



**FACULTEIT ECONOMIE
EN BEDRIJFSKUNDE**

**TWEEKERKENSTRAAT 2
B-9000 GENT**

**Tel. : 32 - (0)9 - 264.34.61
Fax. : 32 - (0)9 - 264.35.92**

WORKING PAPER

Carrying the (Paper) Burden: A Portfolio View of Systemic Risk and Optimal Bank Size

Jaap W.B. Bos, Martien Lamers, Victoria Purice

May 2014

2014/882

Carrying the (Paper) Burden: A Portfolio View of Systemic Risk and Optimal Bank Size

including Internet Appendix

Jaap W.B. Bos* Martien Lamers† Victoria Purice‡

Abstract

We examine the relationship between bank size and financial stability by viewing the supervisor of a banking system as an ‘investor’ holding a portfolio of banks. Based on this view, we investigate the role of large banks in determining the systemic risk in this portfolio. Our results, based on book data of U.S. banks and Bank Holding Companies, indicate that the largest banks are consistently overrepresented in the current portfolio compared with the minimum variance portfolio. Moreover, the risk level of the portfolio can be reduced by limiting concentration without sacrificing returns.

JEL classification: G21; C61; E44; E63.

Keywords: Systemic risk; Modern Portfolio Theory; U.S. banking.

We thank Geert Bekaert, Bob DeYoung, Iftekhar Hasan, Lukas Menkhoff, Christophe Moussu, Koen Schoors, Peter Schotman, David Veredas, Rudi Vander Vennet and seminar participants at Ghent University, Maastricht University School of Business and Economics, the International Conference on Money, Banking and Finance (2012), the Annual Conference of the Royal Economic Society (2013), the Spring Meeting of Young Economists (2013), the International Conference of the Financial Engineering and Banking Society (2013) and the European meeting of the Financial Management Association (2013) for helpful comments. The computational resources (Stevin Supercomputer Infrastructure) and services used in this work were provided by Ghent University, the Hercules Foundation and the Flemish Government - department EWI. Martien Lamers acknowledges financial support from the Research Fund of Ghent University (Bijzonder Onderzoeksfonds). The usual disclaimer applies.

*Maastricht University School of Business and Economics, P.O. Box 616, 6200 MD, Maastricht, The Netherlands, j.bos@maastrichtuniversity.nl.

†Ghent University, Department of Financial Economics, St. Pietersplein 5, B-9000 Ghent, Belgium, martien.lamers@ugent.be.

‡Ghent University, Department of General Economics, Tweekerkenstraat 2, B-9000 Ghent, Belgium, victoria.purice@ugent.be.

Portfolio. ORIGIN early 18th century: from Italian portafogli, from portare 'carry' + foglio 'sheet of paper' (from Latin folium).

1 Introduction

As a result of the financial crisis, the health and safety of the financial system is at the heart of many policy agendas. Concerns regarding the financial system tend to relate mostly to commercial banks and their parent holding companies. Policy discussions focus either on the riskiness of individual financial institutions, or on what is broadly termed systemic risk. Regarding individual banks, the key question debated is whether some banks are too big: too big to fail, too big with respect to their country's GDP (Bertay et al., 2013; Demirgüç-Kunt and Huizinga, 2013), too big to produce at minimum average costs (see e.g. Wheelock and Wilson, 2012; Hughes and Mester, 2013), or even too big to rescue. Systemic risk discussions are much broader, and may consider the stability of the financial system itself, the macro effects of a shock to that system, or the optimal supervisory setup for dealing with and minimizing the likelihood of such a shock.

In this paper, we combine these two discussions and investigate whether the size of the largest banks in the system has contributed to an increase in systemic risk. We do so by engaging the reader in a thought experiment. We imagine a bank supervisor as an investor holding a portfolio of banks. Each bank aims to maximize profits, but thereby incurs a certain amount of risk. Given that banks' profits are not all perfectly correlated, the risk-return relationship of the portfolio that the supervisor holds is expected to be better than that of the riskiest banks in the system on their own. Taking the long-term view, the bank supervisor not only wants to minimize risk but is certainly also interested in return, as high charter values may boost the stability of individual banks.

Although we consider our view of the bank supervisor a thought experiment, recent events have shown that its experimental nature is closer to the reality of a crisis than one may at first suspect. In theory, the bank supervisor mainly represents the interests of deposit holders and deposit insurance guarantees those interests to a large extent. However, during the recent crisis, most supervisors went above and beyond that objective. In the U.S., the Trouble Asset Relief Program (TARP) initially provided support in terms of bank equity share purchases valued at more than three times the total amounts of deposits in the system, although much of these funds were later reclaimed as shares were sold in the market. Moreover, many assets were purchased well above their actual value, resulting in an implicit subsidy of the banking sector (Office of the Special Inspector General for the Troubled Asset Relief Program, 2013). Finally, the Safe, Accountable, Fair & Efficient (SAFE) Banking Act proposed in 2012 gives

regulators additional powers to limit bank size in order to lower systemic risk.

Nevertheless, unlike the typical investor, the bank supervisor is seriously limited in buying and selling assets in order to reach or remain at the optimal frontier as depicted in Figure 1. As the crisis has shown, even this highly constrained investor can rebalance the weights of the banks in the portfolio, through orderly liquidation and other interventions by the Financial Stability Board such as the capital surcharge for Systemically Important Financial Institutions (SIFIs).

[Insert Figure 1 near here]

Using this scenario, we pose three questions, each related to the situation depicted in Figure 1. First, we ask whether large banks offer attractive investment opportunities for the bank supervisor or in other words, whether large banks are characterized by a risk-return relationship superior to that of the other banks in the system. This should establish whether the inclusion of large banks has brought the supervisory portfolio closer to the efficient frontier, e.g., moving from C to B in Figure 1.

Second, we examine what would happen to the portfolio's return if the bank supervisor held the minimum variance portfolio. In light of the example in Figure 1, would the return move from B to A , thereby requiring the supervisor to give up return in order to hold a less risky portfolio, or would it move from C to A , allowing the choice of a less risky portfolio without sacrificing returns?

Third, we examine whether the bank supervisor should reduce investments in large banks in order to achieve the minimum variance portfolio. In light of the example in Figure 1, we ask whether the supervisor has to increase or reduce the share of large banks (s_L) in the portfolio by moving to A . We also examine whether the differences in risk and return between the original portfolio held by the bank supervisor and the new minimum variance portfolio merely reflect a change in the weights of banks, or whether they are driven by high correlation of the returns of the largest banks in the original portfolio.

To perform our thought experiment, we examine developments in the U.S. banking market since 1984. Using quarterly data on banks' assets and profits, we construct two types of bank supervisory portfolios. For each Federal Reserve District, we include all unconsolidated commercial banks located within the district. For the U.S. as a whole, we construct a portfolio comprising all Bank Holding Companies. In both cases, we are interested in the design of the minimum variance portfolio and how it compares with the actual portfolio. Distinguishing between portfolio optimization at the country and at the district level allows us to contribute to the debate about the optimal allocation of supervisory powers.

Our findings indicate that the current portfolios are not located on the efficient frontier,

as risk can be reduced without sacrificing return in order to attain the minimum variance portfolio. Moreover, we find that the largest banks in the Federal Reserve portfolios consistently have a significantly lower weight in the hypothetical minimum variance portfolio. In addition, the minimum variance portfolio does not allow for large levels of concentration in the first place, with its weights being much more evenly spread across banks. These results hold even after relaxing some assumptions and allowing the correlation structure to change with the size of banks. In obtaining the minimum variance portfolio, we assume that the supervisor is able to alter the relative weights of the individual banks in the portfolio, along the lines of the Safe, Accountable, Fair & Efficient (SAFE) Banking Act proposed in 2012. The Act was designed to make banks small enough to fail without causing global panic, using regulatory caps.

