Executive Stock Options and Bank Risk-Taking*

Chris Armstrong, Allison Nicoletti, and Frank Zhou†

The Wharton School

University of Pennsylvania

May 30, 2019

*We thank seminar participants at Carnegie Mellon University, Lancaster University, and the Wharton School for helpful comments and suggestions. We gratefully acknowledge financial support from the Wharton School and the Jacobs Levy Equity Management Center for Quantitative Financial Research.

†Mailing address: 3620 Locust Walk, Philadelphia, PA, 19103 E-mail address: szho@wharton.upenn.edu

Executive Stock Options and Bank Risk-Taking

Abstract

Section 956 of the Dodd-Frank Act aims to reduce risk-taking by regulating the structure of bank executives' compensation. The efficacy of these regulations presumes that bank executives' compensation contracts *cause* them to undertake risky activities, despite the fact that prior research documents mixed evidence about the nature of this relation. We examine whether bank executives' equity incentives *cause* them to take risk at their banks and, if so, which specific *types* of risk and risky activities. We use a novel identification approach to distinguish between any *causal* effect of bank executives' contracts and any endogenous matching of executives and banks that could also produce an empirical relation. We find that bank executives' equity portfolio $Vega$ has a *causal* effect on their bank's future systemic risk during economic downturns, but not during economic expansions. We also find that $Vega$ leads to greater commercial and industrial lending and investments in non-agency mortgage-backed securities. Collectively, our results suggest that bank executives' contracts *cause* them to take systemic risk, which manifests with a delay during economic downturns, and that riskier lending and investments are two activities that are responsible for this effect.

1 Introduction

Understanding the drivers of risk-taking at banks is critical given that their role as financial intermediaries and important credit providers can cause the results of their risk-taking to spill over into–and have an out-sized effect on–the broader (or "real") economy. A number of influential studies emphasize banks' important role in mitigating credit market frictions and show how they can amplify business cycles and influence aggregate economic activity (e.g., Bernanke and Gertler, 1995; Bernanke et al., 1999). Ivashina and Scharfstein (2010) estimate the economic significance of lending contractions during the most recent financial crisis by documenting a decline in total bank loans issued from approximately $700 billion in June 2007 to $281 billion in June 2008. Given these far-reaching consequences, the extent to which bank executives' compensation contracts encourage risk-taking has become focal in recent academic and policy debate. In response to these concerns, Congress passed the Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank) in 2010. One of the centerpieces of this legislation is Section 956, which requires banking authorities to draft regulations to restrict executive compensation practices that encourage risk-taking. Although this mandate is intended to curb risk-taking that regulators deem inappropriate and excessive, it presumes that banks' risk is *caused by*–rather than simply associated with–their executives' compensation contracts. Moreover, since banks are exposed to numerous different risks and engage in a variety of risky activities, the mandate also assumes that regulators can accurately identify the specific activities and risks influenced by bank executives' contractual incentives.

However, if anything, the evidence is mixed regarding the extent to which bank executives' compensation contracts *cause* risk-taking. Although some studies document evidence consistent with a causal relation (Chen et al., 2006; DeYoung et al., 2013; Larcker et al., 2017), others fail to find such a link (Houston and James, 1995; Fahlenbrach and Stulz, 2011; Chesney et al., 2012; Boyallian and Ruiz-Verdú, 2017). The conflicting evidence may be symptomatic of the endogenous nature of executive compensation contracts and highlights the possibility that the new proposed regulations may not produce their desired effect. Although studies of non-financial firms provide evidence that is somewhat more consistent with a *causal* effect of executives' contractual incentives on risk-taking (Coles et al., 2006; Low, 2009; Armstrong and Vashishtha, 2012), most of these studies deliberately exclude banks and other financial institutions because of the numerous institutional features (e.g., regulatory oversight, large differences in capital structure, different types of agency conflicts, and potential externalities) (John and Qian, 2003) that make them sufficiently distinct from their non-financial counterparts. Therefore, these studies provide little insight into the specific types of risks and and risky activities that bank executives' compensation contracts encourage them to take. More importantly, evidence about the relation between the structure of executive compensation and *systemic* risk–which is an, if not the most, important and pressing concern of bank regulators–is virtually non-existent from studies of non-financial firms.

In light of these important gaps at the intersection of the incentive-compensation, banking, and risk-taking literatures, we study the following two related research questions. First, we examine whether bank executives' incentive-compensation contracts *cause* them to take risk, as implicitly presumed by the Section 956 mandate. Second, we examine the specific *types* of risk–including *systemic*–and specific risky activities that bank executives' compensation contracts encourage them to take. We follow the prior empirical risk-taking literature and focus on bank executives' equity portfolio (i.e., stock and option) holdings, which accounts for the vast majority of their monetary wealth and incentives (Core and Guay, 1999; Core et al., 2003; DeYoung et al., 2013). In particular, we examine *Vega*, which captures changes in the value of executives' equity portfolios to changes in their banks' stock return volatility, since it provides executives with an unambiguous incentive to increase risk (Lambert et al., 1991; Ross, 2004).[1] The clear theoretical prediction for *Vega* facilitates our examination of the specific types of risks and risky activities caused by bank executives' compensation contracts, since risks that are within executives control should exhibit a positive relation with *Vega*. This prediction may interact with important institutional features that are unique to the banking sector, including deposit insurance, which provides a "put-like" payoff that can further encourage risk-taking (Ross, 2004), a high degree of leverage, which amplifies the sensitivity of executives equity holdings to stock returns and volatility (Guay, 1999), and regulatory oversight (Saunders et al., 1990; Buser et al., 1981), which may restrict risk-taking.[2] This provides yet another reason why it is difficult to extrapolate empirical findings from non-financial firms to

[1] *Delta*, the change in the value of the executive's equity portfolio to changes in their bank's stock price has a theoretically ambiguous effect on risk-taking. As such, we focus on *Vega* given the unambiguous theoretical prediction and our interest in documenting a causal effect, if any, of contractual incentives on risk-taking. However, we control for *Delta* in our analyses.

[2] For example, Laffont (1998, p. 249) argues that "a government safety net in the form of deposit insurance or less formal government protection weakens the incentives to monitor banks and exacerbates excessive risk-taking by bank managers."

banks and the importance of studying risk-taking specifically at banks.

Answering our first research question relies on our ability to credibly distinguish between two distinct economic forces. On one hand, since bank executives' compensation depends on the outcome of risky activities that are presumably at least partially under their control, their compensation contracts can *cause* them to take risk. More precisely, we define the *causal* effect as the change in risk-taking induced from varying features of the contract (e.g., Vega) while holding all other factors (e.g., executives' risk-tolerance) fixed. On the other hand, banks are complex and risky and managing them requires a certain set of skills and attributes that are thought to be in short supply (Hubbard and Palia, 1995; Dewatripont et al., 2010). In order to attract or *match with* the right "type" of executives–namely those who have the desired attributes–banks might offer different compensation contracts. For example, banks that are either inherently more risky or that seek to pursue riskier activities may offer relatively "high-powered" (e.g., more option-based) compensation contracts as a way to attract executives who are more tolerant of risk. Cheng et al. (2015) present evidence that is consistent with this scenario. In particular, they find a positive relation between bank managers' total compensation and their banks' risk, which they interpret as evidence that bank managers both demand and receive greater compensation for working at riskier banks. However, in this case, although executives' contractual incentives are *correlated* with their bank's risk, it does not reflect a *causal* relation. Rather, it is an artifact of unobservable differences in executives' risk-tolerance that is correlated with both their compensation contracts and their bank's risk.

To distinguish between any *causal* effect of bank executives' compensation contracts on their risk-taking decisions and endogenous matching, we follow the approach developed by Klein and Vella (2010). We refer to this method as "control function regression" since it builds on the control function approach of Heckman (1976, 1978, 1980) and Heckman and Robb (1985, 2000). The standard approach for identifying the *causal* effect of executives' contractual incentives on their firm's risk is to estimate a two-stage least squares model where the first-stage models executives'–potentially endogenous–contractual incentives and the second-stage models firm risk as a function of executives' predicted incentives from the first-stage. This approach relies on the availability of a valid instrument that is correlated with the first-stage dependent variable (e.g., executives' contractual incentives), but is uncorrelated–or, more technically, mean independent–with the second-stage error. Klein and Vella (2010) recognize the difficulty in finding such an instrument and instead decompose the endogenous component of second-stage (structural) error into two multiplicative components by regressing it on the first-stage error: (i) the correlation between the first- and second-stage errors, and (ii) the ratio of the standard deviations of these two errors. The decomposition illustrates that heterogeneity in the standard deviation ratios provides information about the severity of the endogeneity problem. Because the degree of bias in OLS coefficient estimates depends on the severity of the endogeneity problem, the approach uses differences between the standard deviation ratio across different subsamples to correct for the endogeneity problem and, in turn, identify the causal effect.[3]

The following example illustrates how variation in the standard deviation ratio allows us to identify the causal effect of $Vega$ on bank risk in our setting. Consider banks in two different markets: banks on the east and west coasts. Suppose that east coast banks have a larger standard deviation ratio than do west coast banks.[4] Further suppose that separate OLS regressions of bank risk on $Vega$ in the two markets produce similar coefficient estimates. Since the endogeneity bias–if any–is the product of (i) the correlation between the first- and second-stage errors and (ii) the standard deviation ratio, finding similar OLS coefficients in the two markets implies a small correlation between the first- and second-stage errors. If the correlation were large, because the standard deviation ratios differ across the two markets, a larger correlation between the first- and second-stage errors–i.e., a more severe endogeneity problem–would produce larger differences in the OLS coefficients across the two samples. This example shows how any differences in both OLS coefficients and standard deviation ratios across markets provides information about the correlation between the first- and second-stage errors, which is the empirical manifestation of the endogeneity problem. The innovation is that the decomposition does not require joint-normality as in Heckman (1976), and simply re-expresses the OLS coefficient from regressing the second-stage error on the first-stage error.

