to forecast fund performance.

3 Measuring tail risk using put option sensitivities

In this section we present our methodology for measuring banks' tail risk exposures. We define the latter to be the bank's exposure to a general market crash (that is, a severe downturn in the economy). If the market crashes, a bank may suffer large, simultaneous losses on its assets, which may push it close to or into bankruptcy. Crucially, the extent to which it is exposed to crashes may differ from its normal market sensitivity. Consider two banks, A and B. Bank A invests mostly in traditional

banking assets such as, for example, loans to businesses and households. Moreover, it invests in assets that are mainly exposed to normal period risk, such as, for example, junior tranches of securitization products (which lose value for modest increases in defaults, but are insensitive to defaults that go beyond the first loss level). In addition to these assets, bank A insures itself against default by buying protection on its assets (such as by buying credit default swaps on its loans). Bank A's equity value will thus depend less on the market in times of crisis, compared to normal times.

Bank B follows a different business strategy. It invests in traditional assets as well. However, in addition, it also follows investment strategies that return a small and steady payoff in normal periods but incur large losses when the market crashes. Examples of such a strategy is selling protection in the credit default swap (CDS) market or buying senior tranches of securitization products (which lose value only when all other tranches have already incurred a total loss). Thus, even though bank B's equity value may behave similarly to bank A's in normal periods, it tends to fall relatively more when the market crashes.

We next describe our method. For this we consider the economy's representative firm (the "market"). We suppose the firm exists for one period only and that its (stochastic) next period equity value is denoted with x. Similarly, we consider a bank with next period equity value y. We assume for the relationship between the equity values of the bank and the market:

$$y(x) = \begin{cases} x^{\beta} & \text{if } x \geq \overline{x} \\[2ex] \dfrac{x^{\beta}}{\left(\frac{\overline{x}-x}{\overline{x}} + 1\right)^{\gamma}} & \text{if } x < \overline{x} \end{cases} \tag{1}$$

When $x \geq \overline{x}$, the bank's equity value has thus a market dependence equal to the one of a firm with a beta of β (that is, the bank's return is β times the market return). However, for $x < \overline{x}$, the bank's equity value additionally depends on the relative shortfall of the market to $\overline{x}$, $\frac{\overline{x}-x}{\overline{x}}$ ($\in [0, 1]$). For $\gamma > 0$ its equity value will be more sensitive to the market, hence the bank has tail risk over and above the normal period exposure (as expected for Bank B), while for $\gamma < 0$ we have the opposite case (Bank A). Only in the case of $\gamma = 0$ does the bank's tail not differ from its normal period risk.

Since tail risk realizations ($x < \overline{x}$) are rarely observed, our estimation relies on changes in perceived tail risk, which we will measure through changes in put options prices. For this consider a put option with strike price $\overline{x}$ that is deep out-of-the-money ($\overline{x}$ is hence a tail risk realization). We have for the pay-off from this put

$$p(x) = \begin{cases} 0 & \text{if } x \geq \overline{x} \\ \overline{x} - x & \text{if } x < \overline{x} \end{cases} \tag{2}$$

Inserting into Eq. 1, totally differentiating with respect to y and dividing by y yields

$$\frac{dy(x)}{y} = \beta \frac{dx}{x} - \gamma \frac{dp}{p + \overline{x}}. \tag{3}$$

Percentage changes in the bank's equity values ($\frac{dy(x)}{y}$) thus relate to percentage changes in the market ($\frac{dx}{x}$), giving the standard β-effect. Additionally, they also

relate to relative changes in the value of the option, $\frac{dp}{p+\overline{x}}$,[4] arising from tail risk exposure.

In our empirical implementation we will identify tail risk sensitivities (γ in Eq. 3) by adding a put option (on the market) to a standard market regression and interpreting the sign of the put option coefficient. Tail risk sensitivities will thus be estimated through changes in put option prices, which (loosely speaking) arise from changes in either the likelihood of a market crash or its severity.[5]

3.1 A discussion of the methodology

We believe that this method has several attractive features. First, the method is forward-looking in nature, that is, it captures *expected* tail risk exposure at banks. This contrasts with other popular methods for measuring tail risk, such as tail risk betas or the CoVAR. These methods essentially compute correlations (or covariation) of banks with the market (or other banks) at days of large share price drops. They thus draw inferences from historical tail risk distributions and hence measure *realized* tail risk. The difference between forward and backward-looking measures is likely to be limited when banks only undergo small changes in their risks over time, but is potentially important in a dynamically evolving financial system.

Second, our measure identifies banks' tail risk exposure through *changes* in expected market tail risk, as measured by put option prices. This has the advantage that for our estimation we do not need tail risk events to materialize. Such events, by definition, occur only very infrequently and hence it is difficult to estimate their properties. Existing measures that rely on the historical distribution of tail risk events reduce this problem by relying on a large time series and by looking at modest tail risk realizations that occur more frequently. Our method allows the measurement of exposure to extreme forms of tail risk (for this one simply includes a very far out-of-the-money put option).