We argue that our results provide important insights into the optimal design of a portfolio of banks, held by a risk-averse supervisor who prefers to incur the least possible amount of risk. The findings suggest reducing the size of the largest banks in the financial system may not only make individual banks safer and easier to fail or rescue, but can also contribute to a reduction in the riskiness of the system as a whole. For these changes to occur, supervisory powers do not have to be relocated to the highest aggregation level: we find that when permitted to act independently, Federal Reserve Districts' portfolio allocations are very close to what would be achieved by a single national supervisor.

This paper proceeds as follows. In Section 2 we discuss our methodology and data, followed by Section 3 which contains our results. Extensions and policy implications are addressed in Section 4, after which we conclude.

2 Data and Methodology

2.1 Methodology

We regard banks as assets having both a return and a risk component. By placing the banks in this risk-return framework, we analyze not only the profitability of the banks but also the risk inherent in profit maximization. When considering a portfolio of these assets and assuming that returns are not perfectly correlated, a supervisor holding this portfolio can diversify and enjoy a better risk-return trade-off. The diversification opportunities of the supervisor are constrained by the variety of banks available. Of course, banks themselves also individually diversify their loan portfolios, business lines or geographical markets served. However, the dimensions along which banks can diversify are correlated with their size, as larger banks can be assumed to have greater geographical reach and a different mix of activi-

ties compared with smaller banks. Moreover, if enough banks diversify along the same lines, the system as a whole becomes more susceptible to common shocks, even though individual banks themselves seem safer from a micro-prudential point of view. This observation is not new, being noted for instance by Rajan (2005), Wagner (2008), Acharya (2009), Ibragimov et al. (2011) and Allen et al. (2012), who also draw on Modern Portfolio Theory (Markowitz, 1952, 1959) and the insight that portfolio risk can be reduced by choosing assets that are not perfectly correlated with each other.

We contribute to this line of thinking by applying portfolio theory to the banking system in order to investigate the role of large banks in determining systemic (portfolio) risk. The supervisor in our case is able to change the portfolio's risk-return trade-off by altering the size of banks with respect to the system. Although this ability goes beyond the existing mechanisms in place (such as 'Cease and Desist' orders and other 'Prompt Corrective Actions'), we argue that this thought experiment can, at the very least, give us insights into the optimal design of a portfolio of banks. Applying portfolio theory to the banking system is not however all that straightforward, due to some of its strict assumptions. First, the return distribution is assumed to be fully characterized by the first two moments and disregards any tail dependence, even though financial market returns are found to be skewed and fat-tailed. To (partially) mitigate this issue, we use lower-frequency quarterly returns. Second, market participants are assumed to have no influence on prices and return structures of assets in their investment universe, regardless of the weight they are given. In reality, if a bank supervisor were to reduce the size of a bank, its risk-return trade-off would be bound to change as well, as would its correlation structure. Third, an investor is assumed to be able to purchase assets in parcels of any size, meaning that bank sizes could fluctuate heavily between the investor's decision moments. Under a more realistic scenario, the supervisor would be able to change the size of a bank in a limited way, e.g. by only a certain percentage of the bank's assets. Although we initially proceed under the strict setup, the last two assumptions are relaxed at a later stage.

In order to apply portfolio theory and build the regulator's portfolio, we first need to define the return and weight of the assets under consideration. Previous studies have relied mainly on market-based measures when assessing systemic risk (see e.g. De Jonghe, 2010; Adrian and Brunnermeier, 2011; Acharya et al., 2012a; Acharya et al., 2012b; Bisias et al., 2012; Brownlees and Engle, 2012; Engle et al., 2014). Unlike these studies, we instead use the return-on-assets from book data for the returns of the banks. We do so for several reasons, the first being that the aggregated risk concerning the supervisor is not based on the returns and risks of the shareholders of banks, but rather on those of the (productive) assets that they hold. In the event that the regulator has to bail out a bank, saving or guaranteeing its

liabilities will be equivalent to saving or guaranteeing its assets. Second, as shown by Allen and Carletti (2008), in financial crises market prices tend to reflect the amount of available liquidity instead of future earnings. Since these episodes are of particular interest to this analysis, market-based measures might not be appropriate as they could capture liquidity risk instead of systemic risk. Third, accounting data enable us to explore a more extensive sample since market data is only available for a small subset of banks. While listed banks do account for a large percentage of the total banking assets, small banks are potentially a source of (liquidity) contagion through the interbank market (see e.g. Furfine, 1999, 2003; van Lelyveld and Liedorp, 2006; Degryse and Nguyen, 2007). Finally, return-on-assets is a cleaner measure of the underlying profitability, as return-on-equity incorporates management choices with regards to leverage. While our baseline results are based on the book value of the return-on-assets, they are robust to using market-based measures.

The weight that the regulator holds in each bank is calculated as the bank-level total assets divided by the sum of all bank-level total assets available in the portfolio.

According to portfolio theory, the investor's return and risk are calculated as:

$$r_{p,t} = \mathbf{w}'_t \boldsymbol{\mu}_t \quad (1)$$

$$\sigma_{p,t}^2 = \mathbf{w}'_t \boldsymbol{\Sigma}_t \mathbf{w}_t \quad (2)$$

where \mathbf{w}_t is a column vector representing the weights of all banks in period t and $\boldsymbol{\mu}_t$ represents the expected return of the banks, usually defined as the average return of the previous quarters. Furthermore, $\boldsymbol{\Sigma}_t$ represents the expected covariance matrix of these returns and is often replaced by its sample equivalent. The average expected return of the portfolio is given by $r_{p,t}$, while the variance of this set of returns is given by $\sigma_{p,t}^2$ and represent the measure of portfolio (or in our case: systemic) risk at time t .

The supervisor is considered to be risk averse, and to prefer to hold the portfolio with the least amount of risk according to the objective function in Equation 2. Moreover, the supervisor is able to influence systemic risk by changing the weights in the portfolio, assuming that this does not impact the matrix $\boldsymbol{\Sigma}_t$.¹ Minimizing the objective function allows us to compare the differences in portfolio design between the initially realized and the hypothetical minimum variance portfolio.

¹As the recent crisis has shown, it is not unusual for supervisors to intervene through liquidation, (hidden) bailouts or forcing banks to sell off assets to maintain a competitive environment.

To achieve the minimum variance portfolio (MVP), the supervisor solves for:

$$\begin{aligned}
 \underset{\mathbf{w}_t}{\operatorname{argmin}} \quad & \mathbf{w}_t' \Sigma_t \mathbf{w}_t \\
 \text{s.t.} \quad & \mathbf{w}_{i,t} \geq 0 \\
 & \mathbf{1}' \mathbf{w}_t = 1 \\
 & \mathbf{w}_t' \boldsymbol{\mu}_t \geq 0
 \end{aligned} \tag{3}$$

with the addition of several further constraints. First, the supervisor cannot go short in a bank, i.e. no bank can have a negative weight. Second, the weights of banks have to add up to 1 as the existing assets are merely reshuffled, without any being created or destroyed. Finally, a supervisor is also assumed to choose the weights such that the portfolio does not have negative returns.² It follows from these non-linear constraints that no analytical solution is possible, and we therefore rely on a numerical solution.