[3] Mathematically, the OLS estimate from regressing Y on X, denoted by $\hat{\beta}$, equals $\beta + \rho\frac{\sigma_\eta}{\sigma_\xi}$, where β represents the *causal* effect, σ_η is the standard deviation of the residuals of Y, σ_ξ is the standard deviation of the residuals from a first-stage model of X, and $\rho = corr(X, Y)$. From the formula $\hat{\beta} = \beta + \rho\frac{\sigma_\eta}{\sigma_\xi}$, regressing Y on X as well as the interaction between $\frac{\sigma_\eta}{\sigma_\xi}$ and X identifies β and ρ.

[4] We remain agnostic about why the standard deviation ratios differ because the structural errors are, by definition, unobservable. However, they may differ because of differences in banks' (i) business models, (ii) local labor market supply, (iii) regulatory incentives and oversight, or (iv) local lending-market conditions. For these reasons, and as discussed later, we define markets based on geographic regions specified by regulatory oversight.

Our second research question regarding the specific *types* of risks and risky activities that bank executives' contracts cause them to take stems from banks' role as financial intermediaries that are exposed to a variety of risks and engage in risky activities that distinguish them from their non-financial counterparts.[5] Accordingly, limiting our analysis to traditional measures of risk (e.g., stock return volatility) and risky activities (e.g., R&D and CAPEX) would undoubtedly paint an incomplete–and, more likely, misleading–picture of banks' risk profiles. Moreover, since the essential functions of banks (e.g., maturity transformation) might require them to bear a certain amount of risk and engage in specific risky activities, bank executives might have little discretion to alter these risks, regardless of their contractual incentives. Since equity portfolio *Vega* provides executives with an unambiguous incentive to take risk, these tests allow us to infer the specific types of risk and risky activities that are "controllable" (e.g., Hölmstrom, 1979, 1982). In other words, to the extent there is no detectable *causal* relation between *Vega* and a particular type of risk (e.g., systemic) or a specific risky activity, it suggests that the risk or activity is not under the executives' control and therefore is not influenced by their equity incentives.

We first examine the effect of $Vega$ on *systemic* risk, which is the risk that is central–and largely unique–to the banking industry. Systemic risk is usually defined as any risk that may affect the financial system as a whole (De Bandt and Hartmann, 2000; Freixas and Rochet, 2008, p.235), and can be driven by spillovers from one institution to another as well as common exposures across multiple banks. We focus on two measures of systemic risk: MES, marginal expected shortfall following Acharya et al. (2017), which is more likely to capture common exposures, and $\Delta CoVaR$ following Adrian and Brunnermeier (2016), which captures both spillover effects and common exposures. We find a positive relation between $Vega$ and systemic risk in OLS regressions, but no significant relation between $Vega$ and systemic risk in the control function regressions following the approach from Klein and Vella (2010). Although the control function regressions produce larger standard errors due to their reliance on group-specific estimates of the standard deviation of the residuals, the coefficient estimates on $Vega$ are significantly different from their counterparts in the OLS regressions across all specifications. The estimated coefficients on $Vega$ in the control function

[5] Freixas and Rochet (2008, p. 265) explain that "the management of risks can be seen as *the major activity* of banks" (emphasis supplied). They further explain that "Banks have to control and select the risks inherent in the management of deposits, loan portfolios, securities, and off-balance-sheet contracts. Like any limited liability firm, banks are subject to both liquidity risk and solvency risk, but the consequences of these risks are much more dramatic for banks than for the other sectors of the economy."

regressions are roughly one-half and two-thirds the size of their OLS counterparts in the $\Delta CoVaR$ and MES specifications, respectively. Since systemic risk pertains to inherently rare events (e.g., financial crises), an annual window–which is typical for measuring risk–may not be appropriate if systemic risk manifests over a relatively low frequency. We therefore examine whether $Vega$ causes systematic risk over a longer-horizon (Acharya and Naqvi, 2012). We continue to find no significant relation between $Vega$ and systemic risk for up to three years after the measurement of $Vega$ based on our control function regressions. Overall, these findings suggest that endogenous matching is responsible for any positive association between $Vega$ and systemic risk observed in the data.

Given the lack of a statistically detectable relation between $Vega$ and systemic risk in our full sample, we next examine whether there is a differential relation during economic expansions and downturns. Acharya and Naqvi (2012) show that risk-taking during expansions can lead to excessive lending, which suggests that the effects of risk-taking incentives on systemic risk might not manifest immediately and only become detectable during economic downturns. Consistent with this conjecture, we find that a one standard deviation increase in $Vega$ increases one-, two-, and three-year-ahead systemic risk (MES) by 0.42, 0.65, and 0.60 percentage points, respectively, if the economy experiences a downturn during the subsequent one, two or three years. In contrast, the increase is 0.10, 0.14, and 0.09 percentage points, respectively, if the economy does not experience a downturn in one, two or three years. Collectively, these findings suggest that bank executives' equity incentives *cause* them to take systemic risk during economic expansions but–perhaps because of the infrequent nature of negative outcomes from systemically risky activities–that this relation is only detectable with a lag and manifests during economic downturns.

To provide evidence on the channels through which $Vega$ gives rise to systemic risk, we examine two primary banking activities that may generate this risk. Specifically, we examine commercial and industrial (C&I) loans and non-agency mortgage backed securities, which are both considered to be relatively riskier activities given the higher likelihood of losses relative to other types of loans and investments (DeYoung et al., 2013). We find that $Vega$ leads to a larger proportion of commercial and industrial loans ($CommLoans$) in banks' lending portfolios as well as a larger proportion of non-agency mortgage backed securities ($MBSNA$) in their investment portfolios. This suggests that $Vega$ affects systemic risk by encouraging managers to increase their banks' exposure to C&I lending and non-agency mortgage backed securities.

Our paper makes multiple contributions to several distinct literatures. First, prior executive compensation studies provide mixed evidence about the effect of bank executives' contractual incentives on their risk-taking decisions. We contribute to this literature by introducing an econometric technique that, coupled with key institutional features of the banking industry, allows us to isolate (i.e., identify) the extent to which bank executives' contractual incentives *cause* them to take risk as opposed to reflecting the *endogenous matching* of executives and banks. By mapping *a priori* industry-specific knowledge about the banking industry and bank executive labor market into the assumptions needed to apply the method outlined by Klein and Vella (2010), our paper demonstrates how this technique can be used in other settings to identify the causal effects of endogenously designed contracts. In addition, regardless of whether they are explicitly stated or implicitly invoked, any method for causal inference relies on inherently untestable identifying assumptions (Cartwright, 1979; Holland, 1986; Pearl, 2001). Our methodology is no exception. We therefore assess the sensitivity of our inferences to our maintained identifying assumption that we have correctly identified distinct banking markets to alternative definitions of markets. In doing so, we contribute to the emerging branch of the causality literature that advocates the use of sensitivity analysis and partial identification techniques (Manski, 2010; Rosenbaum, 2010; Tamer, 2010; Armstrong, 2013).

Second, we examine industry-specific measures of bank risk over different horizons. Both innovations are important in the own right, and together provide new insight into the nature of banks' risk-taking activities and the horizons over which these risky activities manifest in bank risk. Prior risk-taking studies have largely neglected *systemic* risk, presumably because most large-sample studies deliberately exclude banks and other financial institutions. This is an extremely important gap in the literature because banks' exposure to and contribution to the creation of systemic risk is a key distinguishing feature that is intimately related to their role in credit provision. Moreover, analysis of systemic risk is complicated by the fact that it relates to inherently infrequent events. Accordingly, standard research designs from the risk-taking literature are unlikely to be adequate for powerful tests to detect this specific type of risk. Showing that a relation between *Vega* and systemic risk manifests only during economic downturns and does so with a lag highlights the importance of constructing empirical tests that correspond to the appropriate outcome horizon of the examined risk-taking activities.

Third, in virtue of our first two contributions–namely assessing the nature of causality and examining multiple different types of risk and risky activities–our evidence speaks to the potential efficacy of regulations aimed at curbing risk-taking at banks (e.g., the Section 956 mandate). Our findings should therefore be of interest to bank regulators and other stakeholders in the ongoing debate surrounding executive compensation practices in these important institutions. Since a causal effect of vega on systemic risk only manifests during economic downturns–and with a lag–regulators' ability to constrain risk-taking may be more limited than is presumed by the Section 956 mandate and may also require more foresight than previously assumed.

2 Background and Related Literature

2.1 Risk-taking in the Banking Industry

Although executives at both banks and non-financial firms engage in risky activities, the central role of banks in the financial system suggests that the consequences of their risk-taking can have additional–and potentially out-sized–effects on the broader economy. For example, Bernanke and Gertler (1995) and Bernanke et al. (1999) discuss how lending restrictions during economic downturns can amplify macroeconomic contractions and exacerbate financial crises. A particular concern is banks' exposure and contribution to systemic risk, which, at a broad level, refers to the potential for risk to propagate through the financial system. As discussed by Adrian and Brunnermeier (2016), institutions can be systemically risky by themselves–especially those that are large and interconnected–or collectively, which typically results from banks engaging in similar business activities that have similar exposures to financial risks.

Following the most recent financial crisis, the structure of bank executives' compensation contracts has received increased attention and scrutiny as a potential source of their incentives to take risk. Symptomatic of these concerns, Section 956 of the Dodd-Frank Act, bank regulators have been charged with writing rules to restrict compensation contracts that encourage "inappropriate risk-taking." The proposed rule discusses concerns related to systemic risk, stating that "Larger financial institutions in particular are interconnected with one another and with many other companies and markets, which can mean that any negative impact from inappropriate risk-taking can have broader consequences." The proposed rule also cites business decisions related to lending and investments, which are typically the two largest classes of assets on banks' balance sheets. Moreover, these two decisions are particularly susceptible to systemic risk concerns because banks can take actions that create a build up of latent risk in their lending and investment portfolios that does not materialize until an economic downturn.