Since we estimate exposures to market crashes, our measure captures *system* tail risk exposure. This is desirable since externalities from banking failures are typically associated with system events, and not isolated bank failures. It should, however, be kept in mind that a bank that has a low estimated systematic tail risk may still be individually very risky to the extent that it pursues activities that are uncorrelated with the market. In addition, one should also keep in mind that market risk is not identical to banking sector risk. Even though banks' market exposures have probably increased in recent decades, credit risk is still the major source of risk for banks. However, market and credit exposures are highly correlated in practice: when economic conditions deteriorate, the default risk of firms increases and stock values decline at the same time. For example, during our sample period the correlation

---

[4]The correct term here is indeed $\frac{dp}{p+\overline{x}}$ and not, as one might think, $\frac{dp}{p}$. The bank-market relationship consistent with $\frac{dp}{p}$ would be $y = \frac{x^\beta}{(\overline{x}-x)^\gamma}$ for $x < \overline{x}$ as one can easily verify, which is not a sensible one as for $x = \overline{x}$ the denominator would then be infinite.

[5]The estimation of γ is akin to estimating the factor-loadings in the asset pricing literature (see, for instance, Ang et al. (2006) and the references therein). While in the asset pricing literature the factor loadings are often used to predict returns in a second step, we are interested here in the cross-sectional distribution of the factor-loadings. More precisely, we propose using the cross-sectional variation to identify banks that are perceived as being prone to a market crash.

between the S&P 500 and the CDX crossover index was -0.77. Due to this high correlation, our estimates will (indirectly) also capture credit risks at banks.[6]

In our empirical implementation of Eq. 3 we measure tail risk exposures by the (negative) coefficient of a put-option return (γ) in a regression on bank stock returns. This, however, is in a regression where we also separately control for the market return. Conditional on the market, a key driver of put-option returns is market volatility. We can thus expect the γ to give us similar information as the (negative) coefficient in a regression of bank returns on market volatility.[7] This provides us with an alternative interpretation of the γ. If a bank is symmetrically exposed to upwards and downwards movements in the market, an increase in volatility will not affect its value and the γ should be zero. However, if a bank is more exposed to downward than to upward movements (e.g., Bank B), its value will decrease when volatility increases. It then obtains a positive γ. The tail-risk estimate can thus also be interpreted as a measure of how much more a bank is exposed to downturns than to upturns. Total tail risk is then a combination of the symmetric dependence on the market (given by the standard β-risk) and the asymmetric sensitivity to downturns (the γ-risk).

It should be noted that our measure, like other market-based measures, is net of any bailout expectations. If, for example, markets anticipate that governments may bail out certain banks (for example because they are too-big-to-fail) then these banks may have a low perceived tail risk even if their underlying activities are relatively risky (Kane (1985), for example, shows that the expected value of these bail-out subsidies can be significant). Thus, while our estimates are important for regulators and supervisors in that they quantify a bank's effective failure risk, there are less suitable for being used as a base for regulation that aims at reducing risk-shifting (for example, by conditioning capital requirements on tail risk exposures).

4 Empirical analysis

4.1 Data

We collect daily data on bank share prices and the S&P 500 (our proxy for the market) for the period 4 October 2005 until 26 September 2008 from Datastream. Put option data on the S&P 500 for the same period is from IVolatility.[8] In addition,

---

[6]An alternative to using put-options on the market are senior tranches of securitization products. These tranches only lose value in extreme circumstances and hence represent tail *credit* risk. However, the pricing of such tranches in financial markets is rather imperfect at present; hence they are not suitable for estimating tail exposures.

[7]The two coefficients will obviously not provide identical information since put-option prices (conditional on market returns) can also change due to other factors, such as interest rates, dividends and (most importantly in our context) the skewness of the distribution. Overall, it is preferable to use put-options (instead of volatility) as regressor as this will also capture variations in tail risk arising from changes in skewness.

[8]We also considered using put options on a banking index (the BKX index) instead of the market. There are two disadvantages to this. First, the banking sector index by itself will already reflect tail risk in the financial system, thus the interpretation of the γ-estimates is not straightforward. Second, put option prices on the index are fairly illiquid.

various balance sheet data are collected from the FR Y-9C Consolidated Financial Statements for Bank Holding Companies (BHCs). We focus on U.S. BHCs which are classified as commercial banks and for which data is fully available. We focus on the BHC instead of the commercial bank itself, as typically it is the BHC that is listed on the stock exchange. Excluded are those banks whose share price change is zero in more than 10% of the cases in order to mitigate problems arising from illiquidity. Foreign banks (even when listed in the U.S.) and pure investment banks are also excluded. The final sample contains 209 Bank Holding Companies.

An important question is the choice of the option strike price. Ideally we would choose an option such that on each day it represents the same crash probability. Simply taking an option with a constant strike price is hence not appropriate as market prices change over time and hence the moneyness of the option will change. Taking the strike price to be a (fixed) fraction of the S&P500 is also not desirable as this ignores that the likelihood of tail risk realizations is also driven by the volatility. We hence decided to construct a series of options such that their option price does not vary. Specifically, each day we adjust the strike of the option such that the previous day price of the options is fixed over time.

For this we use an option price of 0.5$, which translates into an implied strike that was on average 33% below the S&P 500 during our sample period.[9] We have checked the S&P500 over the last 25 years and have found three periods with stock market declines of this magnitude: the 1987 stock market crash (maximum decline: 32%), the burst of the dotcom bubble (maximum decline: 23%) and the subprime crisis (maximum decline: 41%). Thus, such a decline materialized about once every eight years.

In order to compute the option price change for, say day 1, we proceed as follows. We first identify among all traded options the two strike prices that give day 0 prices closest to 0.5. We then calculate the weight that makes their average price 0.5. Given this weight, we calculate the weighted average of their prices at day 1 and calculate from this the change of the price, dP, from day 0 to day 1. Effectively, we compute price changes of options whose (hypothetical) strike price varies from day to day.