Using this approach, we investigate the following questions. First, we ask whether systemic risk can be reduced and if so, by how much. To do this, we compare the standard deviation of the initial portfolio with that of the MVP. Second, we investigate whether the supervisor would have to sacrifice returns in order to achieve a lower risk, by comparing the average return of the initial portfolio with that of the MVP. Within the framework described in the introduction, we therefore ask if, in order to reach the MVP, the supervisor has to move along the efficient frontier or shift towards it. Third, we compare the dispersion of weights within each of the two portfolios, by looking at the share of the largest 5% of banks in the initial portfolio and in the MVP.³ Finally, we ask whether the largest 5% of banks in the initial portfolio have retained their relative importance in the MVP by comparing their initial share with the weight they receive in the MVP.

2.2 Data

We perform our analysis on the U.S. banking system, which has several regulatory bodies at different levels. Depending on location, membership status and type, a bank can be regulated by the Federal Reserve System (FED), the Office of the Comptroller of the Currency (OCC) and the Federal Deposit Insurance Corporation (FDIC). While the main supervisory task at district level is carried out by the 12 Federal Reserve Banks, depicted in Figure 2,

²The inclusion of this constraint does impact our results as will be shown later.

³An alternative measure of concentration would be the Herfindahl-Hirschman Index, although using the portfolio weights of the largest banks is more intuitive in this setting. The 5% concentration measure is preferred, since it allows for a better comparison in different-sized banking systems (see Alegria and Schaeck, 2009).

at the national level the main regulatory task is performed by the Federal Reserve's Board of Governors. As a consequence of this division, we consider the regulatory portfolio both at national level and at district level. At a later stage this will allow us to compare the potential importance of the regulatory levels in managing systemic risk.

[Insert Figure 2 near here]

Bank data are obtained on a quarterly basis from the Call Reports for Income and Condition provided by the Federal Reserve System. For the national (FED) portfolio, we consider consolidated Bank Holding Companies (BHCs) as the assets in which the regulator can invest. Data for the BHCs are obtained from the FR Y-9C Forms, between 1986Q3 and 2012Q1. We select only Holding Companies and exclude Insurance/Securities brokers, Utilities and other Non-Depository institutions. At the Federal Reserve District (FRD) level, data on unconsolidated Commercial Banks are retrieved from the FFIEC 031/041 Forms between 1984Q1 and 2010Q4, excluding Savings/Cooperative/Industrial banks as well as Non Deposit Trust companies.

We use balance sheet data instead of financial market data, allowing us to consider all banks that are required to file reports and not only those that are listed on an exchange. Moreover, lower frequency returns are preferred to daily or even weekly returns, in order to comply with the assumption of normality of returns. We collect total assets (bhck/rcfd2170) and net income (bhck/riad4340), deflate both to 2005Q1 dollars using the Producer Price Index provided by the St. Louis Federal Reserve Bank, and filter out banks with return-on-assets exceeding +100% or -100%. This leaves us with 4,694 BHCs across 154,577 bank-year observations and 19,225 commercial banks over 1,132,425 bank-year observations.

Summary statistics for the BHCs and commercial banks are shown in Table 1. In any quarter in our sample, there are between 964 and 2,333 holding companies active in the United States. Due to inflation and a wave of consolidation, the total assets reporting threshold for BHCs was raised from \$150 million to \$500 million in 2006, causing a drop in the number of banks in the sample. Banks controlling less than \$500 million in total assets prior to 2006 are kept in the dataset, since, as will be shown at a later stage, their exit does not affect our results. Given that the distributions of assets and returns are highly unequal, we report percentiles instead of means and standard deviations. The median BHC controlled \$500 million in total assets and reported a net income of \$1.3 million. The table shows the skewness in the distribution of total assets, with the largest 5% of BHCs having total assets ranging between \$14.5 billion and \$2.1 trillion. While the net income of the median holding company is \$1.3 million, again there is a large disparity: the highest earning 5% of BHCs recorded profits ranging from \$44 million to \$6.4 trillion. At the other end of the spectrum,

losses are equally large, partly due to the recent financial crisis, with one holding company reporting a net loss of \$15 trillion in the third quarter of 2008. Since the return-on-assets takes into account the size of the BHC, its values are less extreme compared with those of returns and assets separately, with the mean (0.257%), median (0.292%) and mode (0.325%) lying in close proximity.

The number of reporting commercial banks lies between 6,477 and 14,474 over the 12 districts. Regarding net income and total assets, commercial banks follow a similar pattern to that of Bank Holding Companies, although smaller on average. The total assets disparity is even larger than at the national level, with some banks dwarfing their competitors.

[Insert Table 1 near here]

We proceed by placing every bank in its respective FRD portfolio, defining its weight as the total assets of the bank divided by the sum of total assets in the FRD. The BHCs are analyzed at the national level in a similar manner: the weights are calculated as the individual level of total assets divided by the sum of all total assets of the BHCs. As we have quarterly data over a period of 25 years at our disposal, we perform the analysis using a window of 8 consecutive quarters on which we calculate the expected return and sample covariance matrix, thereby taking into account time-varying correlation.⁴ As a consequence, assets need to have posted data in each consecutive quarter of the window to be included in the analysis.

3 Results

In this section we present the results of our analysis. We first examine the risk-return trade-off between the FED portfolio and the MVP. In a second step, we look at the differences in portfolio allocation between the two systems before comparing their other features. We present the results of the analysis at BHC level graphically, referring the reader to the Internet Appendix for the results on FRD level as they are quantitatively similar, and conclude this section with several robustness tests.

3.1 Are large banks more risky?

Before reporting the results however, we first need to establish the similarities in the return structure between large and small banks. Should large banks have (co)variances different from those of small banks, the assumption that the covariance matrix is independent

⁴In Section 3.5 we show that our results are robust to a different window size.

of size means we would impose an unrealistic structure when large banks are reduced in size, or small banks are made larger.⁵ In Figure 3, we show the two dimensions of the covariance matrix by plotting the densities of the average 8 quarter variance and average 8 quarter pairwise covariance for the largest 5% and the smallest 95% of BHCs. From the Figure, it becomes clear that despite the differences in size, there is ample common support in the individual and common riskiness of bank returns as the distributions overlap almost entirely. These results hold for different time samples, and for both BHCs and Commercial Banks.⁶

[Insert Figure 3 near here]

Of course, even if large and small banks share a common support in the covariance matrix, it does not mean that a bank that changes size will maintain its return structure. At a later stage, we therefore look at banks that have seen large increases or decreases in size, and analyze how this changed the elements of the covariance matrix. Using these average changes in turn allows us to alter the covariance matrix during the numerical optimization, leading to a more realistic portfolio allocation. However, since there are only a limited number of cases on which we can base this analysis, we first proceed by assuming that changing a bank’s size does not influence the structure or level of its returns, and later revisit this assumption.

3.2 What role does inequality play in the risk/return trade-off?

It is quite straightforward to obtain the risk, return and weight distributions for the initial portfolios. By contrast, obtaining the respective MVPs is more cumbersome, as a minimum of 964 and a maximum of 2,333 BHCs are present during the sample period. The Chicago, Kansas City and Dallas portfolios typically contain well over 1,800 banks. The solution is computationally intensive, but is nonetheless obtained after a lengthy optimization process.

Figure 4 shows the results of the portfolio optimization based on BHC data: panels 4a and 4b display the risk-return trade-off for the FED portfolio and the MVP. They show that with different weights, portfolio risk is effectively eliminated in the MVP. Whereas the standard deviation of the FED portfolio spikes during the Savings-and-Loans and subprime mortgage crises, that of the MVP remains stable at around zero. One would expect a portfolio with lower (minimized) risk to have a lower return as well. However, this less risky system has positive returns throughout the sample period with values that closely track the

⁵In this respect, our approach is similar to verifying whether the assumption of a ‘common support’ holds in propensity score matching (Heckman et al., 1998). If there is enough ‘common support’ between small and large banks, the assumption that the supervisor could change assets without having these actions leading to a different return structure is more realistic.