2.2 Related Literature

Our study lies at the intersection of–and contributes to–two streams of research. First, our study adds to the literature that examines whether executives' contractual incentives *cause* risk-taking in general, and at banks in particular. Studies in this literature tend to focus on broad, market-based measures of risk that are intended to capture the collective result of the numerous risky activities that are presumably in part, if not entirely, under executives' control. These measures include total, systematic, and idiosyncratic stock return volatility and the literature provides mixed evidence regarding the existence of a causal effect.

Houston and James (1995) find a *negative* relation between bank executives' equity incentives–measured as the percentage of options held and the percentage of ownership–and total risk, which they interpret as evidence that bank compensation policies do not encourage risk-taking. Similarly, Cheng et al. (2015) also fail to find evidence of a *causal* effect and instead conclude that the positive relation between bank executives' total direct compensation and their bank's total risk is the result of a risk-premium that is compensation for working at riskier banks. However, in contrast to these studies, Chen et al. (2006) document a positive relation between the value of bank managers' options and multiple risk measures of bank risk, including total, systematic, and idiosyncratic stock return volatility and interest rate risk. DeYoung et al. (2013) also find a positive association between bank managers' vega and total, systematic, and idiosyncratic risk. Both of these studies conclude that bank managers' contractual incentives encourage risk-taking. Finally, Low (2009) examines changes in the Delaware takeover laws and finds a positive association between vega and total risk in a sample of non-financial firms. Armstrong and Vashishtha (2012) also examine non-financial firms and find a positive relation between vega and systematic risk, but no relations between vega and idiosyncratic risk.

The second stream of literature to which our study contributes examines the influence of contractual incentives–primarily vega–on risky *bank-specific* activities surrounding the most recent financial crisis in 2007 - 2009. Studies in this literature provide somewhat different inferences from those that focus on longer time periods and are not crisis-specific. For example, Chesney et al. (2012) find no relation between *pre-crisis* vega and loan write-downs *during* the crisis. However, they do find a positive association between the sensitivity of executives' wealth to asset return volatility and write-downs. Boyallian and Ruiz-Verdú (2017) find no evidence of a relation between pre-crisis vega and the incidence of bank failure during the crisis, but find a positive relation between pre-crisis delta and failure. Similar with these conclusions, Fahlenbrach and Stulz (2011) find no consistent evidence that poorly aligned pre-crisis contractual incentives–including vega, delta, cash bonuses, and stock ownership–had an effect on bank performance during the crisis. Although these studies largely conclude that bank executives' pre-crisis equity incentives did not influence bank outcomes during the crisis, several studies present evidence to the contrary. Specifically, Gande and Kalpathy (2017) find that pre-crisis vega is positively associated with the receipt of government assistance during the crisis. Larcker et al. (2017) find evidence of a positive relation between vega and systematic risk at banks with retained interest in securitization during the pre-crisis period. Moreover, they document that pre-crisis vega is positively associated with total and systematic risk during the crisis and that aggregate vega in the bank industry as a whole is associated with future declines in economy-wide indicators, including GDP growth.

Our primary contribution to these two literatures is to provide evidence about the extent to which any relation between bank executives' contractual incentives and bank risk represents a *causal* effect. Although the discrepant conclusions in the aforementioned studies may be the result of differences in research designs, sample periods, or measures of incentives and risk, another possibility is the endogenous matching of executives and banks. Another important contribution of our study is our focus on systemic risk, which is the chief concern of bank regulators and other bank stakeholders given the potentially severe economic consequences of bank distress. Moreover, our research design allows for the possibility that contractual incentives induce executives to take systemically risky actions that manifest over a relatively long horizons (e.g., multiple years) by examining multiple different time periods and several distinct activities that are central to banks' business model.

3 Research design

3.1 Variable measurements

We are primarily interested in systemic risk, which we measure two different ways. First, marginal expected shortfall (MES) captures a bank's expected equity loss when the market experiences losses in the extreme left-tail of the distribution (Acharya et al., 2017). Specifically, MES is calculated as the bank's average return during the market's 5% worst days during year t, multiplied by -1 so that larger values of MES correspond to greater systemic risk.

Second, $\Delta CoVaR$ captures dependence in the tail of the loss distribution between the banking system and a particular bank. As discussed by Adrian and Brunnermeier (2016), $\Delta CoVaR$ captures systemic risk created by common exposures as well as spillover effects, and is based on the conditional value at risk ($CoVaR$) of the banking system. We focus on the estimation of $\Delta CoVaR$ which involves the contribution of each bank i to systemic risk of the banking system (sys) (i.e., examining systemic risk of the banking system conditional on bank i being in a certain state) following Adrian and Brunnermeier (2016).[6] Specifically, we first estimate the following bank-level quantile regressions using bank i's full time series of weekly return data over the sample period, requiring at least 260 observations:

$$Ret_{i,t} = \alpha_i^q + \gamma_i^q M_{t-1} + \epsilon_{i,t}^q \tag{1}$$

$$Ret_{sys|i,t} = \alpha_{sys|i}^q + \gamma_{sys|i}^q M_{t-1} + \beta_{sys|i}^q Ret_{i,t} + \epsilon_{sys|i,t}^q \tag{2}$$

$Ret_{i,t}$ represents bank i's weekly return, $Ret_{sys,t}$ represents the value weighted weekly return of the commercial banking sector (three-digit SIC codes 602 and 603), and M is a vector of macroeconomic variables, which includes (1) change in the three-month treasury bill rate, (2) change in the yield curve slope – measured as the difference between the composite long-term bond yield and three-month treasury bill rate, (3) short-term TED spread – measured as the difference between the three-month LIBOR rate and the three-month secondary market treasury bill rate, (4) change in the credit spread – difference between Moody's Baa-rated bond yield and the ten-year treasury rate,

[6] In untabulated analysis, we also examine the $\Delta CoVaR$ measure that reverses the conditioning (i.e., involves systemic risk of each bank i conditioning on the state of the banking system (sys)), referred to as $Exposure\Delta CoVaR$ in Adrian and Brunnermeier (2016), and find similar inferences to those from the main analyses.

(5) rolling 22-day standard deviation of the value-weighted market return, (6) the value-weighted market return, and (7) value-weighted real estate sector (two-digit SIC codes 65 and 66) return.

Using the predicted values from Equation (1) and Equation (2), we construct both VaR and $CoVaR$ as follows:

$$VaR_{i,t}^{q} = \hat{\alpha}_i^q + \hat{\gamma}_i^q M_{t-1} \tag{3}$$

$$CoVaR_{i,t}^{q} = VaR_{sys|Ret_{i,t}=VaR_{i,t}^q}^{q} = \hat{\alpha}_{sys|i}^q + \hat{\gamma}_{sys|i}^q M_{t-1} + \hat{\beta}_{sys|i}^q Ret_{i,t} \tag{4}$$

The final step is to calculate the difference in $CoVaR$ when bank i is in a distressed state (the 1% worst weeks, $Ret_{i,t} = VaR_{i,t}^{q=1\%}$) compared to a typical state (median, $Ret_{i,t} = VaR_{i,t}^{q=50\%}$).

$$\Delta CoVaR_{i,t}^{1\%} = \hat{\beta}_{sys|i}^q (VaR_{i,t}^{q=1\%} - VaR_{i,t}^{q=50\%}) \tag{5}$$

$\Delta CoVaR_{i,t}^{1\%}$ is constructed at the weekly level. Thus, prior to performing our analyses, we sum this measure to the annual level and multiply the measure by -1 so that larger values correspond to greater systemic risk.

We focus on risk-taking incentives from bank executives' equity (i.e., stock and option) portfolios. Our primary independent variable of interest is $Vega$ because of its theoretically unambiguous risk-taking incentives. $Vega_{i,t-1}$ and $Delta_{i,t-1}$ are the portfolio Vega and Delta of the five highest-paid executives of bank i in year $t-1$. Portfolio vega is the change in the risk-neutral (i.e., Black-Scholes) value of the executives' option portfolios for a 0.01 change in the standard deviation of the underlying stock returns. Similarly, portfolio delta is the change in the risk-neutral value of the executives' stock and option portfolios for a one percent change in the value of the underlying stock price. [7]

[7] The parameters of the Black-Scholes formula are calculated as follows. Annualized volatility is calculated using continuously compounded monthly returns over the previous 60 months, with a minimum of 12 months of returns, and winsorized at the 5th and 95th percentiles. If the stock has traded for less than one year, we use the imputed average volatility of the firms in the Standard and Poors (S&P) 1500. The risk-free rate is calculated using the interpolated interest rate on a Treasury note with the same maturity(to the closest month) as the remaining life of the option, multiplied by 0.70 to account for the prevalence of early exercise. Dividend yield is calculated as the dividends paid during the previous 12 months scaled by the stock price at the beginning of the month. This is essentially the same as the method outlined by Core and Guay (2002).

3.2 OLS model

To examine the relation between bank executives' incentives and their firms' risk, we estimate the following specification:

$$
\begin{aligned}
Risk_{i,t+s} = {} & \delta_t + \beta_1 Vega_{i,t-1} + \beta_2 Delta_{i,t-1} + \beta_3 Size_{i,t-1} + \beta_4 BTM_{i,t-1} + \beta_5 Capital_{i,t-1} \\
& + \beta_6 Growth_{i,t-1} + \epsilon \qquad (6)
\end{aligned}
$$

$Risk_{i,t+s}$ represents each of our risk measures for bank i in year $t+s$ where $s = 0$, 1 or 2 to capture risk one-, two-, or three-years ahead of the measurement of bank executives' incentives. Consistent with prior literature, we also include several control variables that are likely to be correlated with both equity incentives and risk-taking. Specifically, larger banks are likely to have different risk profiles and risk-taking incentives, so we control for bank size using the natural logarithm of total assets ($Size$). To capture differences in investment opportunities, we include the book-to-market ratio (BTM) and annual asset growth rate ($Growth$). We also include the ratio of equity capital to total assets ($Capital$) to control for differences in capital structure. Finally, we include year fixed effects to control for secular changes in risk (e.g., across the business cycle).

