Predicting Bank Failures Using a Market-based Measure of Capital

Keith Friend*
Economics Department
Office of the Comptroller of the Currency

Mark Levonian
Promontory Financial Group

August 21, 2013

Abstract

Supervisors rely on a variety of sources of information to identify problem banks, including non-public information acquired through examinations and ongoing supervision. However, information from financial markets can be a valuable supplemental source of information about the condition of banks. In this paper we consider the ability of equity market data to predict which banks are most likely to fail. We focus on bank failures during the recent financial crisis, and the extent to which equity market signals may have provided valuable early warning about the likelihood of bank failure. We find that signals of bank condition based on equity prices are somewhat more accurate in predicting bank failures than are the regulatory ratios reflected in PCA or measures like the Texas ratio, although not markedly so. Perhaps more importantly, we find that market signals identify failing banks much farther in advance of failure, potentially providing more time for responses that would reduce the cost of such failures. We also demonstrate the potential benefits from thoughtful "tuning" of market measures to predict failures, as well as the trade-offs involved in such tuning. Finally, we present preliminary evidence that market-based measures became less accurate predictors during the recent financial crisis, particularly for larger banks.

---

*Corresponding author.

The views expressed are those of the authors, not the views of the Office of the Comptroller of the Currency, the U.S. Department of Treasury, or Promontory Financial Group. This research was conducted while Levonian was Senior Deputy Comptroller for Economics at the OCC. The authors are grateful to Dilip Patro, Ajay Palvia, and Gary Whalen for helpful comments on an earlier draft.

1 - Introduction

U.S. banking regulators closed approximately five hundred FDIC-insured institutions between 1996 and the beginning of 2012. The vast majority of these – over four hundred – failed after 2007, as banks and thrifts felt the impact of the financial crisis. Bank failures are costly, and such episodes of relatively high rates of bank failure help highlight the early identification of potential failures as a key aim of banking supervision and regulation. Prompt Corrective Action (PCA) and other aspects of supervision are intended to identify and address problems as they emerge, to avert failure when possible, and to make failures less costly when they do occur.

Supervisors rely on a variety of sources of information to identify problem banks, including non-public information acquired through examinations and ongoing supervision. However, information from financial markets can be a valuable supplemental source of information about the condition of banks. Financial market participants have strong incentives to gather and evaluate information to support decisions, and the information reflected in their actions (as evident, for example, in prices or interest rates) reflect the combined views of a large number of analysts. Although market participants may lack access to some of the confidential information available to bank supervisors, as long as their actions convey at least some information that is not already reflected in the information set of bank supervisors, the market information can be used to improve banking supervision.

Flannery (1998) provides a review of the evidence on market information in prudential bank supervision; Flannery notes that the weight of evidence supports the proposition that information possessed by analysts and other market participants about the condition of financial firms can be useful. Gunther, Levonian, and Moore (2001) focus specifically on equity-based measures, and find that information from equity markets has valuable predictive power even after taking into account the information bank regulators already have.

In this paper we consider the ability of equity market data to predict which banks are most likely to fail. We focus on bank failures during the recent financial crisis, and the extent to which equity market signals may have provided valuable early warning about the likelihood of bank failures. We find that signals of bank condition based on equity prices are somewhat more accurate in predicting bank failures than are the regulatory ratios reflected in PCA or measures like the Texas ratio, although not markedly so. Perhaps more importantly, we find that market signals identify failing banks much farther in advance of failure, potentially providing more time for responses that would reduce the cost of such failures.

2 - Data Description

2.1 - Overview

Option pricing theory provides a method for computing a market-based measure of a bank's capital position.[1] This measure, referred to in this paper as the "market capital ratio" (MCR), incorporates market knowledge as expressed in the level and volatility of an individual bank’s share price to adjust the book value of assets to reflect a market perspective on the true value of the bank’s assets on a going concern basis.[2] Because it is based on market information, this market capital ratio may be more forward-looking than more traditional capital measures based on reported financial data. This paper extends previous work demonstrating that the MCR is a meaningful measure of a bank's capital position.[3]

Use of a market-based indicator requires that the analysis be limited to entities for which market data are available, in this case publicly traded bank holding companies (BHCs) whose stock market information is available through the Center for Research in Security Prices (CRSP). Equity data from CRSP are used to calculate quarterly stock return volatilities, quarterly dividends, and end-of-quarter market capitalizations for each BHC in the sample. This information is then merged with regulatory balance sheet data to calculate MCR for each quarter. Since top-tier holding companies are the stock-issuing entities, the regulatory data are sourced from the Federal Reserve's quarterly Y9-C report. Thus, data requirements restrict the sample to banking entities that both file a Y9-C and also issue publicly traded stock covered in the CRSP database.

The resulting sample comprises 902 distinct bank holding companies that reported to the Federal Reserve during the sixteen year period beginning in 1996 and continuing through 2012.[4] The holding companies range in asset size from just under $100 million to over $2.3 trillion. The sample covers a large share of the assets of the U.S. banking industry, ranging between 79 percent and 85 percent of total BHC assets depending on the sub-period (Table 1).

[1] For a discussion of methods and the underlying theory see Appendix.
[2] The ratio is equal to one minus the ratio of book debt to market implied value of assets. The market implied asset value is found by assuming the bank's equity is a call option on assets (Merton model). Stock market data provide measures of the value of equity and equity volatility.
[3] Friend and Levonian, "The Relationship between Market Implied Asset Volatility and Capital Ratios in U.S. Bank Holding Companies" (draft manuscript)
[4] Public holding company data as of April 2012.

Time Period	*BHC Assets in Sample*
1996 - 2000	85%
2001 - 2002	82%
2003 - 2004	80%
2005 - 2006	79%
2007 - 2008	79%
2009 - 2010	83%
2011 - 2012	84%
Total	82%

Table 1 Total assets held by bank holding companies in the sample as a percentage of all BHC assets reported to the Federal Reserve.

Of the roughly 500 banks closed by the FDIC since January 1996, 69 were owned by holding companies in the sample.[5] Information on failure and the date of closure for the holding companies' insured subsidiaries is available from the FDIC. Although the sample covers 69 failed banks, these comprise only 55 instances of failure,[6] since some of the failed banks are affiliates that were closed at the same time. Because stock market data from the parent are used to compute the market capital ratio, these multiple failures are combined into one instance for each BHC.

Time Period of Failure	*Number of Failures in Sample*	*Total Number of Failures*
1996 - 2000		25
2001 - 2002	2	15
2003 - 2004	1	7
2007 - 2008	3	28
2009 - 2010	49	297
2011 - 2012	14	132
Total	69	504

Table 2 Incidence of failure in sample over time.

2.2 - Signal Variable Construction

The market capital ratio is an indicator of bank condition that, as noted above, may have desirable properties due to the information it likely embodies. The MCR could be used as an early warning signal

[5] Failed institution data comes from the FDIC website as of August 2012.

[6] The data reflect BHC ownership at the time each insured institution failed, and do not account for changes of ownership that may have occurred at earlier points in time.

for firms that are at risk of failure. In theory, if a bank holding company's MCR is less than or equal to zero for a particular quarter, one interpretation is that the market is indicating that the bank is insolvent. More generally, if the MCR falls below some threshold, this may serve as an indicator of a market view that the firm is troubled. We begin with a signal threshold of zero for the MCR, then consider other possible values of the threshold in a later section. This market-based signal can be compared to signals provided by some of the more traditional risk-based capital measures used by regulators.

Prompt corrective action thresholds are defined according to three risk-based solvency ratios: tier 1 leverage ratio, tier 1 capital ratio, and total capital ratio (Table 3).[7] For purposes of this paper, a PCA signal occurs when a BHC falls into one of the "Undercapitalized" categories.