⁶Full results available upon request.

actual returns. Therefore, we conclude that the initial portfolio does not lie on the efficient frontier, as risk is reduced while the level of returns has been maintained. Panel 4c shows that the lower risk is achieved in the MVP through a markedly lower concentration than in the FED portfolio. While in reality the weight of the largest banks lies between 65% and 90%, the concentration in the MVP is on average 13% and at most 52%. More interestingly, the largest banks in the FED portfolio see their cumulative weight reduced to at most 15% of total assets in the MVP.⁷

[Insert Figure 4 near here]

The same picture as for the BHC data emerges if we look at the separate FRD portfolios, reported in the Internet Appendix. Table A.1. summarizes the differences in risk and return between the FRD portfolios and their MVPs. Evidence of returns over the whole period is mixed: some FRDs outperform their MVPs, whereas others exhibit lower returns. One interesting fact is that in the boom period of 1994Q1 - 2006Q4, we find that all FRDs outperform their MVP counterparts in terms of returns. Regarding size disparity, Table A.2. shows that the MVPs consistently have a much lower level of concentration compared to their FRD portfolios, the difference ranging on average between 44% and 78% throughout the sample period.

Two further remarks are in order. First, the portfolio standard deviation in Figure 4 seems to be higher during the S&L crisis than in the subprime crisis. This result can be explained by the fact that a standard deviation, unlike a correlation coefficient, is not a dimensionless number and can only be interpreted as a function of its mean. Given that the returns of the FED portfolio and the MVP are of similar size throughout the sample, we can compare their standard deviation in each quarter, but not between quarters. To allow for comparison over time, we de-mean the risk measure, although this leads to a loss in direct interpretation.⁸ Figure 4d displays the de-measured risk measure, and portfolio risk now shows a higher peak in the subprime crisis than in the S&L crisis.

The second issue we want to address is the MVP's high concentration during 2008 and 2009. This spike can be explained by the fact that up to 87% of BHCs reported lower average returns than in the previous quarter and 40% of all BHCs recorded losses. Given the number of banks involved, it is possible that concentration rose because of this increase in correlation between average returns. However, it is also possible that the MVP weights

⁷Although not reported in the Figure, the Gini coefficient exhibits the same pattern, as the average coefficient of the FED portfolio is 0.9 compared to 0.2 for the MVP.

⁸We do so by calculating the normalized range of returns. Since the portfolio standard deviation is the standard deviation of the weighted average returns during the last 8 quarters, we define the normalized range as the (maximum-minimum)/median of this set.

are chosen to avoid violating the no-loss constraint. To test the latter possibility, we ran the analysis excluding the no-loss constraint but still find the same spike, indicating that, indeed, higher weights are given to banks that share a lower correlation. Since these are few in number, they therefore have to receive a higher weight in order to minimize portfolio risk.

The results here suggest that inequality and concentration play an important role in the risk-return trade-off with which a regulator is faced. In this simple exercise, reducing inequality drives down risk without significantly affecting returns at both FED and FRD level. These findings indicate that regardless of the regulatory level, supervisors need to be concerned when looking at the optimal design of their portfolio not only with a bank's individual size but also its size relative to the system. Moreover, the rebalancing of weights does not appear to be random. We find that in order to obtain a less risky portfolio, a supervisor has to reduce holdings of the currently largest banks and create a more equal system. In reality, the largest banks have had a much higher share in the portfolio compared with that in the MVP, and even increased their weight from 65% to 90% during the sample period. Moreover, the financial industry as a whole has also grown in relation to GDP, to the extent that Carvalho and Gabaix (2013) attribute the recent rise in macroeconomic volatility mainly to this growth in combination with idiosyncratic shocks to the largest banks. Indeed, the share of the largest banks in the current portfolio relative to GDP has increased from 35% at the beginning to 75% at the end of the sample. Given the evidence presented here, combined with the finding of Carvalho and Gabaix (2013), large banks do seem to play an increasingly important role not only in the banking system, but also in the broader economy and its volatility.

3.3 What are the other features of the Minimum Variance Portfolio?

Given that we have seen that lower portfolio risk is achievable in a less concentrated banking system, this raises the question of what causes the largest banks in the FED portfolio to have such consistently low weights in the MVP, and how they differ from the largest banks in the MVP. The two components that determine the weighting decision are on the one hand individual bank risk, as measured by the standard deviation of the returns, and on the other hand the correlation between these returns.⁹ Figure 5a shows the average standard deviation of return-on-assets for the largest 5% of banks in the FED portfolio and MVP, while Figure 5b displays the average pairwise correlation among the largest 5% of banks in their respective portfolios. We see that the standard deviations of the returns of these largest

⁹In this section, we show the correlation of returns instead of the covariance as it is more easily interpretable.

banks are on average twice as high in the MVP, yet the average correlation coefficient is much more stable compared to the FED portfolio. In the two crisis periods, the FED portfolio pairwise correlation spikes to average values of 0.6, almost three times larger than the MVP. Nonetheless, even when the individual risk and correlation are lowest in the FED portfolio, its risk is still higher than that of the MVP.

[Insert Figure 5 near here]

These observations add to the evidence that in this context, weight plays a significant role in determining the level of risk, as it magnifies the effect of increased correlation and individual riskiness. When both components have low values, systemic risk is low, even in a highly concentrated market. However, when they increase in crisis periods, systemic risk increases dramatically if size inequality is high. As already shown in Gabaix (2011), individual shocks to firms have the potential to lead to aggregate volatility when the size distribution of an economy is heavy tailed, something that also holds for the banking system (see e.g. Janicki and Prescott, 2006; Blank et al., 2009). In terms of portfolio theory, the variances will dominate the covariances in crisis periods due to the large disparity in weight. Since the movements in correlation can be extremely volatile and difficult to control or even predict, the best tool for keeping systemic risk low in this context is to limit concentration.

We note, however, that these results do not necessarily imply a cap on the size of banks. On the one hand, a system with only small banks is subject to the Too-Many-To-Fail problem since they could herd, thereby acting as one large entity (Acharya and Yorulmazer, 2007; Brown and Dinç, 2011). In this setting, herding would be picked up via an increase in correlation of returns, posing a systemic threat despite a lack in concentration. Our results, however, indicate that the correlations of small banks are relatively stable over the sample period, and would therefore not pose a systemic threat. On the other hand, the system can also be diversified by limiting activities that banks can undertake and/or markets it can serve, provided they operate in their own (uncorrelated) niche. This point was also touched on by Loutskina and Strahan (2011), who found that increased geographic diversification went hand-in-hand with a decline in loan monitoring by lenders prior to the financial crisis.

To determine the characteristics of banks which have been heavily reweighed, we construct a crude industry level balance sheet for both the FED portfolio and the MVP. We use the weights allocated to each bank to construct this weighted average balance sheet, which is shown in Figure 6. The allocation of assets in the FED portfolio shows the increasing importance of trading assets at the expense of loans, whereas this trend is less evident in banks favored in the MVP. While the FED balance sheet has less than 40% of assets invested in loans at the end of the sample, that of the MVP remains close to 50%. On the liabili-

ties side, the FED portfolio is more reliant on non-deposit funding than the MVP balance sheet. We observe that at the end of the sample, the FED portfolio uses about 10% more of these non-deposit liabilities than the MVP, although this difference was much more apparent before the recent crisis.

As a reflection of the industry asset composition, the non-interest income/total income ratio for each portfolio is shown in Figure 6e. We observe that with exception of the crisis years, there has been a significant increase in reliance on non-interest income in the FED portfolio. On the other hand, banks favored in the MVP have a more constant share of non-interest income throughout the sample period. Notwithstanding the financial crisis, the gap between the portfolios has been steadily increasing since 1990.