Model (6) assumes that, given the control variables and year fixed effects, $Vega$ is exogenous with respect to unobserved factors that drive bank risk. However, as noted in the contracting literature in general, and the banking literature in particular (Cheng et al., 2015), unobserved bank and executive characteristics that influence the bank-executive match can affect both executive compensation and risk-taking. For example, a bank that has a relatively risky loan portfolio might offer compensation contracts with greater risk-taking incentives to attract certain types of executives. Alternatively, executives who are prone to risk-taking might match with banks that offer more risk-taking incentives. In these cases, $Vega$ has no *causal* effect, meaning that holding the executive-bank match fixed, there would be no change in risk-taking in response to a change in vega.

The selection problem is two-sided in the sense that both unobserved executive and bank characteristics can affect the outcome. Dealing with such endogeneity can be nontrivial because a valid instrument has to be exogenous to both unobserved bank and executive characteristics. For example, an exogenous shock to bank executives' risk-taking incentives alone may not be a valid instrument if executives can select into banks that are less exposed to the shock or find ways to neutralize the shocks. Specifically, banks can hire executives who are more tolerant of risk–and are therefore more prone to risk-taking–which could either mute or completely undo the effect of regulations that constrain bank executives' contractual incentives. Moreover, as shown by Cheng et al. (2015), both the executive-bank match and risk-taking tend to be persistent. Consequently, identification strategies that rely on time-series (i.e., within-firm and within-executive) variation are unlikely to yield powerful tests. To deal with the two-sided selection problem, we adopt the novel identification strategy developed by Klein and Vella (2010), which we now discuss in more detail.

3.3 Control function regression

The control function regression starts with a typical two-stage regression model. Recall that $Vega_{it}$ represents the $Vega$ of an executives' equity portfolio, and $Risk_{it+s}$ represents the risk of bank i in year $t+s$. $Vega_{it}$ and $Risk_{it+s}$ are specified by the following first- and second-stage models:

$$Vega_{it} = \alpha X_{it} + \xi_{it} \tag{7}$$

$$Risk_{it+s} = \beta_1 Vega_{it} + \beta X_{it} + \eta_{it+s} \tag{8}$$

where β_1 captures the effect of $Vega$ on banks' risk, X_{it} represents the same vector of observable characteristics as in (6) including intercepts, and ξ_{it} and η_{it+s} are unobserved factors that affect $Vega_{it}$ and Risk_{it+s}, respectively. If η_{it+s} is correlated with ξ_{it}, an OLS regression of $Risk$ on $Vega$ will be biased. Let the correlation coefficient between η_{it+s} and ξ_{it} be ρ.

To solve the endogeneity problem, we follow Klein and Vella (2010) and decompose the error term η_{it+s} as follows:

$$\eta_{it+s} = \rho \frac{\sigma_\eta}{\sigma_\xi} \xi_{it} + \omega_{it+s} \tag{9}$$

where $\rho = \dfrac{cov(\eta_{it+s}, \xi_{it})}{var(\xi_{it})}$. The decomposition is achieved by a simple regression of η_{it+s} on ξ_{it} and does *not* assume that η_{it+s} and ξ_{it} are jointly normal. Klein and Vella (2010) show that ρ is identified under the following assumption:

1. The standard deviation ratio of the residuals of equations (7) and (8), $\frac{\sigma_\eta}{\sigma_\xi}$, varies across observations.

To understand the intuition behind this identification strategy, substituting (9) into the risk-taking equation (8) gives

$$Risk_{it+s} = \beta_1 Vega_{it} + \beta X_{it} + \eta_{it+s} \tag{10}$$

$$= \beta_1 Vega_{it} + \beta X_{it} + \rho \frac{\sigma_\eta}{\sigma_\xi}(Vega_{it} - \alpha X_{it}) + \omega_{it} \tag{11}$$

Equation (11) shows that when $\frac{\sigma_\eta}{\sigma_\xi}$ is a constant, γ_1 is not identified because the term $Vega_{it} - \alpha X_{it}$ is collinear with the regressors $Vega_{it}$ and X_{it}. However, when $\frac{\sigma_\eta}{\sigma_\xi}$ varies across observations and its interaction with ξ_{it} is not collinear with $Vega_{it}$, the variation in $\frac{\sigma_\eta}{\sigma_\xi}(Vega_{it} - \alpha X_{it})$ identifies both ρ and γ_1.

The main advantage of this approach is that identification does not rely on the usual functional form assumption of joint normality when researchers do not specify an instrumental variable for the endogeneous regressor. The endogenous component of the structural error is, in turn, decomposed into the following two components: the correlation coefficient ρ and the standard deviation ratio $\frac{\sigma_\eta}{\sigma_\xi}$. The decomposition simply regresses the second-stage error on the first-stage error and does not assume joint normality. Another advantage is that variation in the standard deviation ratio is due to heteroskdasticity, which is a testable feature of linear regression models.

3.3.1 Modeling σ_η and σ_ξ

As illustrated in the previous section, identification of the causal effect relies on variation in the standard deviation ratio, $\frac{\sigma_\eta}{\sigma_\xi}$. To avoid imposing assumptions on the determinants of the standard deviation ratio and to maintain a parsimonious structure, we assume that the standard deviaition ratio varies across "markets" as follows:

$$\frac{\sigma_\eta}{\sigma_\xi} \in \left\{ \frac{\sigma_{\eta,m}}{\sigma_{\xi,m}} \mid m \in \{1, ..., M\} \right\},$$

where m denotes a market.

The definition of markets should be justified on *a priori* theoretical grounds and requires that the unobserved factors that affect both bank executives' risk-taking incentives and their bank's risk exhibit different variances across markets. Although we are agnostic about the specific factors that affect the standard deviation ratios, we argue that differences in the local labor market supply and demand, regional economic conditions, and regulatory incentives are all likely to generate variation in the standard deviation ratio. Since most of these factors stem from differences in geographic and regulatory exposure, we assign banks to four geographical regions based on the Office of the Comptroller of the Currency (OCC) supervision structure. We then consider each region-year combination to be a distinct market. In other words, we assume that $\frac{\sigma_\eta}{\sigma_\xi}$ varies across geographical markets and time. We provide evidence on the nature of variation in standard deviation ratios across markets in section 4.

3.3.2 Discussion

The control function method differs from an IV regression in that it does not explicitly isolate the exogenous variation that induces changes in risk-taking only through $Vega$. Identification instead relies on differences in the standard deviation ratios across subsamples (e.g., markets), because the standard deviation ratio provides information about the severity of the endogeneity problem. For a given correlation coefficient ρ, a smaller standard deviation ratio implies a larger amount of exogenous variation in $Vega$ that is unrelated to $Risk$, which, in turn, means that endogeneity is less of a concern (and vice versa). In other words, the standard deviation ratio is inversely related to the amount of exogenous variation in the first-stage error relative to the second-stage error. Variation in the severity of the endogeneity problem coupled with differences in OLS regression coefficients across different subsamples (e.g., markets) identifies the correlation between the first- and second-stage errors. Once this correlation is known, it can be used to "correct for" any endogenous relation and obtain consistent estimates of the *causal* effect of the endogenous variable on the dependent variable (e.g., *Vega* and bank risk, respectively). For example, finding similar OLS coefficients across markets that have different standard deviation ratios implies that the correlation coefficient is not large. If the correlation coefficient were large, OLS coefficients across markets would have been very different.

For the identification argument to hold, the control function method estimates a constant causal effect and a single correlation coefficient between the two structural errors for all markets. Assuming a single common correlation coefficient is not innocuous and imposes structure on the nature of the endogeneity problem. Klein and Vella (2010) illustrate a few examples. For example, $\xi = \omega_\xi \xi^*$ and $\eta = \omega_\eta \eta^*$. This particular structure implies that a larger variance of the first-stage error relative to that of the second-stage error indicates a smaller endogeneity problem because the first-stage error is more likely to contain exogenous variation that is unrelated to risk-taking. Unlike IV that solves the endogeneity issue using explicit exogenous variation, the control function approach by Klein and Vella (2010) relies on implicit exogenous variation embedded in the heteroskadasticity of the error terms.

Although we argue that identification using control function regressions has a number of advantages given the nature of the endogeneity problem in our setting, there are several caveats to note. First, identifying ρ requires variation in $\frac{\sigma_\eta}{\sigma_\xi}$, as shown in (10). Although we argue that this is the case in our setting, it may not be so in other settings, which would preclude the use of this technique. Second, although the control function approach does not rely on joint normality of the errors, at a minimum, researchers need to specify how the standard deviation ratio varies. Relatedly, we assume that ρ does not vary across markets, although this assumption is not necessary for identification. If there are reasons to expect ρ to systematically vary with certain observables, the researcher would need to specify *a priori* how ρ varies, which is also required by alternative econometric methods, including OLS. Specifically, this idea is similar to allowing for heterogeneous treatment effects in that researchers need to specify how β_1 varies across either individual or groups of observations. Third, the control function approach requires estimating the variance of ξ and η in each market. More observations in a market allow for more accurate variance estimates, but leave fewer variances to identify ρ. Consequently, standard errors can be large, which biases against finding significant coefficients. As such, we perform statistical tests across the OLS and control function regressions to assess the power of our tests.