PCA Indicator	*PCA Category*	*Leverage Ratio*	*Tier 1 Ratio*	*Total Ratio*
None	Well Capitalized	≥ 5%	≥ 6%	≥ 10%
	Adequately Capitalized	≥ 4%	≥ 4%	≥ 8%
Signal	Undercapitalized	< 4%	< 4%	< 8%
	Significantly Undercapitalized	< 3%	< 3%	< 6%

Table 3 Prompt Corrective Action Classifications and Indicator Variable Definition.

3 - Comparative Analysis of MCR and PCA Signals

3.1 - Relation between Capital Signals and Failure

Signals provide potentially useful information if they are related to subsequent survival or failure experience. Applying the criteria discussed above for failure signals – undercapitalized or worse for PCA, and a threshold of zero for MCR – reveals that 89 of the 902 bank holding companies registered a PCA signal in the sample time period, while 180 of the 902 registered at least one MCR signal.[8] The contingency tables in Table 4 below classify the sample companies according to the presence or absence of a signal from MCR or PCA, and whether the company subsequently failed during the sample period or survived (that is, did not fail). Of the 55 bank holding companies that failed, 50 of them had at least one MCR signal, while 43 had at least one PCA signal. Statistical tests reject independence between survival status and either of the signals, implying that the presence of a signal of either type is related to failure (Table 5). Both frequency tables suggest that the odds of a bank failing are much higher for the group that experienced a signal.

[7] OCC Banking Circular on Prompt Corrective Action: <http://www.occ.gov/static/news-issuances/bulletins/pre-1994/banking-circulars/bc-1993-268.pdf>

[8] Several banks registered a second or third MCR signal if the ratio fell below zero, then rose above zero before falling again.

Table 6 shows that the odds of a bank failing are approximately fifty times greater if one or more MCR signals are observed than if no signal is seen. The MCR signal detected 91 percent of failures, while PCA signaled only 78 percent. However, the MCR signal had a 15 percent false positive rate, compared to 5 percent for PCA. The estimate for the odds ratio for the PCA signal is somewhat higher than for MCR, but the difference is not statistically significant, as could be surmised from the wide confidence intervals for the odds ratios shown in Table 5.[9] Thus, for the banks in this sample, the market capital ratio and risk based capital measures are fairly similar in their ability to distinguish between failures and survivors.

Signal by Survival Status Contingency Tables

Market Capital Ratio

	Failed		*Survived*		*Total*	
	Number of BHCs	Column Percent	Number of BHCs	Column Percent	Number of BHCs	Column Percent
No Signal	5	9%	717	85%	722	80%
One or more signals	50	91%	130	15%	180	20%
Total	55	100%	847	100%	902	100%

Prompt Corrective Action

	Failed		*Survived*		*Total*	
	Number of BHCs	Column Percent	Number of BHCs	Column Percent	Number of BHCs	Column Percent
No Signal	12	22%	801	95%	813	90%
One or more signals	43	78%	46	5%	89	10%
Total	55	100%	847	100%	902	100%

Table 4 Contingency tables for each type of capital signal. Columns in each table separate banks based on survival status, while rows distinguish between signal status.

Independence Tests

Signal Type	*Pearson's χ^2*	*Fisher's Exact*
MCR	<.0001	3.324E -32
PCA	<.0001	1.604E -37

Table 5 P-values for statistical tests of the null hypothesis that signal and survival status are independent.

Signal Type	*Odds Ratio*	*95% Confidence Limits*	
MCR	54.73	21.59	140.92
PCA	62.47	30.81	126.34

Table 6 Odds ratio estimates for each capital signal.

3.2 - Signal Lead Time

Identification of troubled banks becomes more valuable as the time between a signal and a bank's failure increases. Greater signal lead time provides both bankers and regulators with more opportunity

[9] An explicit statistical test confirms that the null hypothesis of the signals having equal odds ratios cannot be rejected.

to work through problems and prevent closure. In this section we consider differences in the lead time provided by MCR and PCA signals.

Calculating lead time for the MCR signal is complicated by the fact that troubled institutions may experience multiple signals as the market capital ratio fluctuates around zero. Figure 1 provides two examples of banks with multiple signals.

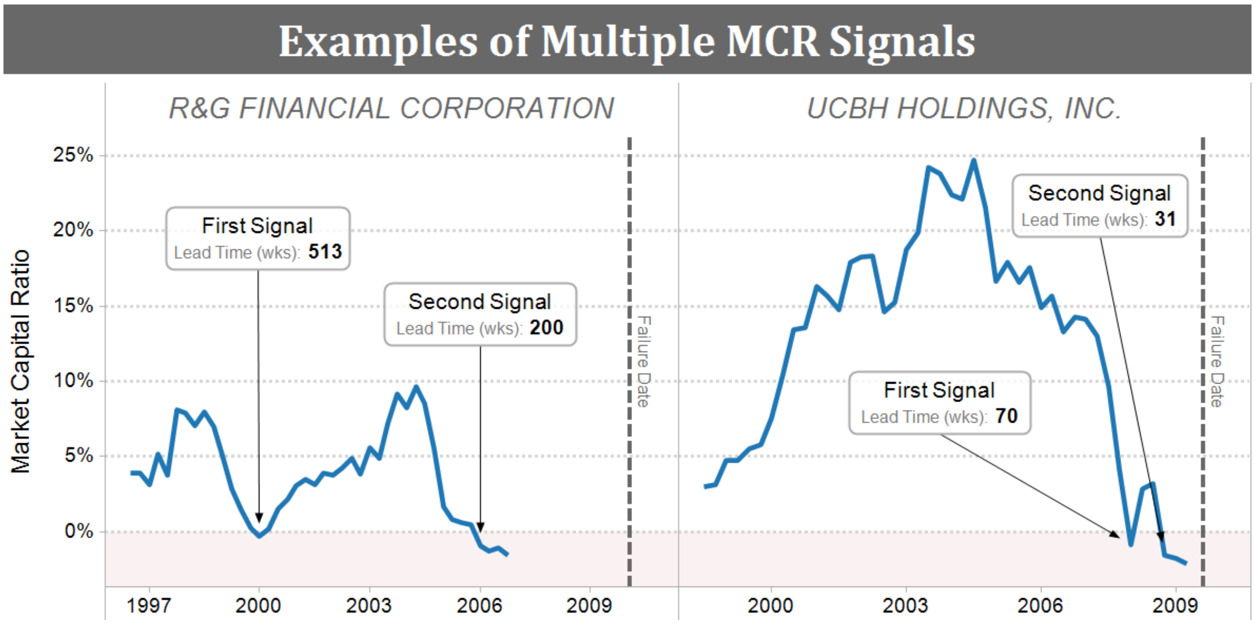

Figure 1 **Each of the banks experiences at least two MCR drops below zero before the date of failure.**

Examples like R&G Financial Corporation on the left in Figure 1 favor using the last signal date, since the first signal precedes it by almost six years and is likely unrelated. The example on the right in Figure 1 shows UCBH, for which the MCR signals are separated by roughly three quarters, making it seem more likely that both are being driven by the same fundamentals. In this case the first signal might appropriately be viewed as an early signal of the ultimate failure.

Distinguishing signals that are unrelated from those that should be grouped together creates a difficult challenge in trying to determine a correct measure of signal lead time. In practice, setting a minimum time between distinct signals (grouping those that are "close together" in time) could help. Alternatively, use of a "recovery" threshold above which MCR must rise before the signal is deemed to be "switched off" might better allow for potential volatility in the market-based measure.