We initially considered all out-of-the-money puts. A first inspection, however, revealed that the 100er strikes (i.e. 500, 600, 700 etc.) are much more liquid than put options with other strike prices. We therefore use only these puts. For each day an option's strike price and its price change are then calculated according to the procedure described above. In order to mitigate the influence of changes in the remaining time to maturity on our analysis, we use for this an "on-the-run" series, where each quarter we jump to more recently issued options with longer maturity. As a result, the remaining time to maturity is limited to an interval of between three and six months.

4.2 Estimated tail risk exposures

We estimate Eq. 3 for each bank using the following specification:

$$\frac{\Delta y_t}{y_t} = \alpha_0 + \alpha_1 \frac{\Delta x_t}{x_t} + \alpha_2 \frac{\Delta p_t}{p_t + \overline{x}_t} + \varepsilon_t. \tag{4}$$

[9]In the more tranquil (low volatility) times of 2006, the average implied strike was around 28% below the S&P 500 while after June 2007 it was on average around 38% below the S&P 500.

In Eq. 4 the subindix t indicates time. The market value of the bank and the S&P500 index are denoted by y_t and x_t, respectively ($\frac{\Delta y_t}{y_t}$ and $\frac{\Delta x_t}{x_t}$ are hence the returns on the bank and the market). The term p_t denotes the price of a put-option on the S&P 500 index with a strike price $\overline{x}_t$ that is set such that its previous-day-price is a constant over time (as described in the previous section), $\overline{x}_t$ denotes the (time-varying) strike price of the option, and ε_t is an error term that fulfills the classical OLS-assumptions.

We expect α_1 in Eq. 4 to be close to one in case banks display similar properties as other firms in the market. We do not hold any priors about the sign of α_2 (note that γ in the model (3) relates to $-\alpha_2$ in the estimation; high α_2 thus indicates *less* tail risk). If a bank is similar to the average firm in the market, its γ should be zero. This is because the bank will then react one-to-one to market movements. Its market dependence in the tail is then not different from its market dependence in normal times and hence gamma is zero. The existence of systemic risk in the financial system may suggests that banks display excess market dependence in the tail, that is, γ is positive (negative α_2). However, bail-out expectation may also limit the perceived exposures of the stocks of banks to a market downturn, in which case we obtain a negative γ (positive α_2).

We estimate Eq. 4 using OLS on daily data. For this we winsorize the dependent variables at the 2.5% level. Table 1 reports summary statistics for the coefficients on the 209 bank-level regressions. We can see that the betas are reasonably distributed. The mean beta is 1.56; banks are thus on average riskier than the market. The 25th and the 75th percentiles are 1.39 and 1.93, respectively. The γ ($= -\alpha_2$) is only significant in 16.3% of the cases, which is not surprising considering that the γ of an average firm in the market should be zero. However, we can see that there is substantial cross-sectional variation in the γ: the 25th and 75th percentiles are -12.2 and -2.8. The mean γ is negative (-7.8), suggesting that that overall factors that reduce tail risk (relative to the market) dominate.

What can be said about the economic significance of the γ-estimates? For the β-estimates it is straightforward to interpret their value since a drop in the market by $x\%$ translates into a drop in the stock price by $\beta x\%$. Such a simple relationship does not exist for the γ because the return on a put-option is not proportional to the return on the market. However, in order to get a sense of the economic significance of the γ-estimates one can do the following exercise. We can consider different scenarios for (instantaneous) drops in the market index of, say 5, 10, 15, 20%. For these drops, we can calculate implied put-option price changes (using an option price formula). We can then use Eq. 3 to calculate the share price return for the average gamma implied by the put-option change, which is given by the term $-\gamma \frac{dp}{p+\overline{x}}$ ($= \alpha_2 \frac{dp}{p+\overline{x}}$). Finally, we can compare this return to the share price return implied by the average beta.

Table 1 Summary statistics for 2-factor model. The table summarizes the coefficients from the 209 bank-level 2-factor regressions given by Eq. 4

	Mean	Stdev	25th perc.	50th perc.	75th perc.	Percentage of banks for which coeff. is significant
Intercept (α_0)	−0.00010	0.00072	−0.00049	0.00007	0.00035	1.4%
Beta (α_1)	1.562	0.562	1.387	1.691	1.928	95.7%
Gamma ($-\alpha_2$)	−7.832	8.375	−12.200	−7.721	−2.811	16.3%

Significance level in the last column is 5%

Some caveats apply to this method. First, when we compute the implied option price change we keep constant all the other determinants of the option price, while in reality a significant drop in the index value, for example, is likely to be associated with a change in volatility as well (most likely an increase in the volatility). Thus, our calculations may over- or underestimate the real impact on the put-option price. Second, the OLS coefficients for beta and gamma relate to small changes in the explanatory variables, while we simulate large changes in these variables. For the beta-coefficient this issue may be relatively innocent—as it only requires the relationship between the bank and stock return to be linear over a wide range of index returns. However, for the gamma this assumption is more problematic. This is because the price of an out-of-the-money put-option responds non-linearly to changes in the market. In particular, it becomes more sensitive to the market as the market gets closer to the strike price (less out-of-the money). Our OLS estimates of the gamma are based on the less sensitive range (where the index value is far away from the strike) but for the simulations we make inferences about the more sensitive range (where the put-option is less out-of-the-money). This may introduce an additional source of error in our exercise.