4 Sample selection and descriptive statistics

4.1 Sample selection

Our sample is comprised of observations with required data at the intersection of Compustat, CRSP, Execucomp, and FR Y-9C regulatory reports during the period 1994 - 2016. We use

Execucomp data to construct our contractual incentive measures, $Vega$ and $Delta$, and Compustat and CRSP data to measure the control variables ($Size$, BTM, $Capital$, and $Growth$). For the risk measures, we obtain daily return data from CRSP and macroeconomic variables from the Federal Reserve Economic Data (FRED) database, the U.S. Treasury Department, and CRSP.[8] In additional analysis, we require lending and investment specific variables, which are obtained from the FR Y-9C regulatory reports. Given that we use one-, two-, and three-year ahead measures of risk, our final sample period for analysis is 1994 - 2013 and includes 1,339 bank-year observations.

4.2 Descriptive statistics

We present descriptive statistics for the pooled sample in Table 1. The first panel of the table presents the distribution of the systemic risk measures and two specific business activities, commercial and industrial (C&I) lending ($CommLoans$) and investments in non-agency mortgaged-backed securities ($MBSNA$), through which $Vega$ potentially affects systemic risk. These variables are presented as percentages, with the exception of $\Delta CoVaR^{1\%}$. The table reveals that the average loss on the worst 5% days for the banking system during the year is 2.9% (MES), while the mean of $\Delta CoVaR^{1\%}$ is 1.7. Moreover, approximately 21% of the loan portfolio is comprised of $CommLoans$ while approximately 6% of the investment portfolio is comprised of $MBSNA$, on average.

[Insert Table 1]

The second panel provides descriptive statistics for the compensation variables. For interpretation purposes, we present descriptives for the compensation variables prior to taking the log. The descriptive statistics for $Vega$ indicate that a 0.01 increase in stock return volatility results in an approximately $282,000 increase in the average risk-neutral value of bank executives' option portfolio. Similarly, a 1% increase in stock price increases the option portfolio value by $939,000 on average ($Delta$). The final panel presents descriptive statistics for several bank characteristics. The average book-to-market ratio is 66.8%, capital ratio is 9.2% and asset growth rate is 12.3%. Untabulated analyses indicate that the mean (median) of total assets is $74 billion ($12 billion)

[8] Market volatility, the market return, and the real estate sector return are constructed using CRSP data. The three-month treasury bill rates, three-month LIBOR rate, 10-year treasury rate, and Moody's *Baa*-rated bond yield are taken from FRED. The long-term composite bond yield is taken from FRED prior to 2000 and the U.S. Treasury department thereafter.

and as such, we take the log of total assets prior to conducting the analyses in order to measure size.

4.3 Portfolio sorts of $Vega$ and systemic risk

Figure 1a illustrates the relation between $Vega$ and systemic risk. Each year, we rank banks into deciles by $Vega$ and compute the average MES for each $Vega$ decile. To account for a potential lead-lag relation, we lead the measurement of MES by that of $Vega$ by two years. The figure shows a positive relation between $Vega$ and MES. Unreported figures indicate that the same pattern exists when using $\Delta CoVaR$ to measure systemic risk and when the measure of systemic risk leads $Vega$ by either one or three years. The positive relation between $Vega$ and systemic risk is consistent with regulators' concern that $Vega$ may contribute to systemic risk. However, it is possible that unobserved executive-bank matching may be responsible for this pattern rather than a causal effect.

[Insert Figure 1]

Because systemic risk manifests primarily during economic downturns, Figure 1b presents the relation between $Vega$ and systemic risk separately during economic contractions and expansions to investigate whether risk-taking incentives further exacerbate systemic risk during economic downturns. We identify economic downturns as the following years: 2001, 2008, and 2009. The dashed blue line presents the relation between $Vega$ and MES when MES is measured during economic expansions, while the solid red line presents the same relation when MES is measured during economic downturns. The figure shows a robust positive relation between $Vega$ and MES during *both* expansions and contractions. However, the magnitude of the relation is much larger during economic downturns: moving from the lowest to the highest decile of $Vega$, systemic risk during economic expansions increases from 2.2 to 2.6. This increase is much smaller, both in absolute and relative terms, than the increase in systemic risk from 5.7 to 7.6 during economic downturns. In other words, although systemic risk is, on average, higher during economic downturns than during expansions, risk-taking incentives seem to further amplify systemic risk during economic downturns. Also note that, by construction, risk-taking incentives are *not* measured during economic downturns because $Vega$ is measured two years prior to the measurement of systemic risk. These results are consistent with the predictions from Acharya and Naqvi (2012) that risk-taking activities during economic expansions "sowed the seeds" of systemic risk that manifest during economic downturns.

4.4 Variation in the standard deviation ratio

As discussed in section 3.3, the control function regression includes the ratio of the standard deviations of the first- and second-stage residuals interacted with the first-stage residuals as an additional regressor. The OLS regression coefficient on the interaction term equals the correlation coefficient between the two residuals. In theory, the correlation coefficient is identified as long as variation in the standard deviation ratio is not zero. In practice, as with any regression, sufficient variation in a regressor is crucial for a high-powered test.

To gauge the extent to which there is heterogeneity across markets, we examine the standard deviation ratios of each risk measure by region in Panel A and by year in Panel B of Table 2. Because the standard deviation ratio of the residuals ultimately depends on the standard deviations of the variables of interest (i.e., $Vega$ and the risk measures), we report the standard deviation ratios of each risk measure and $Vega$ before estimating the control function regression. These descriptive statistics help us gauge the heterogeneity across markets. If the standard deviation ratio of a risk measure and $Vega$ exhibit variation across markets, the likelihood that the residuals also exhibit variation will increase, which can be used as a preliminary diagnosis test.

[Insert Table 2]

Table 2 shows that the ratios of the standard deviations of each risk measure and $Vega$ exhibit significant cross-sectional and time series variation. Panel A presents the standard deviation ratios for each of the OCC regions, where the ratio is defined as the standard deviation of each of the risk measures to the standard deviation of $Vega$. Panel A shows that Region 1 has the largest standard deviation ratio and Region 2 has the smallest. The pattern is consistent across all risk measures, which suggests that the risk measures capture related constructs.

Panel B presents the standard deviation ratios for each year. The economic magnitude magnitude of the standard deviation ratio differs from that in Panel A. For example, the standard deviation ratio for MES is around 1.1 to 1.4 across regions in Panel A but varies from 0.18 in 2012 to 1.33 in 2008 in Panel B. The result suggests that, relative to $Vega$, MES exhibit large time-series variations and is consistent with MES capturing systemic risk, which manifests primarily during economic downturns. $\Delta CoVaR^{1\%}$ also exhibits a similar pattern.

5 Results

5.1 Vega and systemic risk

5.1.1 Main results

We present estimates of the effect of $Vega$ on systemic risk, measured using either MES or $\Delta CoVaR^{1\%}$, in Table 3. We use one-, two-, and three-year-ahead measures of systemic risk to allow for the possibility of a lag between when executives make risky decisions and when the resulting risk manifests.[9] Columns (1), (3), and (5) of Panel A provide estimates from the OLS model using MES as the dependent variable for one-, two-, and three-years ahead, respectively. In each of the three columns, the coefficient on $Vega$ is positive and significantly different from zero.

[Insert Table 3]

To separately identify the causal effect of bank executives' incentives and the endogenous matching of executives and banks, we also present results from the control function regressions. Columns (2), (4), and (6) of Panel A present results using the control function regression and show that the coefficient on $Vega$ is no longer statistically significant for any of the windows. Moreover, the coefficient magnitudes are substantially smaller than those of their OLS counterparts, and all of the differences between the corresponding coefficients across the two methods are statistically significant.

Panel B of Table 3 provides estimates from the same specification using $\Delta CoVaR^{1\%}$ as the dependent variable. Similar to MES, we find positive and significant coefficients on $Vega$ in each of the three time periods examined. However, using the control function regressions, the estimate on $Vega$ becomes insignificantly different from zero for the two- and three-year ahead periods.

[9] Prior research primarily examines one-year ahead risk-taking measures. Although this is sensible for typical measures of risk in non-financial firms (e.g., stock return volatility, R&D expenditures, leverage), it may not be appropriate for systemic risk, which is a "lower-frequency" variable and negative realizations might only be empirically detectable over a sufficiently wide window. This is analogous to the "peso problem" in asset pricing literature.

Moreover, the magnitude of the estimates in the control function regression is smaller than its counterpart in the OLS regression and, as with MES, the difference is statistically significant. Collectively, the estimates from the control function regressions provide no evidence of a causal relation between systemic risk, measured with MES and $\Delta CoVaR$, and $Vega$. Instead, the evidence suggests that the positive correlation between $Vega$ and systemic risk up to three years ahead largely reflects the endogenous matching of executives and banks rather a causal effect.[10]

Because the standard deviation ratios are estimated for each market with an average of 20 observations, one concern may be that standard errors are large and the power of the test is low. However, we find that the OLS and control function regression coefficients are significantly different in the majority of the specifications, which suggests that lack of power is not likely to be responsible for the insignificant coefficient estimates. Nevertheless, to gauge the power of the control function regression, we plot the distribution of the coefficient estimates for OLS versus control function regressions. The idea is that if a control function regression introduces noise, which is entirely plausible, the coefficient estimates will appear random and cover a much larger area than those from an OLS regression.

[Insert Figure 2]

Figure 2 shows that the standard errors for the control function regression are larger than those from the OLS regression. However, it is also clear from the figure that the OLS coefficient significantly differs from the coefficient from the control function regression.[11] The results alleviate the concern that low power is the main reason why the coefficient is insignificant.