Either approach would require arbitrary decisions. For this analysis, where there are cases of multiple signals for a given failure, signal lead time is defined as the shortest lead time among a bank's multiple signals. This approach may tend to understate the actual lead time regulators would have in practice, but eliminates the risk of overstating lead time by disregarding stale signals that precede failure by an unreasonably large interval.

Eight of the 55 failures were preceded by more than one MCR signal. Figure 2 shows each of the fifty-five failures and compares the minimum signal lead time given by the prompt corrective action indicator to the lead time provided by the market capital ratio signal. Figures to the right of each bar show the difference in lead time (in weeks) provided by MCR over PCA.

The results suggest that the signal lead time is longer for MCR than for PCA. The mean lead time for PCA signals is 39 weeks (median: 30 weeks), whereas the mean lead time for MCR signals is 73 weeks (median: 64 weeks).[10] Focusing more directly on the 43 failures that were signaled by both PCA and MCR, the mean difference in lead time is 32 weeks, with a 95 percent confidence interval of 24 to 40 weeks (the median is 26 weeks, with a 95 percent confidence interval of 26 to 39 weeks). This difference implies that the MCR provides a signal of trouble roughly eight months earlier than the PCA ratios. A paired T-test on the 43 banks that have lead times for both types of signals reveals that this average difference in lead times is significantly different from zero (see Table 7). The non-parametric sign and Wilcoxon signed rank tests also confirm that the lead times for MCR are larger than those for PCA.[11]

Paired Difference Tests

Null Hypothesis	*Test Name*	*Assumption*	*Test Stat*	*P-Value*
Mean diff. = 0	Paired T-Test	Normal Dist.	7.955	<.0001
Median diff. = 0	Signed Rank	Symmetric Dist.	333	<.0001
	Sign	Continuous Dist.	18	<.0001

Table 7 Both the paired T test and the non-parametric tests confirm that the mean difference between the signal lead time for MCR and PCA signals is significantly different than zero.

[10] Interestingly, both types of signals tend to have longer lead times for more recent failures. Failures that occurred after the second half of 2010 are almost always associated with lead times that are above average.

[11] The distribution of the differences between the signal lead times is not symmetric around zero since the MCR signal lead time is strictly greater than the PCA lead time, i.e. all of the differences are positive. Although the T-test and Wilcoxon test are shown in Table 7, the sign test is most appropriate given the skewed nature of the data.