We proceed as follows. In order to calculate the expression $-\gamma \frac{dp}{p+\bar{x}}$, we assume an index value equal to the average of the S&P 500 during our sample period. From this we calculate the strike of the option using the average discount used in our analysis. We then calculate the implied volatility (using the Black–Scholes formula) which makes the option price equal to 0.5 (the price used in our regressions). Holding this volatility constant we can then compute the option price change if the market drops by a certain amount. Using the sample mean γ, we can then calculate the term $-\gamma \frac{dp}{p+\bar{x}}$, which gives us the share price return induced by the γ-risk. The share price return induced by the β-risk of the average bank is simply given by the (mean) β times the drop in the market.

Table 2 summarizes the results for the various scenarios about index drops. The first column shows the stock return implied by the β-risk, the second column the return implied by the γ-risk. The third column shows the total implied return. We can see that for modest index drops, the γ-exposure does not matter a lot. For example, for an index drop of 5%, the average share price change implied by the gamma is only 0.6%, compared to an average share price change implied by the beta of −7.8% (recall that the average beta in our sample is larger than one). However, for larger

Table 2 Economic significance of γ-estimates. The table calculates the share price reaction to various market crash scenarios. The first column gives the share price reaction implied by the average beta (= 1.56). The second column gives the share price reaction implied by the average gamma (= −7.82). The third column gives the combined share price reaction. The last column gives the difference in the share price response of a bank at the 25% quantile and at the 75% quantile of the gamma-distribution

Scenario index return	Stock price return implied for average β	Stock price return implied for average γ	Total stock price return for average β and γ	Difference in stock price return between 25% and 75% γ-quantile bank
−0.05	−0.078	0.006	−0.072	0.007
−0.1	−0.156	0.020	−0.136	0.023
−0.15	−0.234	0.048	−0.186	0.056
−0.2	−0.312	0.102	−0.210	0.120

index drops the importance of the gamma raises. For example, for an index drop of 20%, the change implied by the gamma is 10.2%, while the beta-implied change is −31.2%. The increasing importance of the γ-exposure for larger index drops reflects the non-linear dependence of the put-option price on the underlying: as we get closer to the strike, the sensitivity (delta) of the (out-of-the money) put increases.

Besides this exercise (which compares the average γ-risk with the average β-risk), it is also informative to study how important the cross-sectional variations in γ-risk are in economic terms. This matters for a regulator who wants to know whether a bank that has a high γ-risk relative to its peers really has much more tail risk. For this we have calculated in the last column of the table the difference in the implied share price return for a bank that has a gamma equal to the 25% quantile of the distribution with the one of a bank at the 75% quantile of the distribution. As before, it turns out that for smaller index drops the gamma does not matter a lot (for a 5% drop the difference in the return is about 0.7% between the two banks). However, for larger drops, the difference becomes important. For example, when the market drops by 20%, a bank with a gamma at the 75% quantile of the γ-distribution drops by 12% more than a bank at the 25% quantile.

Figure 1 shows next the gammas plotted against bank size. It can be seen that there is considerable variation. There also seems to be a pattern of large banks having lower tail risk.

An important question is whether our tail risk measure really adds anything in terms of informational content to the normal beta. For example, it may simply be that the banks with large tail risk are also banks that have a large beta. In this case, estimating the tail risk beta separately is of little value. To shed light on this question we study how gammas relate to betas. In order to avoid potential interdependencies between beta and gamma arising when they are estimated in the same regression, we estimate for this betas that are obtained from a standard one-factor model (that is, without put-options) of the following form:

$$\frac{\Delta y_t}{y_t} = \alpha_0 + \alpha_1 \frac{\Delta x_t}{x_t} + \varepsilon_t. \tag{5}$$

Fig. 1 Gamma and bank size
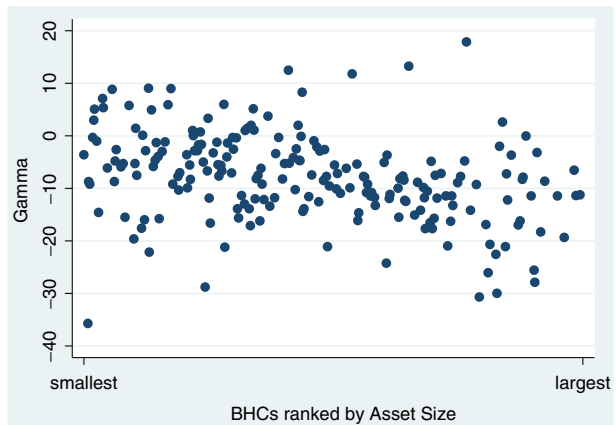
Table 3 Summary statistics for 1-factor model. The table summarizes the coefficients from the 209 bank-level 1-factor regressions given by Eq. 5

|	Mean	Stdev	25th perc.	50th perc.	75th perc.	Percentage of banks for which coeff. is significant
Intercept (α_0)	−0.00016	0.00072	−0.00055	0.00002	0.00031	1%
Beta (α_1)	1.476	0.529	1.329	1.613	1.832	97%

Significance level in last column is 5%

Table 3 provides summary statistics for β's estimated from this 1-factor model, showing that although the mean beta is now somewhat smaller (1.48 instead of 1.56), the overall distribution is similar to the one from the 2-factor model.