[10] In untabulated analyses, we also examine the effect of $Vega$ on next year's systematic risk, which is an approach commonly used by prior studies on both banks (DeYoung et al., 2013) and non-financial firms (Armstrong and Vashishtha, 2012). We find a statistically significant positive effect of $Vega$ on systematic risk when using OLS regressions, but an insignificant coefficient estimate on $Vega$ in the matching regression, suggesting that the relation between $Vega$ and systematic risk is also driven by matching between executives and banks. We focus primarily on systemic risk given its importance and specificity to the banking industry. Moreover, the proxy for systematic risk relies on realized risk, which may capture additional features of the bank (e.g., information environment) that are not necessarily related to risk-taking. In addition, the measure presumes that market participants can adequately decipher risk-taking activities in real time, which may not be a valid assumption in the banking industry given the greater opacity of their activities (Morgan, 2002).

[11] Inferences are similar for $\Delta CoVaR$ as well as different horizons of both systemic risk measures.

5.1.2 Boom vs Bust

This section examines whether the relation between $Vega$ and systemic risk varies over the business cycle. Prior research discusses how different aspects of banks' operating environment can result in increased risk-taking during "boom" periods relative to "bust" periods. Acharya and Naqvi (2012) illustrate how excess liquidity leads to loosened lending standards, which subsequently results in increased risk-taking in the lending portfolio. Similarly, Dell'Ariccia and Marquez (2006) show how information asymmetry between borrowers and banks can lead to loosened credit standards and lending booms during which banks take on additional risk in their lending portfolios. They also illustrate how lending booms increase the probability of a banking crisis, which is consistent with the eventual realization of negative outcomes stemming from the boom period actions (e.g., increased lending to lower quality borrowers). Finally, Ruckes (2004) shows that competition can influence lending standards, leading to more risk-taking in boom periods when credit standards are lower. Collectively, these studies suggest that risk-taking opportunities–especially those related to systemic risk–can vary across the business cycle. Accordingly, we separately estimate the relation between vega and systemic risk during boom and bust periods.

[Insert Table 4]

Table 4 presents results from estimating the control function regression using MES as the dependent variable and allowing the coefficient on $Vega$ to vary according to whether MES belongs to a "bust" year.[12] Specifically, $Bust$ is equal to one if MES corresponds to the years 2001, 2008, or 2009, and is equal to zero otherwise. We find a positive and significant coefficient on the interaction, $Vega*Bust$, across all three columns. This suggests that $Vega$ leads to increased systemic risk that manifests during economic downturns. In terms of economic magnitude, a one standard deviation increase in $Vega$ corresponds to an additional 0.42, 0.65, and 0.60 percentage points of systemic, respectively, if the economy experiences a downturn one-, two-, or three-years ahead. Alternatively, during boom periods, the economic magnitude of the effect of $Vega$ on systemic risk is 0.10, 0.14, and 0.09 percentage points for one-, two-, and three-year ahead systemic risk, respectively.

[12] For the sake of brevity, we do not tabulate results for the remaining analyses using $\Delta CoVaR^{1\%}$, but note that we obtain similar inferences using this alternative measure.

5.2 Vega and bank operations

This section examines why a relation between $Vega$ and systemic risk exists (i.e., the operational decisions that bank executives make that bring about the systemic risk that we documented in the previous sections). We first examine whether Vega causes banks to change their economic activities. We focus on two types of activities, namely non-agency mortgage backed securities ($MBSNA$) and commercial and industrial (C&I) loans ($CommLoans$). C&I loans generally result in higher realized loss rates compared to other loan types (e.g., mortgages) and as such, are viewed as riskier (DeYoung et al., 2013). Non-agency mortgaged backed securities are those that are not guaranteed by a government agency or government sponsored enterprise (i.e., Ginnie Mae, Frannie Mae, Freddie Mac) and were the center of focus in the recent financial crisis. We examine the relation between these activities and $Vega$ in Table 5, using control function regressions.

[Insert Table 5]

Table 5 demonstrates significant positive effects of $Vega$ on both non-agency mortgage backed securities and C&I loans. A one standard deviation increase in $Vega$ increases $CommLoans$ and $MBSNA$ by 3.6 percentage points and 2.3 percentage points, respectively. These results suggest that a higher $Vega$ leads to more activities that are considered to be risky.

6 Conclusion

Our paper provides evidence on whether bank executives' contractual incentives *cause* risk-taking using a novel approach developed by Klein and Vella (2010). We find that bank executives' portfolio vega *causes* systemic risk that manifests with a lag only during economic downturns. We also examine the specific risky activities that are motivated by contractual incentives and find that vega is associated with future commercial and industrial (C&I) lending and investments in non-agency mortgage-backed securities. Collectively, our results suggest that vega *causes* systemic risk in downturns and that two potential mechanisms through which this effect occurs are C&I lending and non-agency mortgaged backed securities.

We contribute to prior research in several ways. First, we provide evidence on whether executives' contractual incentives *cause* risk-taking at banks. Prior literature provides mixed evidence regarding the association, leaving open the possibility that any previously-documented associations are driven by matching between executives and banks. Second, we focus primarily on systemic risk, which is of primary concern to bank regulators, and several bank-specific activities that may give rise to future systemic risk. Finally, we examine several windows over which the effect of vega on systemic risk may manifest, consistent with prior theoretical studies indicating that the outcomes of risk-taking behavior may not arise until economic downturns.

These findings also have implications for the compensation guidelines issued by regulators under Section 956. Our evidence suggests that the regulator presumption that contractual incentives *cause* risk-taking is appropriate, but that this effect may only manifest during economic downturns. Thus, the relation may be more nuanced than previously thought.

References

Acharya, V. and H. Naqvi (2012). The seeds of a crisis: A theory of bank liquidity and risk taking over the business cycle. *Journal of Financial Economics 106*(2), 349–366.

Acharya, V. V., L. H. Pedersen, T. Philippon, and M. Richardson (2017). Measuring systemic risk. *The Review of Financial Studies 30*(1), 2–47.

Adrian, T. and M. K. Brunnermeier (2016). Covar. *American Economic Review 106*(7), 1705–41.

Armstrong, C. S. (2013). Discussion of "CEO compensation and corporate risk-taking: Evidence from a natural experiment". *Journal of Accounting and Economics 56*(2-3), 102–111.

Armstrong, C. S. and R. Vashishtha (2012). Executive stock options, differential risk-taking incentives, and firm value. *Journal of Financial Economics 104*(1), 70–88.

Aubuchon, C. P., D. C. Wheelock, et al. (2010). The geographic distribution and characteristics of us bank failures, 2007-2010: do bank failures still reflect local economic conditions? *Federal Reserve Bank of St. Louis Review 92*(5), 395–415.

Bernanke, B. S. and M. Gertler (1995). Inside the black box: the credit channel of monetary policy transmission. *Journal of Economic Perspectives 9*(4), 27–48.

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of Macroeconomics 1*, 1341–1393.

Boyallian, P. and P. Ruiz-Verdú (2017). Leverage, CEO risk-taking incentives, and bank failure during the 2007–2010 financial crisis. *Review of Finance 1*, 1–43.

Buser, S. A., A. H. Chen, and E. J. Kane (1981). Federal deposit insurance, regulatory policy, and optimal bank capital. *The journal of Finance 36*(1), 51–60.

Cartwright, N. (1979). Causal laws and effective strategies. *Nous*, 419–437.

Chen, C. R., T. L. Steiner, and A. M. Whyte (2006). Does stock option-based executive compensation induce risk-taking? An analysis of the banking industry. *Journal of Banking & Finance 30*(3), 915–945.

Cheng, I. H., H. Hong, and J. A. Scheinkman (2015). Yesterday's heroes: Compensation and risk at financial firms. *The Journal of Finance 70*(2), 839–879.

Chesney, M., J. Stromberg, and A. F. Wagner (2012). Risk-taking incentives and losses in the financial crisis. *Working Paper*.

Coles, J. L., N. D. Daniel, and L. Naveen (2006). Managerial incentives and risk-taking. *Journal of financial Economics 79*(2), 431–468.

Core, J. and W. Guay (1999). The use of equity grants to manage optimal equity incentive levels. *Journal of Accounting and Economics 28*(2), 151–184.

Core, J. and W. Guay (2002). Estimating the value of employee stock option portfolios and their sensitivities to price and volatility. *Journal of Accounting research 40*(3), 613–630.

Core, J. E., W. R. Guay, and R. E. Verrecchia (2003). Price versus non-price performance measures in optimal ceo compensation contracts. *The Accounting Review 78*(4), 957–981.

De Bandt, O. and P. Hartmann (2000). Systemic risk: A survey. *Working paper*.

Dell'Ariccia, G. and R. Marquez (2006). Lending booms and lending standards. *The Journal of Finance 61*(5), 2511–2546.

Dewatripont, M., J.-C. Rochet, and J. Tirole (2010). *Balancing the banks: Global lessons from the financial crisis*. Princeton University Press.

DeYoung, R., E. Y. Peng, and M. Yan (2013). Executive compensation and business policy choices at us commercial banks. *Journal of Financial and Quantitative Analysis 48*(1), 165–196.

Engelberg, J. E. and C. A. Parsons (2011). The causal impact of media in financial markets. *The Journal of Finance 66*(1), 67–97.

Fahlenbrach, R. and R. M. Stulz (2011). Bank CEO incentives and the credit crisis. *Journal of Financial Economics 99*(1), 11–26.

Freixas, X. and J.-C. Rochet (2008). *Microeconomics of banking*. MIT press.

Gande, A. and S. Kalpathy (2017). Ceo compensation and risk-taking at financial firms: Evidence from us federal loan assistance. *Journal of Corporate Finance 47*, 131–150.

Guay, W. R. (1999). The sensitivity of CEO wealth to equity risk: an analysis of the magnitude and determinants. *Journal of Financial Economics 53*(1), 43–71.

Haugen, R. A. and L. W. Senbet (1981). Resolving the agency problems of external capital through options. *The Journal of Finance 36*(3), 629–647.

Heckman, J. J. (1976). Simultaneous equation models with both continuous and discrete endogenous variables with and without structural shift in the equations. *Studies in Nonlinear Estimation, Ballinger*.