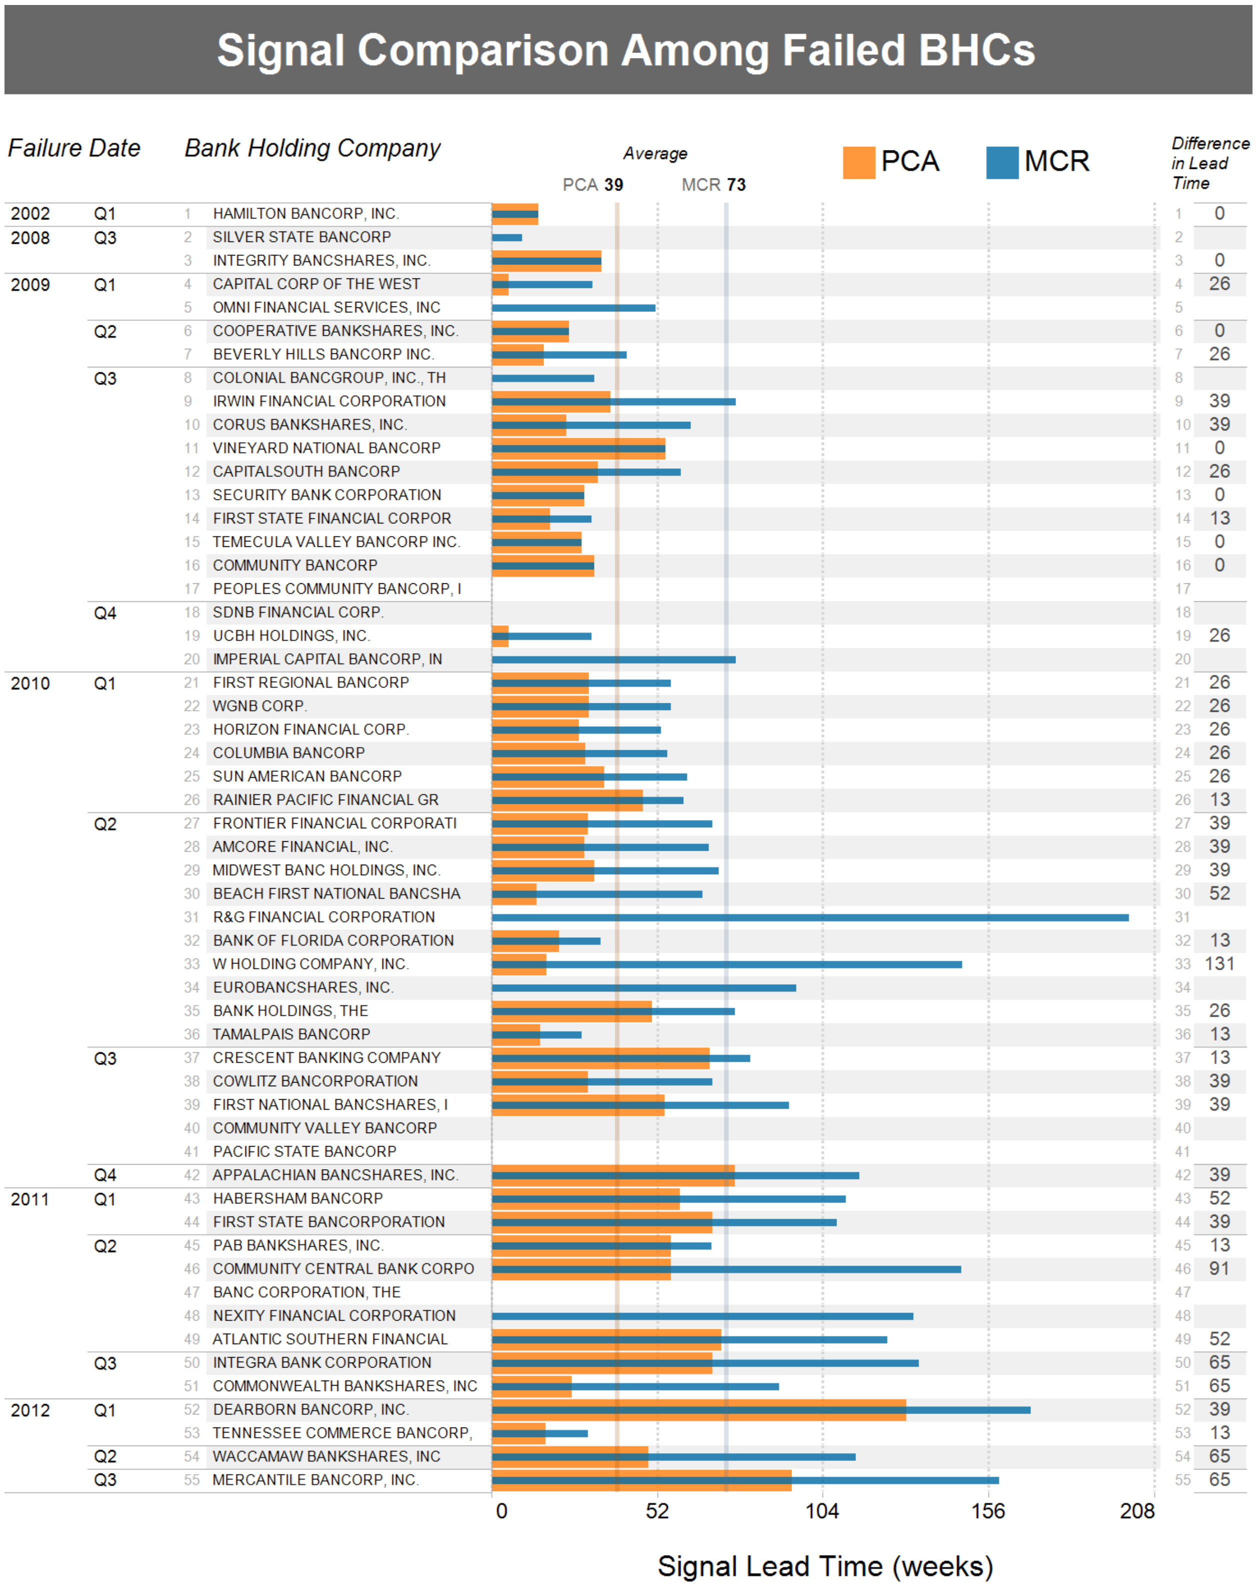

Figure 2 **Lead time for each signal type by failed bank (figures to the right of the bars show the difference, in weeks). Not all of the BHCs in the sample had signals before they failed.**

4 - Texas Ratio Comparison

The Texas ratio (TR) provides an alternative financial-statement-based measure for comparison to MCR. With its focus on the relation between troubled assets and the resources available to absorb losses, it is plausibly related to the probability of bank failure.

Signal by Survival Status Contingency Tables

Market Capital Ratio

	Failed		Survived		Total	
	Number of BHCs	Column Percent	Number of BHCs	Column Percent	Number of BHCs	Column Percent
No Signal	5	9%	717	85%	722	80%
One or more signals	50	91%	130	15%	180	20%
Total	55	100%	847	100%	902	100%

Texas Ratio

	Failed		Survived		Total	
	Number of BHCs	Column Percent	Number of BHCs	Column Percent	Number of BHCs	Column Percent
No Signal	10	18%	804	95%	814	90%
One or more signals	45	82%	43	5%	88	10%
Total	55	100%	847	100%	902	100%

Table 8 Contingency tables for the two early warning signals.

The TR signal catches 82 percent of the failures, less than the MCR signal, but with a lower false positive rate. Repeating the statistical tests from section 3.1 shows not only that the TR signal is associated with survival status, but also that such a signal increases the odds of observing a failure 84.14 times. As was the case with PCA, the TR signal's odds ratio is not significantly different from the MCR signal's odds ratio, meaning that both are similarly able to distinguish survivors from failures in the sample.[12]

Comparing lead time between MCR and TR signals yields results similar to section 3.2 for PCA, although less pronounced. Unlike PCA, TR signals did have longer lead times than MCR in some cases. Still, lead time was on average 18 weeks longer for MCR among the 45 failed institutions that were picked up by both MCR and TR.[13] Paired T, sign, and signed rank tests all confirm that MCR lead times are significantly longer than those for TR.

[12] Ninety-five percent confidence interval for the TR odds ratio is 39.71 to 178.26, which overlaps with the other signals' intervals.

[13] Ninety-five percent confidence interval for the mean difference is 10 to 27 weeks; the median difference is 13 weeks with a confidence interval of 0 to 26 weeks.

5 - Signal Threshold Calibration

Up to this point the MCR signal has been based on whether or not a bank holding company's MCR drops below zero during any quarter. Using zero as the threshold is a natural choice since, at least in theory, banks below this level could be thought of as insolvent. However, the choice of zero as the threshold is to a certain extent arbitrary, and other thresholds might better differentiate between failures and survivors.

A more general approach takes the signal threshold as a parameter, denoted *k*. Banks whose MCR drops below *k* trigger a signal of future failure, whereas banks that never go below *k* do not. While in theory the parameter space for *k* is semi-infinite – that is, *k* could take any value less than 1 – for estimation only *k* values within the range of MCR values observed in the sample need to be considered.[14] As the threshold increases, more banks in the sample trigger a positive signal at some point during the sample period, until the threshold exceeds the minimum MCR of every bank in the sample. At each threshold, a contingency table can be constructed similar to the left panel in Table 4 above. The odds ratio for each of these tables is a relatively natural measure of how well each alternative threshold performs in identifying banks that fail. The threshold with the highest odds ratio could be considered optimal.

[14] If the minimum MCR value for the sample across banks and over time is negative 5 percent, then a threshold value of negative 10 percent will yield the same contingency table as one of negative 20 percent. Both thresholds are too low to ever trigger a signal for any bank in the sample. Setting the threshold to just above this minimum value will result in one positive signal for the bank with the lowest MCR value; all the other banks will still have no signal.

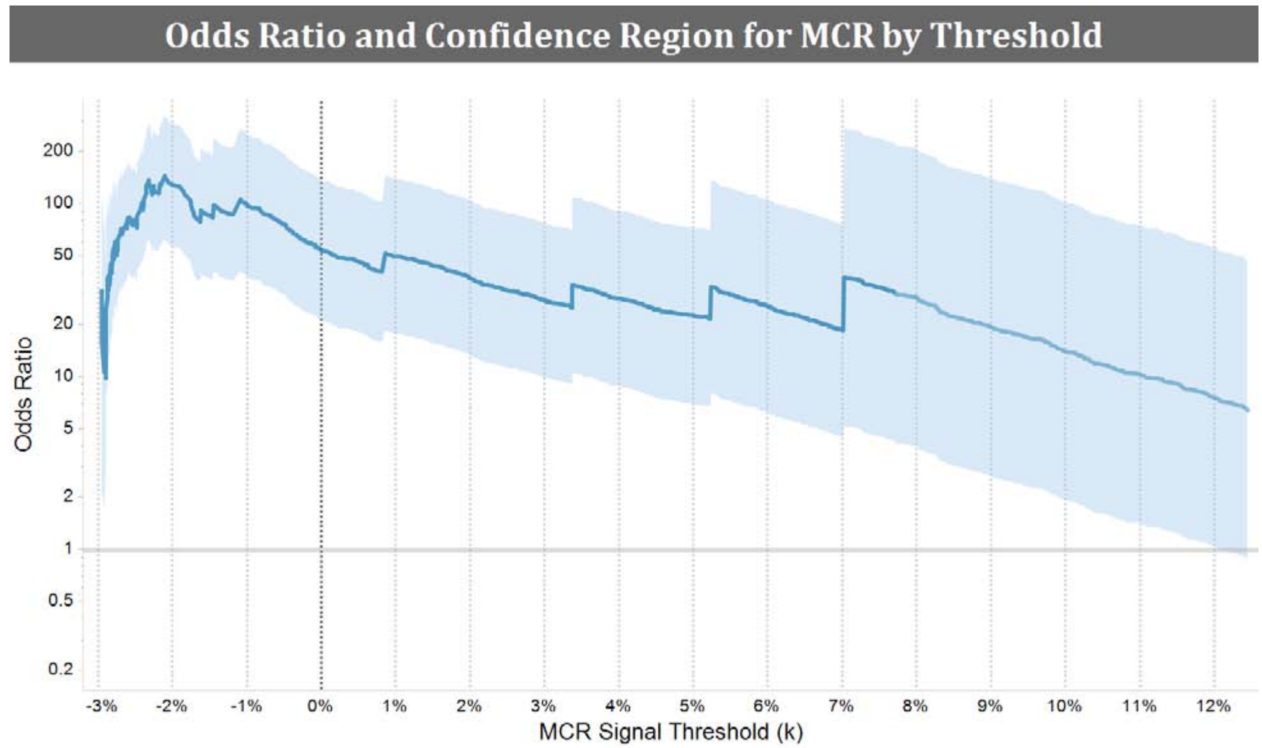

Figure 3 **The odds ratio of the MCR signal for different possible threshold values, k. The 95 percent confidence region is shown by the shaded area.**

Figure 3 displays the odds ratios that result from alternative values of the MCR signal threshold, with a 95 percent confidence interval indicated through shading. The odds ratio reaches its highest point at a threshold value around negative 2 percent. As the threshold increases beyond this value, the odds ratio falls. This gradual decline is interrupted by instances of failed banks that had relatively high minimum MCRs prior to failure; as the signal threshold is raised enough to capture these previously undetected failures, the odds ratio jumps. However, as the threshold continues to rise, more and more surviving banks are classified as failures by the signal, resulting in the overall downward trend. By the time the threshold is as high as 12 percent the confidence region for the odds ratio estimate overlaps with one, suggesting that a hypothesis that the odds of failure for the signal group are the same as the odds of failure for the no-signal group would not be rejected.