[Insert Figure 6 near here]

Given these results, we conclude that the fictitious banking industry in the MVP is characterized by retail banking, as higher weights are given to banks that are mainly funded by deposits, make loans and therefore rely less on non-interest income. Since banking activities are not a direct input in the minimization of the portfolio risk, it is their influence on the behavior of the returns which drives these findings. Indeed, this is in agreement with a growing literature emphasizing the role of income diversification in financial instability. For instance, Stiroh (2004) and Stiroh and Rumble (2006) find that non-interest income reduces aggregate profits while increasing risk. More recent evidence by De Jonghe (2010) shows that systemic risk is exacerbated by banks diversifying into activities other than lending, due to increasing correlations between income streams. This finding was also corroborated by Adrian and Brunnermeier (2011), Brunnermeier et al. (2012), DeYoung and Roland (2001) and DeYoung and Torna (2013). Huang and Ratnovski (2011), meanwhile, argue that wholesale lenders have lower incentives for costly monitoring, leading to large (and inefficient) fluctuations of loans on negative public signals, a problem not encountered in relationship banks. Finally, Boot and Ratnovski (2012) find that although there are initial benefits for banks from starting trading activities, beyond a critical point inefficiencies dominate and trading becomes increasingly risky. On the funding side, Fahlenbrach et al. (2012) emphasize that banks with increasing balance sheets through the use of short term non-deposit liabilities performed poorly during the last crises.

3.4 How easily is the Minimum Variance Portfolio obtained?

In our baseline scenario, the supervisor is able to switch assets rapidly from one bank to another on a quarterly basis to obtain the MVP. Although reweighing also occurs naturally

in the FED portfolio via mergers and acquisitions, bank entry and exit or bailouts, the MVP would not be a realistic approximation if reweighing was much higher than in reality. In order to assess how stable the MVP is over time compared with the FED portfolio, we therefore calculate both of their turnovers. Turnover is defined as the sum of absolute weight changes in the portfolio between period $t-1$ and t , taking values ranging from zero (no change) to two (where all assets that were held are sold, and all assets that were not held bought). Figure 7 plots the turnover for both portfolios. We observe that the MVP turnover is on average 3 times as high as that of the FED portfolio. The spike in turnover in 2006Q1 is due to changes in the reporting threshold, as the banks that reported in 2005Q4 were considered to be sold in 2006Q1 and proceeds reinvested in the other banks. The MVP can only achieve low risk through a higher level of reweighing, especially in the crisis period. For the district portfolios, the average MVP turnover is around 3 times higher than the actual portfolio, ranging from 2 times for the least concentrated to 6 times for the most concentrated districts. In Section 4, we therefore explore alternative MVPs where the reweighing is restricted, in order to achieve a more realistic turnover.

[Insert Figure 7 near here]

3.5 Robustness

To find out how robust our results are, we test several of our assumptions. The results of these tests are summarized in Table 3, where we evaluate how well risk was reduced while limiting concentration. To this end, we define the ratio $(\sigma_{\text{FED}} - \sigma_{\text{MVP}})/\sigma_{\text{FED}}$, which measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements were possible. A test is regarded as successful when this ratio averages 0.9 or higher throughout the sample period, and when the level of concentration of the largest 5% of banks in the MVP is on average below 50%.

Covariance matrix – 1 We begin with the assumption that the covariance matrix is independent of the size of banks. In Section 3.1 we showed that large and small banks share a common support in the variance and average pairwise covariance. This however, does not imply that a bank which changes in size will maintain the same level of returns or the same structure with regard to other banks. If we knew how the return structure changes due to a change in size, we would be able to adjust the covariance matrix in each iteration of the optimization. To this end, we have identified 15 cases in which BHCs experience a negative

jump in bank size, and 287 where they experience a positive jump in bank size.¹⁰ A jump is defined as an increase/decrease of bank assets of 25% or greater from one quarter to the next, provided that the preceding and following 8 quarters did not show jumps larger than 10% in each of the quarters, nor a cumulative change in the preceding and following 8 quarters of 25%. These last two conditions are imposed to make sure that bank size before and after the jump was relatively stable and that the change in the elements of the covariance matrix can be chiefly attributed to the one-time jump. The banks receiving a negative shock lost 40% of their total assets on average, while banks receiving a positive shock gained 60%. The average changes in the return and covariance matrices are displayed in Table 2. T-tests show that banks experiencing a positive jump in assets have a statistically significant lower return-on-assets, which is likely due to the construction of the variable, and a higher average variance and covariance of these returns. On the other hand, banks experiencing a negative jump do not see changes in their average variance and covariance, and only see a marginally significant higher return-on-assets.

[Insert Table 2 near here]

Using the statistically significant changes for positive and negative jumps, we incorporate the effect in the sample covariance matrix and the return matrix, such that these matrices change dynamically with the weights of the banks.¹¹ We perform two additional robustness tests. In the first test, the vector of expected returns is adjusted according to the changes in weight, and a new sample covariance matrix is estimated with which portfolio risk will be minimized. In the second test, we adjust both the return and covariance matrix based on the changes in weight. Changes in the return and covariance matrices are interpolated if the proposed change in portfolio weight lies between -40% and +60%. As there was no data on changes in assets larger than these bounds, we use any proposed changes beyond them as if they were -40% or +60%, i.e. while a bank can receive an increase in weight higher than +60%, its variance and covariance terms are adjusted as if the weight has only been increased by +60%.¹² In both cases, however, this adds more complexity to the optimization and indeed we find that there are cases in which no improvement in portfolio risk is found. Fortunately an improved solution is still possible in most of the sample period, as can be seen in Figure 8. Interestingly, the findings in the baseline specification seem to be robust and are not influenced by our original assumption. The concentration in the MVP remains much

¹⁰Full details on this identification can be found in the Internet Appendix.

¹¹As the change in return-on-assets for the banks experiencing a negative jump is marginally significant at 11%, we also regard this change as significant.

¹²A scenario under which reweighing of each bank was limited to -40% and +60% yielded no improvements over the actual portfolio. Results are available upon request.

lower than that of the FED portfolio, while simultaneously maintaining a smaller portfolio risk and returns of a similar level.

[Insert Figure 8 near here]

Covariance matrix – 2 Another issue in our baseline setup is that optimization can be quite unstable when using the sample covariance matrix: small changes in the return structure can lead to large differences in the outcome of the portfolio choice. To counter these unstable solutions, Brodie et al. (2009) use regularization of the optimal mean-variance portfolio by including a penalty in the objective function. They show that while introducing this penalty can lead to a sparse choice of weights, opting for a high penalty is equivalent to a constraint that does not allow shorting. Since our supervisor is already assumed not to be able to go short in any banks, we therefore consider the optimization problem stabilized. However, as Jobson and Korkie (1980) have pointed out, a considerable amount of noise is introduced in a sample covariance matrix in a small T and large N setting like ours. In response, Ledoit and Wolf (2003, 2004) proposed a shrinkage based estimator of the covariance matrix to reduce this noise. We use their proposed estimator of the covariance matrix instead of its sample equivalent and re-run the analysis.¹³ The full results are reported in the Internet Appendix and are similar to the baseline specification. We find that the standard deviation of the MVP is now higher compared with the baseline specification, although still lower than that of the FED portfolio, while returns are at a similar level. The levels of concentration in the MVP are basically unchanged, as the largest banks on average still have a weight of 10%, while the largest banks in the FED portfolio have a weight in the MVP of 5%.¹⁴

Starting values Third, we explore the optimization starting values and choice of the length of the moving window. Given the fact that we are dealing with many banks, the minimization of the portfolio risk is likely to be a complex, highly nonlinear problem comprising multiple minima/solutions. The starting values, which are selected as the weights in the original portfolio, can have a substantial impact on whether a global or local minimum is found and in what direction the distribution of weights will move. To account for this possible bias, we run two robustness tests. In the first, we choose starting values based on an equally weighted portfolio. In the second, we run, for the BHC data only, 100 repetitions per quarter using

¹³The code for estimating the covariance matrix is obtained from:
http://www.ledoit.net/honey_abstract.htm

¹⁴As portfolio risk is not minimized by at least 90%, however, we do not consider this test to be successful in Table 3.

randomized starting values.¹⁵ Both tests show that the results are almost identical to the baseline specification, and we therefore refer the reader to the Internet Appendix for the full results.