Figure 2 plots the banks' gammas against market betas. The scatter plot shows that higher gammas cannot simply explained by higher betas. In fact, there is a strong negative relationship between beta and gamma.

This negative relationship has the following interpretation. Since the gamma is estimated from a regression that also includes the market, it measures tail risks *over and above* the tail-dependence implied by the beta. In other words: a positive gamma implies that a bank has more tail risk than its beta would suggest. A negative relationship between beta and gamma thus means that low beta banks have more tail risk than their beta would indicate (and vice versa for high beta banks). Thus, there is convergence in the tail: banks' risk exposures in the tail are more similar than the ones in normal times.

A potential explanation for the negative relationship is so-called tail risk strategies, which produce steady returns in normal periods but actually expose the banks to severe downturns. For example, an institution that writes protection in the CDS market receives in normal periods a relatively safe stream of insurance premia. However, in a large recession many exposures will simultaneously default and large losses may materialize for this institution. Many trading strategies, such as the ones exploiting apparent arbitrage relations, create similar pay-off distributions. If banks that have low beta (and hence also low tail risk) tend to source additional tail-risk through such means, this could explain our finding. Another explanation for

Fig. 2 Gamma vs. beta
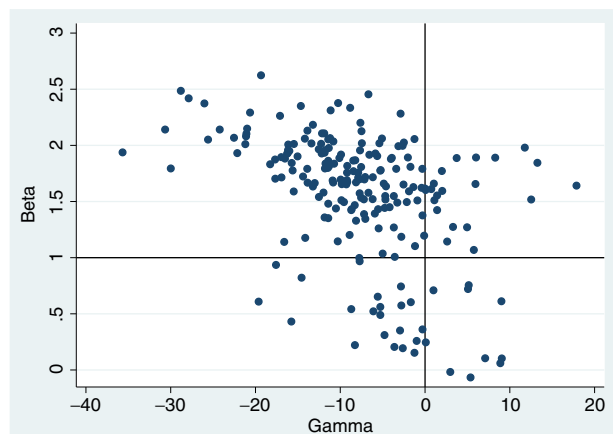
this negative correlation is that highly profitable institutions that operate in risky environments protect their franchise, for example by buying protection in the CDS market or by imposing a less fragile capital structure. Yet another interpretation is that it is simply difficult for banks to avoid exposure to systemic events. Thus while banks may differ substantially in their normal business risk, their tail exposure may be rather similar.

We can classify banks into different risk profiles using the two dimensions of Fig. 2. Specifically, we allocate banks to four groups relative to a benchmark bank with $\beta = 1$ and $\gamma = 0$. This creates the following categorization in Fig. 2:

1. Upper-right quadrant (Quadrant I): Banks with high normal times ($\beta > 1$) and high tail risk ($\gamma > 0$).
2. Upper-left quadrant (Quadrant II): Banks with high normal times ($\beta > 1$) and low tail risk ($\gamma < 0$).
3. Lower-left quadrant (Quadrant III): Banks with low normal times ($\beta < 1$) and low tail risk ($\gamma < 0$).
4. Lower-right quadrant (Quadrant IV): Banks with low normal times ($\beta < 1$) and high tail risk ($\gamma > 0$).

Table 4 lists the top-10 banks (in terms of their average asset size during the sample period) for each of the four risk profiles. We can see that all the very large

Table 4 Classification of banks according to their risk profiles. The table lists the ten largest banks for each quadrant in Fig. 2

Quadrant I ($\beta > 1$ and $\gamma > 0$)		Quadrant II ($\beta > 1$ and $\gamma < 0$)	
Name of BHC	Log(Assets)	Name of BHC	Log(Assets)
UNIONBANCAL CORP	17.80888	CITIGROUP INC	21.38418
W HOLDING CO INC	16.65328	BANK OF AMERICA CORP	21.16085
SANTANDER BANCORP	16.00912	JPMORGAN CHASE & CO	21.11819
AMCORE FINANCIAL INC	15.47797	WACHOVIA CORP	20.33326
ITLA CAPITAL CORP	15.06647	WELLS FARGO & CO	20.09249
FIRST FINL BANCORP INC	15.03965	U S BANCORP	19.23489
INTEGRA BANK CORP	14.91511	SUNTRUST BANKS INC	19.00746
REPUBLIC BANCORP INC	14.90969	NATIONAL CITY CORP	18.79059
CITY HOLDING COMPANY	14.73645	STATE STREET CORP	18.67159
WASHINGTON BANC. INC	14.72756	CAPITAL ONE FINC. CORP	18.65735
Quadrant III ($\beta < 1$ and $\gamma < 0$)		Quadrant IV ($\beta < 1$ and $\gamma > 0$)	
Name of BHC	Log(Assets)	Name of BHC	Log(Assets)
FIRST CITIZENS BANCSH	16.58151	EUROBANCSHARES INC	14.75665
PENNSYLVANIA COMM BANC.	14.45603	FARMERS CAPITAL BANK	14.44821
CADENCE FINANCIAL CORP	14.41014	MBT FINANCIAL CORP	14.26376
FNB UNITED CORP	14.37156	FIRST UNITED CORP	14.16776
STATE BANCORP	14.30592	TIB FINANCIAL CORP	14.11398
SOUTHERN CMNTY FINL CORP	14.22843	FIRSTBANK CORP	13.99682
BEVERLY HILLS BANCORP INC	14.20897	PAB BANKSHARES INC	13.95141
BANCORP RHODE ISLAND INC	14.2045	SMITHTOWN BANCORP INC	13.93047
GATEWAY FINANCIAL HLDGS	14.19157	GERM. AMER. BANCORP	13.90569
CAPITAL BANK CORP	14.16501	CAMCO FINANCIAL CORP	13.86312