Heckman, J. J. (1978). Dummy endogenous variables in a simultaneous equation system. *Econometrica 46*(4), 931–959.

Heckman, J. J. (1980). Addendum to sample selection bias as a specification error. *Evaluation Studies Review Annual 5*.

Heckman, J. J. and R. Robb (1985). Alternative methods for evaluating the impact of interventions: An overview. *Journal of Econometrics 30*(1-2), 239–267.

Heckman, J. J. and R. Robb (1986; reprinted 2000). Alternative methods for solving the problem of selection bias in evaluating the impact of treatements on outcomes. *Drawing Inferences from Self-selected Samples, Lawrence Erlbaum Associates.*

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association 81*(396), 945–960.

Hölmstrom, B. (1979). Moral hazard and observability. *The Bell journal of economics*, 74–91.

Hölmstrom, B. (1982). Moral hazard in teams. *The Bell Journal of Economics*, 324–340.

Houston, J. F. and C. James (1995). Ceo compensation and bank risk is compensation in banking structured to promote risk taking? *Journal of Monetary Economics 36*(2), 405–431.

Hubbard, R. G. and D. Palia (1995). Executive pay and performance evidence from the us banking industry. *Journal of Financial Economics 39*(1), 105–130.

Ivashina, V. and D. Scharfstein (2010). Bank lending during the financial crisis of 2008. *Journal of Financial economics 97*(3), 319–338.

John, K. and Y. Qian (2003). Incentive features in ceo compensation in the banking industry. *Federal Reserve Bank of New York Economic Policy Review 9*(1), 109–121.

Klein, R. and F. Vella (2010). Estimating a class of triangular simultaneous equations models without exclusion restrictions. *Journal of Econometrics 154*(2), 154–164.

Laffont, J.-J. (1998). *Competition, information and development.* World Bank.

Lambert, R. A., D. F. Larcker, and R. E. Verrecchia (1991). Portfolio considerations in valuing executive compensation. *Journal of Accounting Research*, 129–149.

Larcker, D. F., G. Ormazabal, and D. J. Taylor (2017). Risk-taking incentives in bank holding companies and the financial crisis. *Working paper*.

Low, A. (2009). Managerial risk-taking behavior and equity-based compensation. *Journal of Financial Economics 92*(3), 470–490.

Manski, C. F. (2010). Partial identification in econometrics. In *Microeconometrics*, pp. 178–188. Springer.

Morgan, D. P. (2002). Rating banks: Risk and uncertainty in an opaque industry. *American Economic Review 92*(4), 874–888.

Pearl, J. (2001). Causal inference in statistics: a gentle introduction.

Rosenbaum, P. (2010). *Observational Studies, 2nd ed.* Berlin: Springer Series in Statistics.

Ross, S. A. (2004). Compensation, incentives, and the duality of risk aversion and riskiness. *The Journal of Finance 59*(1), 207–225.

Ruckes, M. (2004). Bank competition and credit standards. *Review of Financial Studies 17*(4), 1073–1102.

Saunders, A., E. Strock, and N. G. Travlos (1990). Ownership structure, deregulation, and bank risk taking. *The Journal of Finance 45*(2), 643–654.

Sørensen, M. (2007). How smart is smart money? A two-sided matching model of Venture Capital. *The Journal of Finance 62*(6), 2725–2762.

Tamer, E. (2010). Partial identification in econometrics. *Annu. Rev. Econ. 2*(1), 167–195.

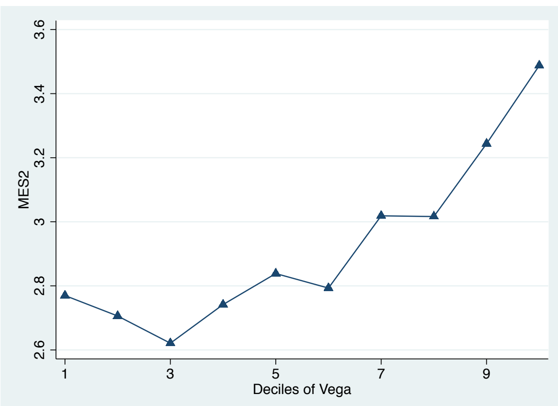

(a) Full Sample

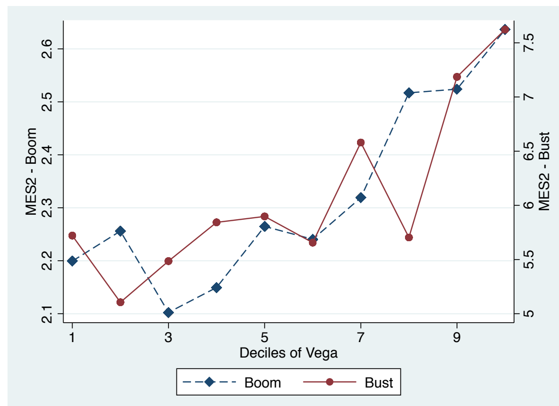

(b) Boom vs. Bust

Figure 1: Vega and Systemic Risk

This table presents the relation between of MES and $Vega$. The vertical axis represents MES, the marginal expected shortfall, defined in Section 3.1. The horizontal axis represents deciles of $Vega$. The measurement of MES leads the measurement of $Vega$ by two years. The left panel presents the relation for the full sample (Figure 1a). The right panel presents the relation separately for periods of economic booms, $Boom$, and economic downturns, $Bust$, where $Bust = 1$ if MES is measured during 2001, 2008, or 2009, and zero (one) otherwise. The solid red line captures economic downturns and the dashed blue line captures economic booms.

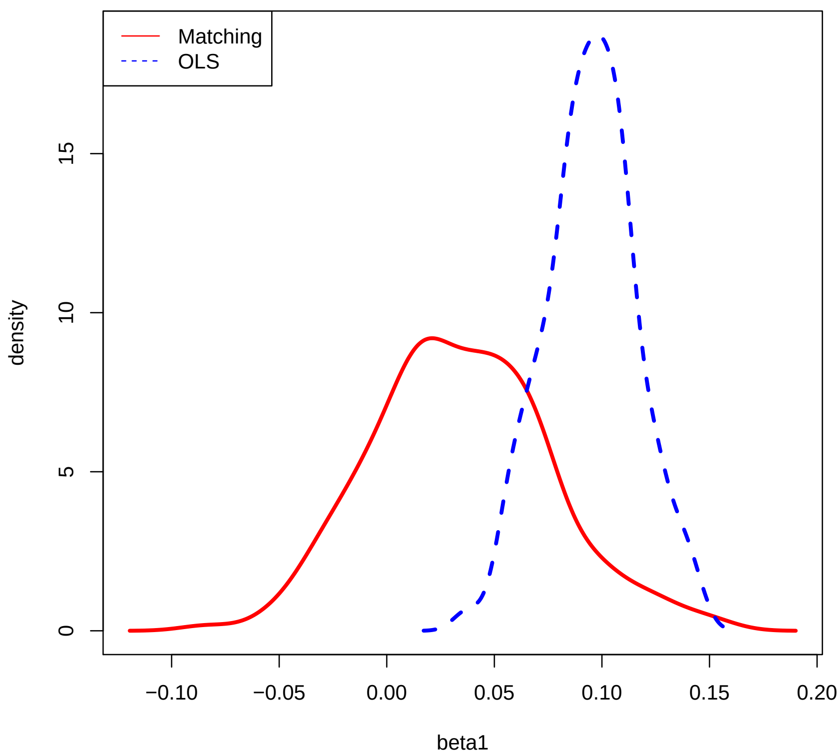

Figure 2: OLS vs Control Function

This figure presents the coefficient distribution from OLS and control function regressions using model (8) and MES_t as the dependent variable. The distribution of the coefficient estimates is computed from bootstrapping. The solid red line represents the density of coefficient estimates from the control function regression. The dotted blue line represents the density of coefficient estimates from the OLS regression. Both regressions include control variables and year fixed effects.

Table 1: Descriptive Statistics

This table presents descriptive statistics for all variables included in each of the analyses. The sample period is 1994 - 2013. $Vega$ and $Delta$ are the log of portfolio level vega and delta, respectively, measured following Core and Guay (2002), for the top five executives. MES is marginal expected shortfall, the average bank return on the worst 5% days for the market during the year, following Acharya et al. (2017), multiplied by 100. $\Delta CoVaR^{1\%}$ is the contribution of bank i to the systemic risk of the banking system, following Adrian and Brunnermeier (2016). Both MES and $\Delta CoVaR^{1\%}$ are multiplied by -1 so that larger values correspond to greater systemic risk. $MBSNA$ is non-agency mortgage backed securities scaled by total available-for-sale investments, multiplied by 100. $CommLoans$ is commercial and industrial loans scaled by total loans, multiplied by 100. $Size$ is the log of total assets. BTM is the ratio of book equity to market value of equity. $Capital$ is the ratio of equity to total assets. $Growth$ is the annual growth rate in total assets. Continuous variables are winsorized at the 1st and 99th percentiles.

	N	Mean	SD	25%	50%	75%
Systemic risk measures and business activities						
MES	1339	2.875	2.187	1.424	2.151	3.574
$\Delta CoVaR^{1\%}$	1339	1.739	0.758	1.195	1.595	2.160
MBSNA	1339	5.907	10.268	0.000	1.266	7.483
CommLoans	1339	21.370	11.869	13.762	19.710	26.981
Compensation variables						
Vega ($000s)	1339	281.7	521.2	30.7	79.2	282.0
Delta ($000s)	1339	939.4	1357.6	142.0	415.5	1074.6
Bank characteristics						
Size	1339	9.712	1.453	8.614	9.405	10.566
BTM	1339	0.668	0.387	0.409	0.573	0.814
Capital	1339	0.092	0.023	0.076	0.090	0.106
Growth	1339	0.123	0.195	0.017	0.074	0.161

Table 2: Standard Deviation Ratio by Region and Year

This table presents the standard deviation ratio between each of the risk outcome measures and $Vega$ by region in Panel A and by year in Panel B. MES is marginal expected shortfall, the average bank return on the worst 5% days for the market during the year, following Acharya et al. (2017). $\Delta CoVaR$ is the contribution of bank i to the systemic risk of the banking system, following Adrian and Brunnermeier (2016).