Length of rolling window Finally, we explore alternative lengths of the rolling window. So far we have taken an 8 quarter time frame to estimate the sample covariance matrix. However, it could be argued that using more data to estimate it would be less noisy and less prone to outliers. Taking this into account, we rerun the analyses using a 16 quarter window and report the full results in the Internet Appendix. For both the BHCs and commercial banks, results follow similar patterns to those using 8 quarter windows: the largest banks are still shown to be consistently overweighted compared with their MVP counterparts, where lower risk is achieved while keeping returns at a comparable level. Table 3 summarizes the results of the robustness tests in this section.

[Insert Table 3 near here]

4 Extensions and Policy implications

As we have seen in Section 3.4, the amount of turnover needed to lower systemic risk is three times higher in the MVP. In this section, we therefore explore some more realistic scenarios, and discuss implications for policy resulting from the analysis. We first look into several weighting alternatives. Besides analyzing these other weighting methods, we repeat our analysis on a smaller and more realistic sample of banks. Finally, we discuss whether optimization at district level also results in a lower countrywide systemic risk. Similar to Section 3.5, we summarize all results in Table 4, where we again define a test successful if the ratio $(\sigma_{\text{FED}} - \sigma_{\text{MVP}})/\sigma_{\text{FED}}$ averages 0.9 or higher throughout the sample period.

4.1 Can portfolio risk be minimized while limiting portfolio turnover?

Equally weighted portfolio Given that the MVP seems to favor a more equal distribution of assets, a natural course of action would be to analyze an equally weighted portfolio. As noted by DeMiguel et al. (2009), equally weighted portfolios still outperform many optimizing portfolio choice models and have a very low turnover. The turnover in our setup would indeed be lower than that of the FED portfolio, albeit not zero as bank entry and exit

¹⁵The random starting values are drawn from a half-normal distribution and then divided by its sum, such that they add up to 1.

would still take place. Figure 9 shows the risk-return trade-off that the equally weighted portfolio (EWP) would have in comparison with the other two. In terms of returns, the EWP performs similarly to the MVP and FED portfolio, except for the last crisis period in which it records losses. Regarding risk, the EWP has levels similar to that of the FED portfolio, albeit marginally lower. All in all, this suggests that there is an optimal level of concentration, as neither a highly concentrated nor an equally weighted portfolio are able to significantly reduce systemic risk in the same way as the MVP.

[Insert Figure 9 near here]

Limited reweighting – 1 Since a high turnover is costly for the supervisor and therefore not very desirable, we consider several alternatives involving limited reweighing that could reduce turnover. We do this by setting lower and upper boundaries to the weights banks can take, conditional on their true weights. First, we allow banks to grow/shrink by 10% and 20% of their initial weight. Second, we construct a measure of asset growth in the previous quarter and allow changes equal to either the mean or standard deviation of this growth measure. Whereas the first constraint is static in nature, the second allows for business cycle effects to determine how much reweighing can take place. To ensure the no-loss constraint does not influence the results, we run the limited reweighing scenarios with and without this requirement. However, regardless of the specification, the risk in the MVP is practically unchanged compared with the FED portfolio in all time periods.¹⁶

Limited reweighting – 2 Another way of reducing turnover would be to keep the largest 5% of banks at their current cumulative size, allowing unlimited reweighing of the remaining banks while still adhering to the no-shorting and no-loss constraint. The results of this exercise are shown in Figure 10, where we see that for most of the sample period it was possible to decrease systemic risk significantly while maintaining large banks. However, during the S&L and subprime crises periods, this MVP variant has a risk which is barely below that of the initial portfolio. This finding can be explained by the high correlation of returns for the largest banks, as seen in Figure 5b. During sudden increases in correlation between these largest banks, the high concentration of assets in a few banks will inevitably affect risk. In terms of optimal portfolio design, it seems to point to a trade-off between concentration and correlation: if the supervisor wants to keep the large banks at their current

¹⁶Since the standard deviation of asset growth was extraordinarily large in 1997Q4, this allowed the optimization to apply larger changes to the banks and therefore managed to reduce risk. A dynamic approach was also considered for the scenarios where banks are allowed to grow/shrink by 10% and 20% of their weights in the MVP in time $t - 1$; however, portfolio risk was not significantly reduced. Full results can be found in the Internet Appendix.

size, it would be necessary to ensure that correlation between them remains relatively low to avoid the Too-Many-To-Fail problem. This could, for instance, be achieved by limiting the geographical markets in which a bank can be active or the activities it can engage in, as used to be the case prior to e.g. the Riegle-Neal and the Gramm-Leach-Bliley acts.

[Insert Figure 10 near here]

Repeating this analysis on the FRD portfolios shows that there are 3 districts which are able to reduce systemic risk to a minimum in each of the three sub-samples.¹⁷ Interestingly, a common feature is that they have the lowest levels of concentration among all districts. This relationship is shown in Figure 11, where we plot the average concentration ratios against the extent to which they are able to reduce systemic risk in three time periods. The horizontal axis represents the amount by which they reduce risk and is again constructed such that 1 stands for a reduction of risk effectively to zero, and 0 indicates that no improvements in risk are possible. The Figure clearly shows the trend in consolidation, with most districts becoming more concentrated over time. As they become more concentrated, they find themselves less able to achieve low risk while maintaining their largest banks.

[Insert Figure 11 near here]

Limited reweighting – 3 Since the weight of the 5% largest banks in the three successful districts never exceeds 60%, this leads us to a final test using limited reweighting. Is it possible to reduce portfolio risk while maintaining the cumulative weight of the largest 5% of banks between 50% and 60%? The results for the FED portfolio are plotted in Figure 12. We observe that under this limited reweighting scheme, it is possible to effectively eliminate risk while maintaining similar returns. Concerning the concentration in the portfolio, the share of the largest banks always hits the lower bound of 50%. Consequently, the level of concentration in the MVP also lies close to 50% except for 2008 and 2009, with similar results for the analysis at commercial bank level as shown in the Internet Appendix. As in the previous scenarios, the high concentration in the MVP in 2008 and 2009 is likely due to increasing correlation of returns. To minimize portfolio risk, a small number of banks need to receive a higher weight, such that their concentration approaches that of the FED portfolio.

[Insert Figure 12 near here]

We have shown in Section 3 that minimizing systemic risk requires an extremely powerful and active regulator, who would have to intervene three times more than is currently the case. In

¹⁷Full results are shown in the Internet Appendix.

practice, this could only be achieved by increasing the regulator’s discretionary power. Such a proactive position was also included in the proposed SAFE Banking Act of 2012, under which a maximum bank size relative to the system would be imposed. Notwithstanding a range of limited reweighing schemes, our results indicate that in our setup, systemic risk could not be reduced while maintaining the current size of the largest banks. However, in terms of optimal portfolio design, bringing their cumulative weight down from 90% to 50% yielded a significant improvement in systemic risk.