The odds ratio is a convenient measure of predictive ability in that it is easily interpretable. Unfortunately, it suffers from high variance as the number of errors (either false positives or false negatives) becomes small. This high variability in the estimation can be seen from the wide confidence region when the threshold value is around negative 3 percent or greater than positive 7 percent. At these thresholds, there is either only one false positive or only one false negative.

Alternative performance measures commonly used in binary classification problems are sensitivity and specificity. Sensitivity, or the 'hit' rate, is the percentage of failures correctly classified; it is calculated as the number of true positives (banks that failed and were predicted to fail) divided by the sum of true positives and false negatives (all banks that failed, both those with and without a signal). High sensitivity means that the signal is correctly identifying most of the failures. Specificity summarizes performance among the banks that survived, through the percentage of surviving banks that were *not* classified as failures. It is calculated as the number of true negatives (banks with no signal that survived) divided by the total number of survivors. High specificity occurs when a signal rarely generates false alarms.

A natural trade-off exists between sensitivity and specificity. Increasing the threshold will necessarily increase sensitivity since more banks will trigger signals, while simultaneously decreasing specificity because of the rise in false positives.

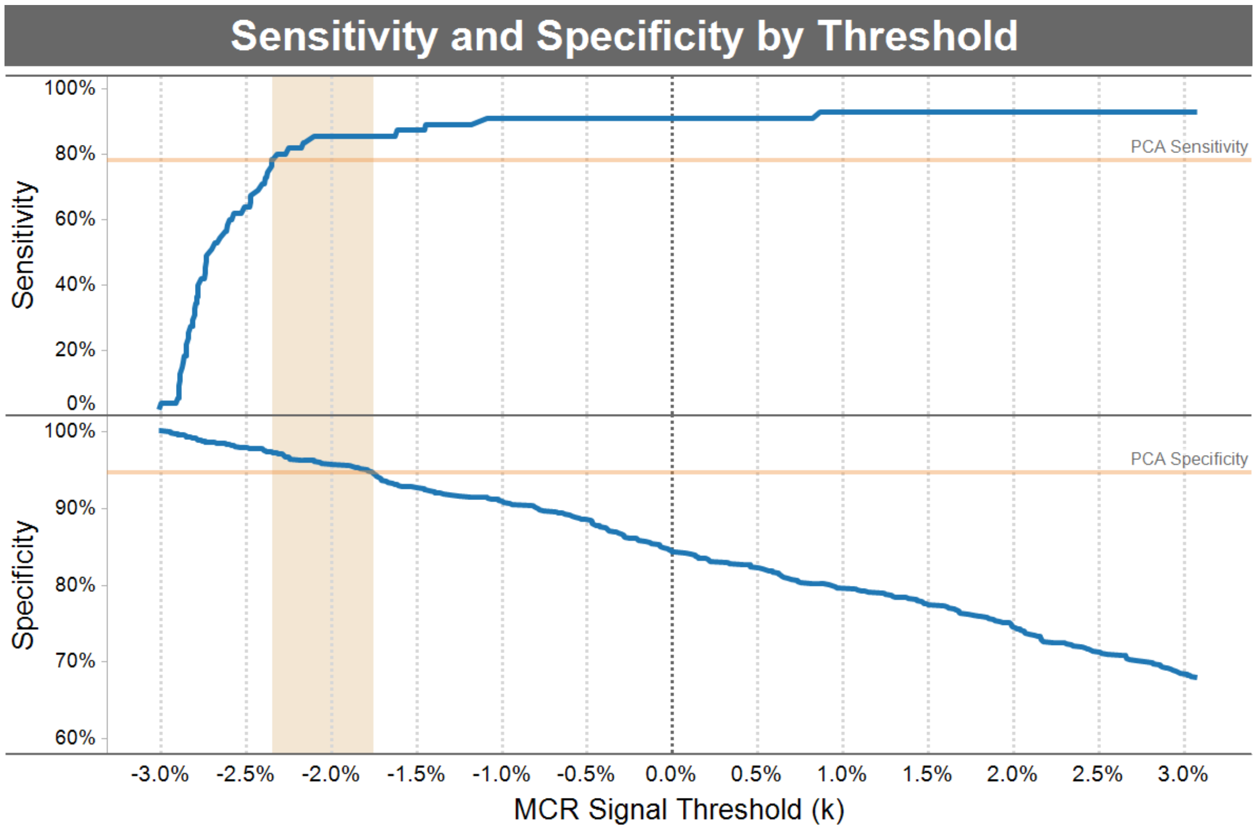

Figure 4 **Sensitivity and specificity as a function of threshold. The shaded region includes all the threshold values for which both sensitivity and specificity are higher for the MCR than they are for the PCA signal.**

The top panel of Figure 4 shows the increasing sensitivity of the MCR signal in blue and the constant sensitivity of the PCA signal in orange. When the threshold is greater than approximately -2.3 percent the MCR signal has higher sensitivity than PCA, i.e. the MCR signal is more sensitive to failures. The lower panel, which plots the relationship between specificity and the MCR signal threshold, shows that the specificity of the MCR signal is superior when the threshold is less than about -1.7 percent. As long as the MCR threshold is below this value its signal will generate fewer false alarms than the PCA signal. The shaded region between these two points represents thresholds for which MCR signals distinguish failure better than PCA in terms of both sensitivity and specificity.[15] There is no region in which the PCA signal outperforms MCR in both sensitivity and specificity.

[15] The cost of false negatives may differ from the cost of false positives; any such difference is not taken into account in this analysis.

Another important statistic that measures the degree of association between the signal and failure for each of the possible MCR thresholds is the phi coefficient, derived from the Pearson's chi-squared statistic used in Table 4. The phi coefficient for a 2x2 contingency table, also known as the Matthew's coefficient, is:

$$\varphi = \sqrt{\frac{\chi^2}{n}}$$

An MCR signal threshold may be considered optimal if it produces a higher value of this statistic than all other possible thresholds.