banks are in Quadrant II (high beta − low gamma). This is good news to the extent that it suggests that many of the large banks have tail-risks that do not exceed their normal times risk. In addition, due to their high-normal times risk these banks are likely to be already on the radar screen of regulators. We can also see that some of large banks are in Quadrant I (high beta − high gamma). These banks have excess tail-risk, but are probably under (well-deserved) regulatory scrutiny because they also have large betas. The probably most interesting quadrant—from a regulatory perspective—is Quadrant IV (low beta − high gamma). These are the banks were normal times risk understates their tail risk and which may hence fall through the cracks. Regulators should pay more attention to these banks. It is, however, comforting that the largest banks do not belong to these category. In fact, the banks in this category tend be much smaller than in the previous two quadrants. Finally, in Quadrant III (low beta − low gamma) we have the relatively unproblematic banks. These banks tend to be also small in size.

5 Determinants of bank tail risk

In this section we study whether and how a bank's business activities relate to its tail risk. The most obvious way to do this is by regressing (estimated) gammas upon a number of balance sheet variables that represent various banking activities. This two step method, however, has the disadvantage that the estimation is not efficient as information from the first step (estimating the gammas) is not used in the second step.

For this reason we estimate the relationship in one step.[10] We amend Eq. 3 to allow a bank's put option sensitivity to vary with a certain bank activity, say B, relative to its sample mean ($\hat{B}$). In addition, we also interact the S&P 500 return with the balance sheet variable B to take into account that general market sensitivities may also differ depending on bank activities. We obtain the modified equation:

$$\frac{dy(x)}{y} = \alpha + (\beta + \theta(B - \hat{B}))\frac{dx}{x} - (\gamma + \delta(B - \hat{B}))\frac{dp}{p + \overline{x}}. \qquad (6)$$

The coefficient δ in this equation gives us the relationship between a bank's gamma and activity B (the equivalent of the coefficient of a regression of estimated gammas on B), evaluated at the mean. Since we are interested in several determinants of bank tail risk, we estimate a multivariate variant of Eq. 6:

$$\frac{\Delta y_{i,t}}{y_{i,t}} = \alpha_0 + \left(\alpha_1 + \sum_j \alpha_2^j \left(B_i^j - \hat{B}_i^j\right)\right)\frac{\Delta x_t}{x_t} + \left(\alpha_3 + \sum_j \alpha_4^j \left(B_i^j - \hat{B}_i^j\right)\right)\frac{\Delta p_t}{p_t + \overline{x}_t} + \varepsilon_{i,t}. \qquad (7)$$

where j denotes a bank activity. Our coefficients of interest in this estimation are the α_4^j. A positive α_4^j implies that balance activity j decreases tail risk, while a negative coefficient implies that it increases tail risk.

We estimate Eq. 7 by means of pooled OLS. Table 5 presents the implied coefficients for the determinants of gamma, δ^j $(= -\alpha_4^j)$. Note that a positive

---

[10]The two-step method, however, yielded very similar results.

Table 5 Relationship between gamma and bank characteristics. This table reports the negative coefficients of the interaction terms between the adjusted put option and the respective balance sheet variables. It represents the effect of the respective balance sheet item on a bank's tail risk exposure where a positive value implies a larger exposure to tail risk

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log(TA)	−2.722***	−2.934***	−2.681***	−2.708***	−2.530***	−3.237***	−2.571***
	(0.371)	(0.384)	(0.541)	(0.492)	(0.415)	(0.476)	(0.427)
Debt/TA	0.727	36.53	−21.45	−19.68	2.117	−6.373	−3.771
	(29.25)	(30.54)	(33.01)	(30.79)	(30.08)	(29.85)	(29.97)
Loans/TA	−23.12***	−12.46***	2.485		−23.02***	−21.28***	−23.16***
	(4.433)	(4.128)	(18.17)		(4.440)	(4.637)	(4.785)
Non-performing loans/TL		−106.3					
		(75.26)					
Loan loss allowance/TL		348.2					
		(225.6)					
Loan growth		19.10**					
		(9.178)					
Interest loans/TL		−432.5***					
		(37.47)					
ROA		569.3***					
		(57.88)					
Held-to-maturity securities/TA			25.62	23.33*			
			(23.37)	(13.38)			
For-sale securities/TA			31.25	28.64***			
			(21.01)	(6.360)			
Trading assets/TA			66.68*	63.06***			
			(34.97)	(20.20)			
C&I loans/TA				−0.0419			
				(7.882)			
Total securitization/TA					−3.396		
					(6.303)		

Table 5 (continued)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Loans sold/TA					−48.42		
					(41.47)		
Int. cred. exposure from sec. & sales					0.953		
					(1.486)		
Ext. cred. exposure from sec. & sales					147.4**		
					(64.64)		
Derivatives held for trading/TA						0.361***	
						(0.108)	
Derivatives not for trading/TA						3.092	
						(3.605)	
Time dep. > 100^7/domestic dep.							10.44**
							(4.746)
Deposits/liabilities							3.889
							(7.265)
Observations	154,242	154,242	154,242	154,242	154,242	154,242	154,242
R^2	0.230	0.239	0.230	0.230	0.230	0.230	0.230