4.2 Does the result hold for a system where only listed banks are considered?

One assumption we have consistently made is that the regulator is able to move substantial amounts of assets from large banks to very small ones. However, small banks might not be able to sustain such an increase in assets in the first place. Moreover, previous research has shown that start-up banks only behave as mature banks after their first nine years of existence (see e.g. DeYoung and Hasan, 1998). Because these small banks might not be realistic investments for the supervisor, we select only those BHCs which have publicly traded equity, using the CRSP-FRB link provided by the Federal Reserve Bank of New York (2013). We thereby also remove those banks that do not file reports after 2006 due to the increase in the reporting threshold. The selected banks are considered to be the entire portfolio in which the supervisor can invest. The results in the Internet Appendix show that removing these banks does not quantitatively or qualitatively change our results: It is possible to minimize portfolio risk by relocating assets from the largest listed banks to smaller listed ones. Moreover, in this new MVP, the actual largest 5% of banks would still receive a very low weight.¹⁸

Next, we use the CRSP-FRB link to obtain the market valuation of assets for listed BHCs by downloading equity prices as well as the number of outstanding shares from CRSP, and match them to the Call Reports. The quarterly market valuation of assets is obtained by adding the book value of the liabilities to the average market capitalization during that quarter. Compared with book value, which gives information on the past performance of a bank, the market valuation should indicate what market participants believe to be the value of the bank going forward, notwithstanding liquidity considerations during crises (Allen and Carletti, 2008). To perform our analysis on this subsample, we define returns as the quarter-to-quarter percentage changes in the market valuation of assets, and a banks’ weight as the relative share in the portfolio. As in the baseline scenario, we obtain the MVP using an eight

¹⁸A scenario in which the smallest 60% of banks were removed yielded similar results.

quarter moving window. If a merger takes place during this time, the assets of the acquired bank are added to the acquiring bank before the merger takes place for the appropriate quarters, while the acquired bank is removed from the investment universe.¹⁹ The results are reported in Figure 13.

[Insert Figure 13 near here]

Figure 13 shows spikes in the portfolio risk of the FED portfolio at the beginning of the 2000s and during the subprime crisis, whereas that of the MVP is essentially zero. While the returns of the MVP are slightly higher prior to 2005, they are similar to the FED portfolio thereafter. Regarding the weight distribution of both portfolios, we again see a steady increase in the weights of the largest banks in the FED portfolio. The MVP based on market valuation shows a higher concentration than before, as the largest banks on average are assigned 33% of the assets compared with 13% in the baseline scenario. Similar to the baseline scenario, concentration in the MVP peaks in the subprime crisis when the largest banks have a weight of 69%. However, it appears that the largest banks in the actual portfolio are still overweighted, as they see their weight reduced to an average of 4.4% in the MVP.

4.3 What is the appropriate supervisory level?

The last policy implication we discuss is the question of the appropriate supervisory level. So far, we have looked at the Federal Reserve System for the consolidated BHCs and at each Federal Reserve District separately for the unconsolidated commercial banks. However, just as with banks, if all FRDs diversify in a similar manner, the system as a whole might be prone to common shocks between districts. We therefore look at what happens to portfolio risk when we aggregate the FRD portfolios and their respective MVPs to countrywide level. The weights of the FRDs are calculated as the sum of total assets in their district divided by the sum of total assets of all districts together. While some FRDs are larger than others, overall they are relatively evenly distributed, with the weight of the largest at 20% and the smallest district representing 3% of total assets. Figure 14 shows how risk and return are aggregated up to the countrywide level, and we observe that the aggregated FRD risk and return translate almost perfectly to those of the FED portfolio. This should not come as a surprise since the latter is based on BHCs, which themselves consist mainly of commercial banks. The aggregated MVPs follow a similar pattern. The return of the aggregated MVPs

¹⁹Similar to the analyses on book data, banks that report a return below -100% or above +100% are not considered.

is slightly lower on average than that of the aggregated FRD portfolio. However, the risk of the aggregated MVPs is still practically eliminated in this scenario.

[Insert Figure 14 near here]

Figure 14 indicates that local supervision is just as able to reduce systemic risk to a minimum as the national supervisor. These results seem somewhat contradictory with regard to the new wave of macro-prudential risk oversight boards that have been established recently (the Financial Stability Oversight Council in the US and the European Systemic Risk Board for the EU). However, we have to keep in mind that our results are based on a less concentrated system, where the local supervisors are more active in shifting assets from bank to bank at their discretion. Therefore, provided risk is not concentrated in the assets of a few banks, local supervision, by reducing the risk of its own district, simultaneously reduces systemic risk. If supervisors are willing and able to minimize risk, the level at which this is achieved does not appear to be crucial. Table 4 summarizes the results of all scenarios in this section.

[Insert Table 4 near here]

5 Conclusion

The last two decades have seen a major wave of consolidation and concentration of assets in the banking industry. In the same period, the sector has experienced two major crises with a significant impact on the real economy, of which the subprime crisis had global repercussions. As a consequence of recent bailouts and government-forced sales, the sector is now even more concentrated than before the crises. In the light of moral hazard and Too-Big-To-Fail banks, we have investigated how the high concentration in the industry impacts systemic risk.

In the absence of counterfactuals, we have performed a thought experiment in which we view the supervisor as a constrained investor in a portfolio of banks. As profit maximization by banks is inherently risky, but is not perfectly correlated with that of other banks, the portfolio of the supervisor will have a better risk-return profile. By applying elements of Modern Portfolio Theory, we derive a hypothetical distribution of weights that the supervisor should have held to arrive at the minimum variance portfolio in order to give us insights into the optimal design of the banking system.

Our results consistently show that the supervisor would be able to decrease risk by reducing the level of concentration in the portfolio. Moreover, it would not be necessary

to sacrifice returns to achieve this lower risk. The minimum variance portfolio favors more traditional banks as measured by the non-interest income/total income ratio and balance sheet items such as loans, trading assets and deposits. In contrast, an equally-weighted portfolio would perform similarly to the actual, concentrated, system. These findings are robust to different starting values, time windows, covariance matrices and the exclusion of the smallest banks.

However, to achieve lower risk, the supervisor would have to adjust weights in each quarter, leading to a portfolio turnover 3 times higher than that of the real portfolio. Since this might not be possible or even desirable within the current regulatory framework, we tested several alternatives involving limited reweighing which were largely unsuccessful. These findings indicate that in times of crisis, an increase in systemic risk is unavoidable if we want to keep the concentration at current levels. Nonetheless, our analysis did show - *ceteris paribus* - that if the weight of the largest banks were kept at a sufficiently low level, systemic risk would be practically eliminated.

Our findings have two broad policy implications for bank supervisors. First, they demonstrate the importance of a strong supervisor whether at national or local level, able to manage the portfolio through interventions in the banking system. Flowing from this, the second implication is that supervisors should seriously consider the effects of concentration on systemic risk. A reduction in disparity of size could create a more competitive environment, similar to provisions of the proposed Safe, Accountable, Fair & Efficient (SAFE) Banking Act of 2012, which would limit individual banks' funding strategies to 10% of the total industry. Other measures could include imposing higher equity capital demands for the large banks as are currently being implemented, or separating investment from retail banking as proposed by Paul Volcker, and the Vickers and Liikanen reports in Europe. Forcing large banks to hold more capital could lead to a relative reduction in their size only if their assets are 'redistributed' to smaller banks in order to maintain a safer and more competitive environment. Our findings show that we should not only consider the size of each bank individually, but also consider each bank's size with respect to the whole system.