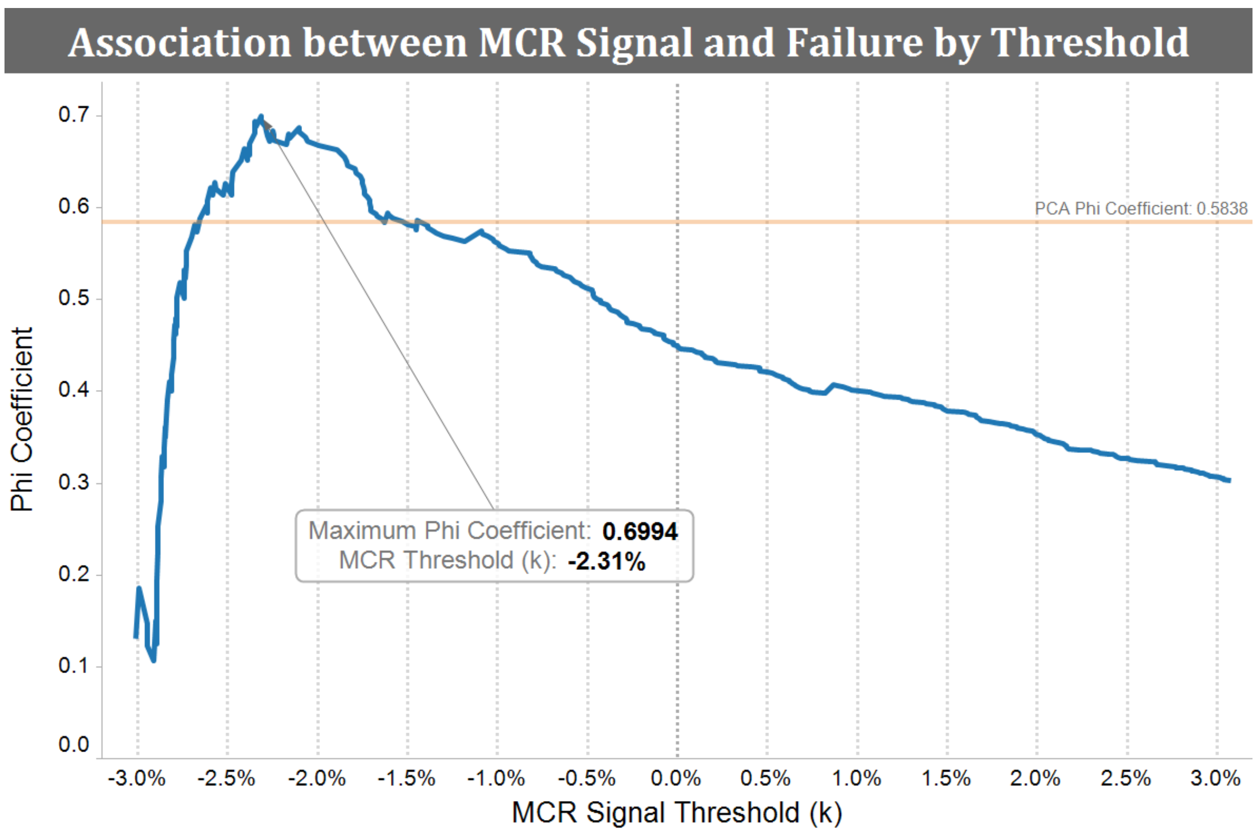

Figure 5 The phi coefficient reaches its maximum value when the MCR threshold is -2.31 percent.

Figure 5 plots the phi coefficient across different thresholds. Local peaks occur when the MCR signal is able to correctly detect a new failed bank. The global maximum of phi for the sample occurs at a threshold value of -2.31 percent, a value contained within the shaded region from Figure 4. A bootstrapping procedure can be used to estimate a confidence interval for this optimal threshold; the 95 percent confidence bounds for the threshold are -2.45 to -2.10 percent.[16]

[16] A distribution of optimal thresholds is simulated by performing the phi maximization procedure on each of ten thousand replicates generated using stratified random sampling with replacement.

This threshold is optimal in the sense that it is highly correlated with whether or not a bank will fail (more so than an MCR signal at any other threshold). The sensitivity and specificity are both higher than they are for PCA. Unfortunately, much of the lead time advantage that the MCR signal had over PCA when the threshold was set to zero disappears. Figure 6 shows average lead time for the MCR signal as the threshold is varied. (This is based on the minimum lead time, using the latest signal date for cases in which the signal is triggered multiple times prior to a failure.)

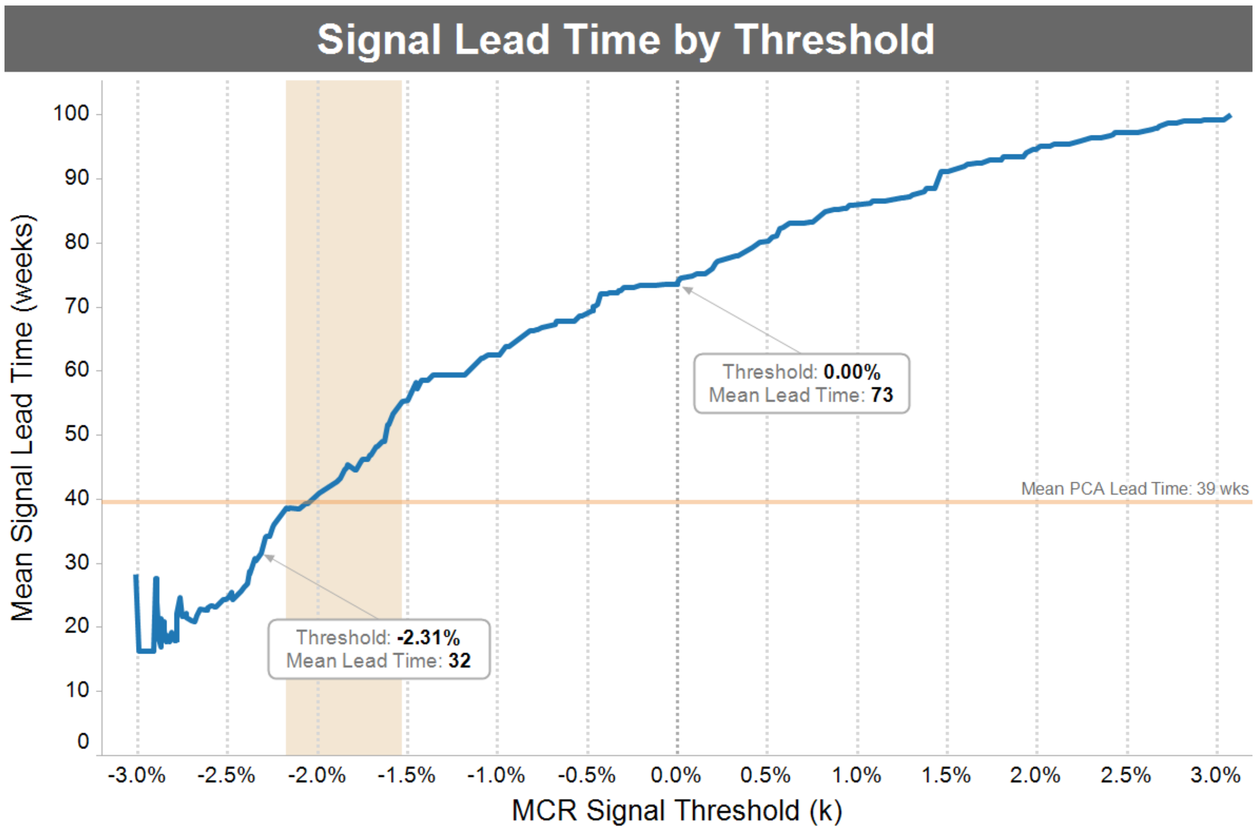

Figure 6 Mean lead time of the MCR signal as a function of its threshold. The shaded region shows thresholds at which the MCR and PCA lead times are not statistically different (paired sign test at 5% significance level).

When the threshold is set to -2.31 percent the average lead time is about 32 weeks, compared to about 73 weeks when it is set to zero. If increasing lead time is an important objective in calibrating the MCR threshold, then the other statistics considered in this section − odds ratio, sensitivity, specificity, phi coefficient − are insufficient for optimization, since these statistics ignore timing effects.