Robust standard errors are reported in parentheses and significance is denoted as follows: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

coefficient in the table implies that the activity increases tail risk.[11] The first column contains the results from a regression with some basic bank characteristics: size (measured by the log of total assets), the loan-to-asset ratio and the leverage ratio (measured by the debt-to-asset ratio). Size is negatively related to tail risk exposure. This may indicate that markets perceive large banks as being too-big-to-fail (TBTF). The loan-to-asset ratio is also negatively related to a bank's tail risk exposure. This finding is in line with other recent findings: both De Jonghe (2010) and Demirgüç-Kunt and Huizinga (2010) find that traditional banking activities are less risky than non-traditional activities. The last variable considered is the leverage ratio. Although a higher leverage ratio is often associated with more default risk, it does not come out significant here (we return to this later).

Column two focuses on banks' lending activities by including proxies for loan quality and profitability. Among the loan quality proxies only the loan growth variable is significant, indicating a positive relationship with tail risk. This is consistent with the idea that a bank may only grow faster at the cost of lowering lending quality, and hence may become more exposed in a downturn.[12] We also find that a higher interest rate on the loans is associated with less tail risk, which can be explained by the fact that this indicates a higher profitability of banks, thus exposing them less to a crash in the market. Additionally, we include the return of assets (ROA) to capture the returns from other (partly non-traditional) asset activities. We find a positive relationship with tail risk, which is consistent with other recent findings (e.g., Demirgüç-Kunt and Huizinga 2010).[13]

Next, we turn to the influence of other assets. In column three we include held-to-maturity securities, for-sale securities and trading assets (all scaled by total assets). Only trading assets turn out significant, and only at the 10% level. At this point, one has to keep in mind that non-traditional activities are likely to be negatively correlated with traditional activities (banks may specialize in either), which may create multicollinearity problems and hence affect the estimates. Therefore, in column four we use the ratio of commercial and industrial loans to total assets (C&I Loans/TA) instead of the loan-to-asset ratio (the traditional activity) as it is less correlated with the non-traditional activities. The result is that trading assets and for-sale securities now contribute very significantly to tail risk. Held-to-maturity securities have a positive coefficient as well, but its magnitude and significance is lower. The C&I-loans-to-asset ratio is insignificant, similar to the loan-to-asset ratio in column three.

It is often argued that non-traditional activities increase (tail) risk exposure. In columns five and six, we will analyze which role financial innovations play among the non-traditional activities. First, we investigate securitization and asset sales activities. In addition to the total value of securitization and asset sales (both scaled by total assets) we also include the internal and external credit exposure arising from these

---

[11] The choice of the explanatory variables is motivated by the empirical literature on bank risk in normal times, see, for example, Stiroh (2006).

[12] This is in line with other studies, which identify loan growth as a main driver of risk (see, for example, Foos et al. 2009).

[13] Note that the interest income from loans is a part of the ROA so that potential multicollinearity issues could affect the results. However, tests in which we split the ROA into returns from loans and returns from remaining assets revealed that this is not a problem.

activities. The internal credit exposure arises from a bank's own securitization or asset sale activities via recourse and other credit enhancing agreements between the bank and its special purpose vehicle (SPV). An external credit exposure can arise if a bank provides credit enhancements to other banks' securitization structures.

Column five shows that only the external credit exposure variable is significant and positive. This is in line with our prior findings as external credit exposure is new credit exposure taken on in addition to existing exposure. Moreover, such exposure (for example, from credit enhancements) only materializes under relatively adverse scenarios, and hence should be related to tail risk. The insignificance of a bank's own securitization and asset sale activities may indicate that opposing forces are at work. On the one hand, securitization and asset sales are, by themselves, of course a mean of off-loading risk to other market participants, making a bank less risky.[14] In particular, if the bank keeps the equity tranche but sells senior tranches it sheds tail risk relative to normal period risk. On the other hand, recent experience has shown that these activities induced banks to take on more risk.[15] In addition, although the credit exposure seemingly disappeared from the balance sheet to the SPV (which is legally independent), the market might expect that this separation does not survive if the SPV encounters large losses. A bank might buy back assets from its SPV in order to protect its reputation and customer base (as happened in the case of Bear Stearns). Therefore, the credit exposure (which is mostly tail risk exposure) may not be effectively removed by securitization.

Column six focuses on banks' derivatives activities. Based on the available data, we can make the distinction between derivatives that are held for trading purposes and derivatives that are held for other purposes (most likely hedging). A priori one would expect that the latter would reduce tail risk. The expected effect for derivative trading is less clear cut. Resulting counterparty risk (which tends to materialize in tail risk scenarios) may, for example, create an increase in tail risk exposure. The results in column six show that derivatives held for trading contribute to tail risk, while the other derivatives do not seem to affect it. The latter is somewhat surprising but may be explained by the fact that only some of these derivatives are used for hedging and that they create counterparty risk as well.

The last column takes a closer look at the importance of capital structure for tail risk. In column one we found that the leverage ratio does not contribute to tail risk exposure. We now include information on the share of deposits and the composition of deposits. In the last column of Table 5, in addition to the variables from column one, we consider the deposit-to-liabilities ratio and the ratio of time deposits above \$100,000 to domestic deposits.[16] Time deposits above \$100,000 were not insured during our sample period, which makes them similar to wholesale funding, as both funding sources might be prone to runs. The results in column seven show that the leverage ratio is again not significant. Insignificance also obtains for the

---

[14]There is evidence from bank stock prices supporting this channel, see for example, Wu et. al. (2011). However, there is also evidence that securitized assets have lower risk (Ambrose et al. 2005); thus the relative riskiness of a bank's portfolio may increase.