Panel A

Region	MES	$\Delta CoVaR^{1\%}$
1	1.381	0.465
2	1.168	0.455
3	1.270	0.432
4	1.260	0.382

Panel B

Year	MES	$\Delta CoVaR^{1\%}$
1994	0.230	0.302
1995	0.560	0.338
1996	0.549	0.470
1997	0.631	0.475
1998	0.410	0.436
1999	0.698	0.652
2000	0.499	0.429
2001	0.531	0.451
2002	0.405	0.402
2003	0.302	0.253
2004	0.267	0.240
2005	0.252	0.291
2006	0.377	0.435
2007	1.080	0.571
2008	1.326	0.514
2009	0.512	0.272
2010	0.544	0.314
2011	0.354	0.240
2012	0.185	0.197
2013	0.193	0.178

Table 3: Vega and Systemic Risk

This table presents results of regressing the systemic risk measures on $Vega$ and control variables. Panel A presents results using MES as the systemic risk measure, and Panel B presents results using $\Delta CoVaR^{1\%}$. The panels present systemic risk at multiple intervals $(t+s)$, including one year ahead (t), two years ahead $(t+1)$, and three years ahead $(t+2)$ of the measurement of $Vega$. Columns (1), (3), and (5) in each panel present results using OLS estimation, while columns (2), (4), and (6) present similar results using the control function regression estimation. 95% confidence intervals are reported in parentheses below the coefficient estimates. Continuous variables are winsorized at the 1st and 99th percentiles.

Panel A

	(1)	(2)	(3)	(4)	(5)	(6)
	MES_t		MES_{t+1}		MES_{t+2}	
	OLS	Control Function	OLS	Control Function	OLS	Control Function
$Vega_{t-1}$	0.095	0.034	0.103	0.028	0.074	0.023
	[0.059, 0.129]	[-0.033, 0.106]	[0.071, 0.133]	[-0.039, 0.088]	[0.046, 0.1]	[-0.046, 0.089]
$Delta_{t-1}$	-0.027	0.032	0	0.076	0.037	0.088
	[-0.08, 0.03]	[-0.048, 0.102]	[-0.048, 0.056]	[0.005, 0.156]	[-0.017, 0.08]	[0.017, 0.159]
$Size_{t-1}$	0.196	0.187	0.201	0.192	0.166	0.159
	[0.145, 0.246]	[0.134, 0.242]	[0.154, 0.248]	[0.145, 0.24]	[0.127, 0.209]	[0.118, 0.204]
BTM_{t-1}	0.719	0.776	0.344	0.405	0.164	0.203
	[0.487, 0.945]	[0.55, 1.006]	[0.133, 0.522]	[0.203, 0.622]	[0.04, 0.296]	[0.067, 0.371]
$Capital_{t-1}$	-2.716	-2.807	-1.239	-1.239	-1.647	-1.658
	[-5.424, -0.281]	[-5.486, -0.397]	[-3.972, 1.262]	[-4.006, 1.368]	[-3.914, 0.67]	[-3.862, 0.678]
$Growth_{t-1}$	0.287	0.261	0.368	0.353	0.21	0.21
	[0.041, 0.551]	[0.03, 0.525]	[0.13, 0.624]	[0.106, 0.605]	[-0.045, 0.479]	[-0.047, 0.465]
Constant	0.064	-0.057	0.851	0.688	1.751	1.641
	[-0.332, 0.386]	[-0.459, 0.35]	[0.477, 1.285]	[0.299, 1.158]	[1.41, 2.112]	[1.27, 2.037]
ρ		0.098		0.121		0.085
		[-0.033, 0.214]		[0.012, 0.236]		[-0.031, 0.216]
Obs	1339	1339	1339	1339	1339	1339
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Panel B

	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta CoVaR_t^{1\%}$		$\Delta CoVaR_{t+1}^{1\%}$		$\Delta CoVaR_{t+2}^{1\%}$	
	OLS	Control Function	OLS	Control Function	OLS	Control Function
$Vega_{t-1}$	0.075	0.043	0.071	0.023	0.065	0.027
	[0.058, 0.091]	[0.002, 0.09]	[0.055, 0.085]	[-0.02, 0.068]	[0.047, 0.081]	[-0.016, 0.066]
$Delta_{t-1}$	0.119	0.151	0.126	0.174	0.127	0.166
	[0.094, 0.144]	[0.101, 0.202]	[0.102, 0.15]	[0.126, 0.222]	[0.098, 0.156]	[0.122, 0.217]
$Size_{t-1}$	0.248	0.243	0.236	0.23	0.227	0.221
	[0.224, 0.269]	[0.218, 0.266]	[0.213, 0.258]	[0.204, 0.254]	[0.201, 0.252]	[0.197, 0.246]
BTM_{t-1}	-0.296	-0.271	-0.292	-0.256	-0.284	-0.256
	[-0.378, -0.207]	[-0.364, -0.177]	[-0.366, -0.219]	[-0.337, -0.178]	[-0.358, -0.207]	[-0.33, -0.17]
$Capital_{t-1}$	1.648	1.594	1.435	1.365	0.813	0.727
	[0.598, 2.778]	[0.544, 2.765]	[0.372, 2.586]	[0.314, 2.505]	[-0.302, 2]	[-0.355, 1.897]
$Growth_{t-1}$	-0.101	-0.107	-0.119	-0.123	-0.183	-0.188
	[-0.256, 0.027]	[-0.261, 0.015]	[-0.258, 0.004]	[-0.271, -0.006]	[-0.304, -0.063]	[-0.317, -0.065]
Constant	0.476	0.413	0.579	0.483	0.996	0.918
	[0.312, 0.647]	[0.205, 0.612]	[0.42, 0.753]	[0.279, 0.667]	[0.782, 1.185]	[0.671, 1.116]
ρ		0.084		0.126		0.104
		[-0.043, 0.208]		[0.004, 0.235]		[-0.015, 0.221]
Obs	1339	1339	1339	1339	1339	1339
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 4: Vega and Systemic Risk: Boom vs. Bust

This table presents results of regressing MES on $Vega$ and the interaction between $Vega$ and $Bust$, which is equal to one if MES is measured during 2001, 2008, or 2009, and zero otherwise. The model is estimated using the control function regression. MES is measured at multiple intervals $(t+s)$, including one year ahead (t), two years ahead $(t+1)$, and three years ahead $(t+2)$. 95% confidence intervals are reported in parentheses below the coefficient estimates. Continuous variables are winsorized at the 1st and 99th percentiles.

	(1)	(2)	(3)
	MES_t	MES_{t+1}	MES_{t+2}
$Vega_{t-1}$	0.055	0.078	0.048
	[-0.016, 0.123]	[0.006, 0.151]	[-0.023, 0.118]
$Vega_{t-1} * Bust_{t+s}$	0.169	0.288	0.29
	[0.055, 0.292]	[0.137, 0.419]	[0.146, 0.429]
$Delta_{t-1}$	-0.014	-0.018	0.026
	[-0.097, 0.065]	[-0.123, 0.07]	[-0.06, 0.098]
$Size_{t-1}$	0.19	0.197	0.165
	[0.138, 0.243]	[0.149, 0.244]	[0.12, 0.21]
BTM_{t-1}	0.729	0.32	0.135
	[0.491, 0.973]	[0.116, 0.513]	[0.002, 0.296]
$Capital_{t-1}$	-2.65	-1.237	-1.654
	[-5.087, -0.247]	[-3.936, 1.677]	[-3.791, 0.597]
$Growth_{t-1}$	0.224	0.392	0.228
	[-0.001, 0.485]	[0.141, 0.643]	[-0.026, 0.51]
Constant	0.15	1.095	1.941
	[-0.257, 0.505]	[0.65, 1.561]	[1.565, 2.325]
ρ	0.02	-0.019	-0.023
	[-0.119, 0.163]	[-0.163, 0.114]	[-0.157, 0.101]
Obs	1339	1339	1339
Year FE	Yes	Yes	Yes

Table 5: Vega and Bank Operations

This table presents results of regressing the business activity measures, $CommLoans$ and $MBSNA$, on $Vega$ and control variables. $CommLoans$ is commercial and industrial loans divided by total loans, and $MBSNA$ is non-agency mortgage backed securities divided by total available-for-sale investments. The model is estimated using the control function regression. 95% confidence intervals are reported in parentheses below the coefficient estimates. Continuous variables are winsorized at the 1st and 99th percentiles.

	(1)	(2)
	$CommLoans_t$	$MBSNA_t$
$Vega_{t-1}$	1.999	1.31
	[0.872, 3.107]	[0.605, 2.021]
$Delta_{t-1}$	-1.019	0.37
	[-2.11, 0.12]	[-0.376, 1.099]
$Size_{t-1}$	1.251	0.566
	[0.595, 1.79]	[0.09, 1.075]
BTM_{t-1}	-1.961	0.413
	[-4.012, 0.067]	[-0.955, 1.83]
$Capital_{t-1}$	-2.454	-1.069
	[-30.064, 22.343]	[-21.244, 19.721]
$Growth_{t-1}$	-2.985	2.902
	[-5.694, -0.401]	[0.139, 6.228]
Constant	25.435	-3.942
	[20.603, 30.238]	[-6.936, -0.453]
ρ	-0.142	-0.214
	[-0.256, -0.011]	[-0.337, -0.089]
Obs	1339	1339
Year FE	Yes	Yes