Throughout this analysis, all classification errors have been treated as equally costly. In practice type I errors (false negatives) likely are more costly than type II errors (false positives, which are "false alarms").[17] A type I error occurs when a bank fails without an advance signal. These unanticipated failures are likely to impose greater costs on the deposit insurance fund, bank customers and other bank stakeholders, and the broader economy than would failures for which there was sufficiently advanced

[17] This follows the conventional error-type definitions used in the bank failure literature, for example Thomson (1991), Gilbert *et al* (2000), and Cole and White (2011).

notice. Conversely, the cost of a type II error primarily stems from costs related to the additional scrutiny a healthy bank would face if regulators wrongly believed it was likely to fail. These false alarms may result in needless expenditure of supervisory (and perhaps bank) resources, but do not result in losses for the deposit insurance fund or other stakeholders, and thus are likely to be less costly than an unexpected failure.

6 - Factors Affecting False Positives

False positives occur whenever the MCR signal indicates a bank will fail and it does not. As the specificity of the MCR signal drops, more and more false positives will occur. However, these false positives do not occur uniformly throughout the sample. By defining the false positive rate as one minus the specificity (or the number of false positives divided by the total number of non-failing banks) cross sections of the sample can be compared to see whether any identifiable factors are associated with higher or lower false positive rates.

Figure 7 shows the false positive rate over time. Banks in the denominator for each quarter comprise only those active within that quarter that do not subsequently fail during the sample period. All time periods exhibit rates under 5 percent except 2008 Q4 and 2009 Q1, periods in which the false positive rate spiked to nearly 8 and 14 percent, respectively.

One possibility is that the significant equity market turmoil during this period made equity price movements less informative, leading to a jump in the number of false MCR signals across all banks. However, extraordinary actions taken by regulators to stabilize the financial system may have saved some banks that would otherwise have failed. This also would lead to an apparent increase in the rate of “false” signals.

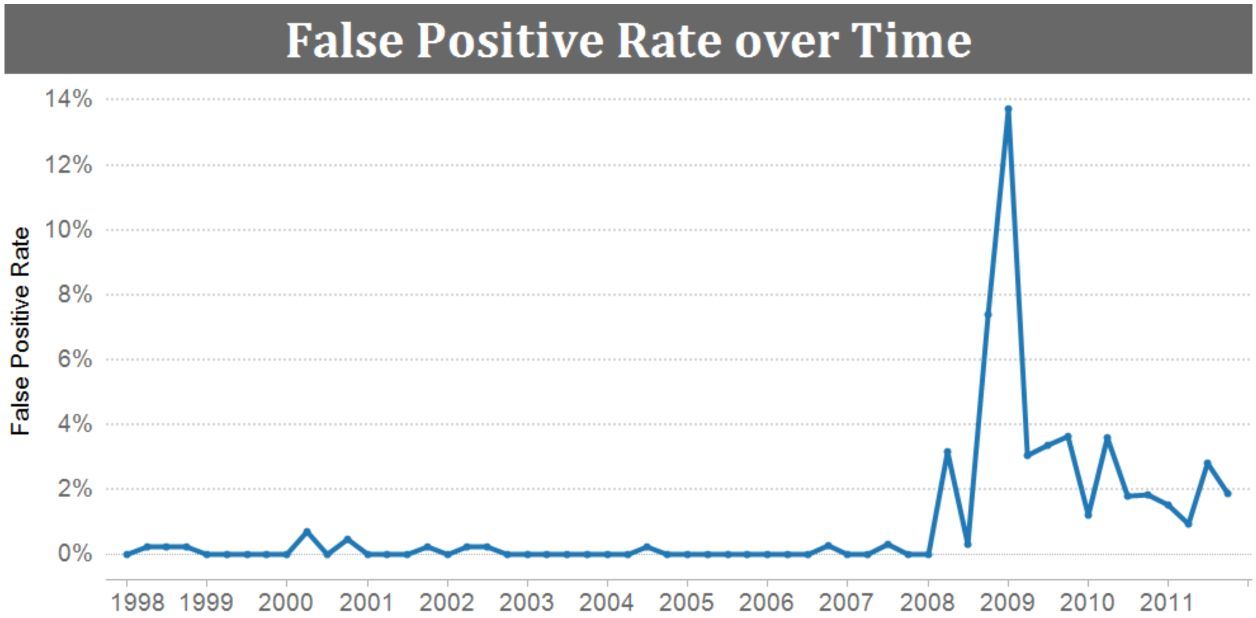

Figure 7 The quarterly false positive rate is the number false signals divided by the total number of non-failing banks active during the quarter.

To investigate the idea that false positives are related to regulatory intervention, we consider whether false signals are correlated with systemic importance. Systemically important financial institutions may have been more likely to receive regulatory support due to the risks they posed to the entire financial system, which could inflate their false positive rate. While statutory designation of systemically important institutions now depends on a range of factors, a rough, simple proxy that can be applied to this historical sample is BHCs with total assets of more than 10 billion dollars.

Comparing the false positive rates for the largest banks alongside all other banks reveals that a disproportionately high number of false positives came from institutions with assets over 10 billion dollars during 2009 Q1 (Figure 8). The false positive rate hit nearly 25 percent for these large banks, while peaking at only 11 percent for the others. This divergence in rates is most prominent during 2009 Q1, suggesting that the timing effect from Figure 7 may be directly related to bank size. Additional analysis of this effect would be a useful avenue for future research.

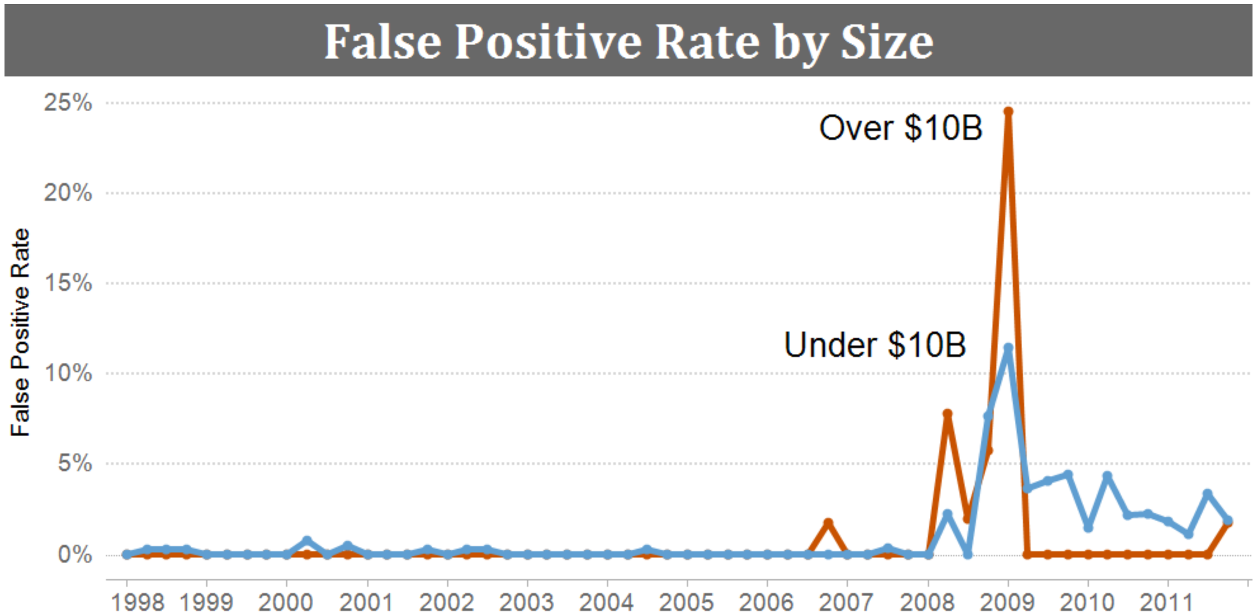

Figure 8 **The false positive rate is higher for bank holding companies over 10 billion dollars during 2009 Q1.**

7 - Practical considerations

Market-based measures face some hurdles for use in supervision (see for example the discussion in Feldman and Levonian, 2001). The main downside to MCR as a predictive risk measure is its availability, since only publicly traded bank holding companies have the liquid equity markets necessary to calculate the ratio. As a result, MCR cannot be the only measure used, since most banks and bank holding companies do not have traded equity. Market-based measures also are likely to be more volatile than measures of financial condition that are based on accounting data. Volatility could be desirable if it is an accurate reflection of variation in bank condition, but is undesirable to the extent it is driven by technical factors or noise in equity market trading activity.