[15]For example, Franke and Krahnen (2007) and Nijskens and Wagner (forthcoming) find that securitization increases a bank's beta.

[16]The FR Y-9C reports do not contain information on deposits in foreign subsidiaries, hence we scale by domestic deposits.

deposit-to-liabilities ratio. However, the time deposits above $100,000 do contribute positively and significantly to tail risk. Since these deposits are subject to withdrawal risks similar to wholesale funding, this result is consistent with Demirgüç-Kunt and Huizinga (2010) who find that wholesale funding increases bank risk.[17]

6 Conclusion

In this paper we propose a forward-looking method to measure tail risk exposures at banks. Tail risk is defined as a bank's exposure to a large negative market shock and it is measured by estimating a bank's share price sensitivity to changes in far out-of-the-money put options on the market, correcting for market movements themselves. Because far out-of-the-money put options on the market only pay out if the market crashes, changes in their prices reflect changes in the perceived likelihood and severity of a crash. The estimated sensitivities, in turn, represent the market's perception of exposures to a hypothetical crash, making them a truly forward-looking measure. Another attractive feature of this measure is that it does not require the actual observation of tail risk events since it identifies banks' tail risk exposure through changes in expected market tail risk. Our measure is also relatively easy to estimate as it basically comes from an amended market regression.

The application to U.S. bank holding companies yields several interesting facts about their tail risk exposures. For example, (excess) tail risk seems to be negatively correlated with the CAPM share price beta. This suggests that banks which appear relatively safe in normal times (that is, have a low beta) are actually riskier than their beta would suggest. We also find that the impact of non-traditional activities on tail risk depends on whether they leave assets on the balance sheets or not. In the former case they increase tail risk, while in the latter they do not. Our results also suggest that leverage itself does not increase tail risk, but will do so if it comes through uninsured deposits.

References

Acharya V, Pedersen LH, Philippon T, Richardson M (2009) Restoring financial stability, 1st edn. Wiley, New York

---

[17]Demirgüç-Kunt and Huizinga do not distinguish between normal times risk and tail risk but focus instead on the Z-score.

Adrian T, Brunnermeier MK (2009) CoVaR. Mimeo, Princeton University

Agarwal V, Naik N (2004) Risks and portfolio decisions involving hedge funds. Rev Financ Stud 17:63–98

Brent A, LaCour-Little M, Sanders A (2005) Does regulatory capital arbitrage, reputation, or asymmetric information drive securitization? J Financ Serv Res 28:113–133

Ang A, Hodrick RJ, Xing Y, Zhang X (2006) The cross-section of volatility and expected returns. J Finance 51:259–299

Bushman R, Williams C (2009) Accounting discretion, loan loss provisioning and discipline of banks' risk-taking. Mimeo, University of North Carolina at Chapel Hill

De Jonghe O.G. (2010) Back to the basics in banking? A micro-analysis of banking system stability. J Financ Intermed 19(3):387–417

Demirgüç-Kunt A, Huizinga H.P. (2010) Bank activity and funding strategies: the impact on risk and return. J Financ Econ 98(3):626–650

Flannery MJ (1998) Using market information in prudential bank supervision: a review of the U.S. empirical evidence. J Money, Credit Bank 30(3):273–305

Flannery MJ (2001) The faces of "market discipline". J Financ Serv Res 20(2–3):107–119

Foos D, Norden L, Weber M (2009) Loan growth and riskiness of banks. Working paper, Bank of International Settlements

Franke G, Krahnen JP (2007) Default risk sharing between banks and markets: the contribution of collateralized debt obligations. In: Carey M, Stulz R (eds) The risks of financial institutions. National Bureau of Economic Research conference report

Hovakimian A, Kane E (2000) Effectiveness of capital regulation at U.S. commercial banks, 1985 to 1994. J Finance 55:451–468

Hasan I, Wall LD (2004) Determinants of the loan loss allowance: some cross-country comparisons. Financ Rev 39:129–152

Huang X, Zhou H, Zhu H (2010) Systemic risk contributions. Working paper, Bank of International Settlements

Huizinga H, Laeven L (2009) Accounting discretion of banks during a financial crisis. IMF Working Paper WP/09/207

Jorion P (2006) Value at risk, 3rd edn. McGraw-Hill, New York

Laeven L, Majnoni G (2003) Loan loss provisioning and economic slowdowns: too much, too late? J Financ Intermed 12:178–197

Standard and Poor's (2005) Chasing their tails: banks look beyond value-at-risk: RatingsDirect. http://www2.standardandpoors.com/portal/site/sp/en/us/page.article/2,1,6,4,1126533408950.html. Accessed 26 Nov 2008

Stiroh K (2006) New evidence on the determinants of bank risk. J Financ Serv Res 30:237–263

Wall LD, Koch TW (2000) Bank loan loss accounting: a review of theoretical and empirical evidence. Fed Reserve Bank Atlanta Econ Rev 85(2):1–19

Wu D, Yang J, Hong H (2011) Securitization and banks' equity risk. J Financ Serv Res 39:95–117