In addition, MCR is constructed using a model based on a particular theoretical framework. As with the result of any model, its value and reliability depends on the degree to which the underlying model is a suitable reflection of reality. Values of MCR are sensitive to choices made in the specification of the model. A more thorough evaluation of the robustness of the results presented here would be an important precondition for use of MCR to make supervisory decisions.

8 - Conclusion

In this paper we develop a market-based indicator of the condition of banks, the "market capital ratio" or MCR. Because the MCR is based on market information, it is likely to incorporate expectations about the future, and may therefore be valuable as an early warning signal that would allow bank regulators and others to identify banks that may fail or otherwise become problems in the future.

We find that the MCR can be used to predict with reasonable accuracy which banks will fail. However, statistical tests indicate that its performance in this regard is not significantly different from more traditional regulatory capital measures such as those incorporated in the Prompt Corrective Action (PCA) framework, or measures such as the Texas ratio. MCR identifies a higher percentage of the banks that ultimately fail than do the PCA or Texas ratio measures; however, it does so in part at the expense of a higher rate of "false alarms," wrongly identifying as problems a larger percentage of banks that ultimately do not fail.

However, there is a significant difference in timing, with the MCR able to signal banking problems earlier than either the PCA capital ratios or the Texas ratio. This timing difference is likely to understate the actual timing advantage of MCR, for two reasons; first, in this paper we have made conservative assumptions with regard to the MCR signal in cases where multiple signals for a single bank are received; second, the share price information needed for MCR can be observed in real time, whereas capital and other financial ratios are reported only with some lag. Thus, while the market capital ratio and traditional financial ratios have similar ability to distinguish problem banks from healthy banks, the market-based signals can identify problems sooner, which may allow more time and scope to reduce the impact of banking problems.

We also consider a refinement to the use of MCR as a signal of problem banks by considering thresholds other than zero. We find that an optimal threshold for the MCR is in the neighborhood of negative 2 percent; that is, the most reliable indicator that a bank is likely to be a future failure rather than a survivor is that the properly adjusted market value of a bank's assets falls below the value of the bank's liabilities by an amount equal to 2 percent of the value of those assets.

Finally, we present preliminary evidence that market-based signals became less informative at the depth of the recent financial crisis, and that the reduction in accuracy at that time was most pronounced for larger banks.

References

Cole, Rebel A., and Lawrence J. White (2011), "Déjà Vu All Over Again: The Causes of U.S. Commercial Bank Failures This Time Around." *Journal of Financial Services Research,* May, vol. 42, pp. 5-29

Feldman, Ron J and Mark Levonian (2001), "Market Data and Bank Supervision: The Transition to Practical Use." *The Region*, Federal Reserve Bank of Minneapolis, September.

Flannery, Mark J. (1998), "Using Market Information in Prudential Bank Supervision: A Review of the U.S. Empirical Evidence," *Journal of Money, Credit, and Banking* 30 (August, pt. 1), pp.273–305.

Gilbert, R. Alton, Andrew P. Meyer, and Mark D. Vaughan (2000), "The Role of a CAMEL Downgrade Model in Bank Surveillance." *Federal Reserve Bank of St. Louis Working Paper Series,* 2000-021A.

Gizycki, Marianne, and Mark Levonian (1993), "A Decade of Australian Banking Risk: Evidence from Share Prices." *Reserve Bank of Australia Research Discussion Paper 9302*.

Gunther, Jeffrey W, Mark E Levonian, and Robert R Moore (2001), "Can the Stock Market Tell Bank Supervisors Anything They Don't Already Know?" *Federal Reserve Bank of Dallas Economic Review*, September, pp.2-9.

Jesswein, Kurt R. (2009), "Bank Failure Models: A Preliminary Examination of the 'Texas' Ratio." *Proceedings of the Academy of Banking Studies,* vol. 9, no. 1, New Orleans.

Levonian, Mark E, and R Keith Friend (2012?), "The Relationship between Market Implied Asset Volatility and Capital Ratios in U.S. Bank Holding Companies" *DRAFT MANUSCRIPT*

Parkinson, M. (1980), "The Extreme Value Method for Estimating the Variance of the Rate of Return'' *Journal of Business*, vol. 53, no. 1, January, pp. 61-65

Thomson, James B (1991), "Predicting Bank Failures in the 1980s." *Federal Reserve Bank of Cleveland Economic Review,* vol. 27, no. 1, pp.9-20.

Appendix

The market capital ratio is constructed by solving for the market-implied value of a bank's assets using a Merton-style option model.[18] The model assumes that a bank's equity capital is analogous to a call option on the company's assets with a strike price equal to the book value of debt. Numerical computation techniques allow for estimation of the price of the underlying "stock" (the bank's assets) based on the market price of the option (the bank's publicly traded equity). Such a model incorporates not only the market share price of the company, but also the stock return volatility. Higher return volatility decreases the market-implied asset value, all else being equal.

Calculating market-implied capital ratios requires collecting data from both a bank's balance sheet and from the stock market. Since top-tier bank holding companies issue stock, balance sheet information must come from the FR Y-9C quarterly call report filed by all domestic top-tier bank holding companies in the United States. Daily stock prices for these holding companies come from the Center for Research in Security Prices (CRSP).[19] The institutions used for the sample are the publicly traded bank holding companies that appear both in CRSP and file quarterly Y-9C reports with the Federal Reserve.

Stock price data for each bank are sampled on a daily basis, while the Y-9C is only available quarterly. The primary market inputs for the option model (market capitalization, dividend rate, and equity volatility) must be aggregated from a daily basis to create quarterly observations that match up with balance sheet data. Since the balance sheet is assumed to reflect the state of the bank as of the end of the quarter, the bank's market capitalization is calculated as the closing stock price multiplied by the total shares outstanding on the quarter's last trading day. The dividend rate is the sum of all dividends paid out during the quarter divided by the book value of assets from the balance sheet. Equity volatility is estimated from the distribution of weekly returns within the quarter using an extreme value estimator that is more efficient than the standard deviation of returns estimator typically used.[20] This estimator uses the highest and lowest closing prices for each week in the quarter to find volatility.

Once the market data has been aggregated to the quarter level, a Merton-style option model can be used to solve for market-implied assets and asset volatility at quarter end. The market capital ratio is defined as the difference between the market-implied assets and book liabilities divided by the market-implied assets.

[18] See Gizycki and Levonian (1993).

[19] The FR Y-9C reporting form collects basic financial data from a domestic BHC on a consolidated basis in the form of a balance sheet, an income statement, and detailed supporting schedules, including a schedule of off balance-sheet items. CRSP is a non-profit research center at the Booth School of Business of the University of Chicago that provides historical stock market data.

[20] See Parkinson (1980).