References

- Acharya, V. V. (2009). A theory of systemic risk and design of prudential bank regulation. *Journal of Financial Stability* 5(3), pp. 224–255.
- Acharya, V. V., R. F. Engle, and M. P. Richardson (2012a). Capital Shortfall: A new

- approach to ranking and regulating systemic risks. *American Economic Review* 102(3), pp. 59–64.
- Acharya, V. V., L. H. Pedersen, T. Philippon, and M. P. Richardson (2012b). Measuring systemic risk. CEPR Discussion Papers 8824, C.E.P.R. Discussion Papers.
- Acharya, V. V. and T. Yorulmazer (2007). Too many to fail – An analysis of time-inconsistency in bank closure policies. *Journal of Financial Intermediation* 16(1), pp. 1–31.
- Adrian, T. and M. K. Brunnermeier (2011). CoVaR. Working Paper 17454, National Bureau of Economic Research.
- Alegria, C. and K. Schaeck (2009). On measuring concentration in banking systems. *Finance Research Letters* 5, pp. 59–67.
- Allen, F., A. Babus, and E. Carletti (2012). Asset commonality, debt maturity and systemic risk. *Journal of Financial Economics* 104(3), pp. 519–534.
- Allen, F. and E. Carletti (2008). Mark-to-market accounting and liquidity pricing. *Journal of Accounting and Economics* 45(2-3), pp. 358–378.
- Bertay, A. C., A. Demirgüç-Kunt, and H. Huizinga (2013). Do we need big banks? Evidence on performance, strategy and market discipline. *Journal of Financial Intermediation* 22(4), pp. 532–558.
- Bisias, D., M. Flood, A. W. Lo, and S. Valavanis (2012). A survey of systemic risk analytics. *Annual Review of Financial Economics* 4(1), pp. 255–296.
- Blank, S., C. M. Buch, and K. Neugebauer (2009). Shocks at large banks and banking sector distress: The banking granular residual. *Journal of Financial Stability* 5(4), pp. 353–373.
- Boot, A. and L. Ratnovski (2012). Banking and Trading. IMF Working Paper 12/238.
- Brodie, J., I. Daubechies, C. De Mol, D. Giannone, and I. Loris (2009). Sparse and stable Markowitz portfolios. *Proceedings of the National Academy of Sciences* 106(30), pp. 12267–12272.
- Brown, C. O. and I. S. Dinç (2011). Too many to fail? Evidence of regulatory forbearance when the banking sector is weak. *Review of Financial Studies* 24(4), pp. 1378–1405.
- Brownlees, C. T. and R. F. Engle (2012). Volatility, correlation and tails for systemic risk measurement. Working paper. <<http://ssrn.com/abstract=1611229>>.

- Brunnermeier, M. K., G. Dong, and D. Palia (2012). Banks' non-interest income and systemic risk. Working paper. <<http://ssrn.com/abstract=1786738>>.
- Carvalho, V. and X. Gabaix (2013). The great diversification and its undoing. *American Economic Review* 103(5), pp. 1697–1727.
- De Jonghe, O. (2010). Back to the basics in banking? A micro-analysis of banking system stability. *Journal of Financial Intermediation* 19(3), pp. 387–417.
- Degryse, H. and G. Nguyen (2007). Interbank exposures: An empirical examination of contagion risk in the Belgian banking system. *International Journal of Central Banking* 3(2), pp. 123–171.
- DeMiguel, V., L. Garlappi, and R. Uppal (2009). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies* 22(5), pp. 1915–1953.
- Demirgüç-Kunt, A. and H. Huizinga (2013). Are banks too big to fail or too big to save? International evidence from equity prices and CDS spreads. *Journal of Banking & Finance* 37(3), pp. 875–894.
- DeYoung, R. and I. Hasan (1998). The performance of de novo commercial banks: A profit efficiency approach. *Journal of Banking & Finance* 22, pp. 565–587.
- DeYoung, R. and K. P. Roland (2001). Product mix and earnings volatility at commercial banks: Evidence from a degree of total leverage model. *Journal of Financial Intermediation* 10(1), 54–84.
- DeYoung, R. and G. Torna (2013). Nontraditional banking activities and bank failures during the financial crisis. *Journal of Financial Intermediation* 22(3), pp. 397–421.
- Engle, R. F., E. Jondeau, and M. Rockinger (2014). Systemic risk in Europe. *Review of Finance*.
- Fahlenbrach, R., R. Prilmeier, and R. M. Stulz (2012). This time is the same: Using bank performance in 1998 to explain bank performance during the recent financial crisis. *The Journal of Finance* 67(6), pp. 2139–2185.
- Federal Reserve Bank of New York (2013). CRSP-FRB Link.
- Furfine, C. H. (1999). The microstructure of the Federal Funds market. *Financial Markets, Institutions & Instruments* 8(5), pp. 24–44.

- Furfine, C. H. (2003). Interbank exposures: Quantifying the risk of contagion. *Journal of Money, Credit and Banking* 35(1), pp. 111–128.
- Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica* 79(3), pp. 733–772.
- Heckman, J. J., H. Ichimura, and P. Todd (1998). Matching as an econometric evaluation estimator. *Review of Economic Studies* 65(2), pp. 261–94.
- Huang, R. and L. Ratnovski (2011). The dark side of bank wholesale funding. *Journal of Financial Intermediation* 20(2), pp. 248–263.
- Hughes, J. P. and L. J. Mester (2013). Who said large banks don’t experience scale economies? Evidence from a risk-return-driven cost function. *Journal of Financial Intermediation* 22(4), pp. 559–585.
- Ibragimov, R., D. Jaffee, and J. Walden (2011). Diversification disasters. *Journal of Financial Economics* 99(2), pp. 333–348.
- Janicki, H. P. and E. S. Prescott (2006). Changes in the size distribution of U.S. banks: 1960-2005. *Economic Quarterly*, pp. 291–316.
- Jobson, J. D. and B. Korkie (1980). Estimation for markowitz efficient portfolios. *Journal of the American Statistical Association* 75(371), pp. 544–554.
- Ledoit, O. and M. Wolf (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance* 10(5), pp. 603–621.
- Ledoit, O. and M. Wolf (2004). Honey, I shrunk the sample covariance matrix. *The Journal of Portfolio Management* 30(4), pp. 110–119.
- Loutskina, E. and P. E. Strahan (2011). Informed and uninformed investment in housing: The downside of diversification. *Review of Financial Studies* 24(5), pp. 1447–1480.
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance* 7(1), pp. 77–91.
- Markowitz, H. (1959). *Portfolio Selection: Efficient Diversification of Investments*. New York: John Wiley & Sons.
- Office of the Special Inspector General for the Troubled Asset Relief Program (2013, October 29). Quarterly report to congress. Technical report, US Treasury.

- Rajan, R. G. (2005). Has financial development made the world riskier? *Proceedings Aug*, pp. 313–369.
- Stiroh, K. J. (2004). Diversification in banking: Is noninterest income the answer? *Journal of Money, Credit and Banking* 36(5), pp. 853–882.
- Stiroh, K. J. and A. Rumble (2006). The dark side of diversification: The case of US financial holding companies. *Journal of Banking & Finance* 30, pp. 2131–2161.
- van Lelyveld, I. and F. Liedorp (2006). Interbank contagion in the Dutch banking sector: A sensitivity analysis. *International Journal of Central Banking* 2(2), pp. 99–133.
- Wagner, W. (2008). The homogenization of the financial system and financial crises. *Journal of Financial Intermediation* 17(3), pp. 330–356.
- Wheelock, D. C. and P. W. Wilson (2012). Do large banks have lower costs? New estimates of returns to scale for U.S. banks. *Journal of Money, Credit and Banking* 44(1), pp. 171–199.

Figures and Tables

Figure 1: The supervisory view: Markowitz efficient frontier, systemic risk and bank size

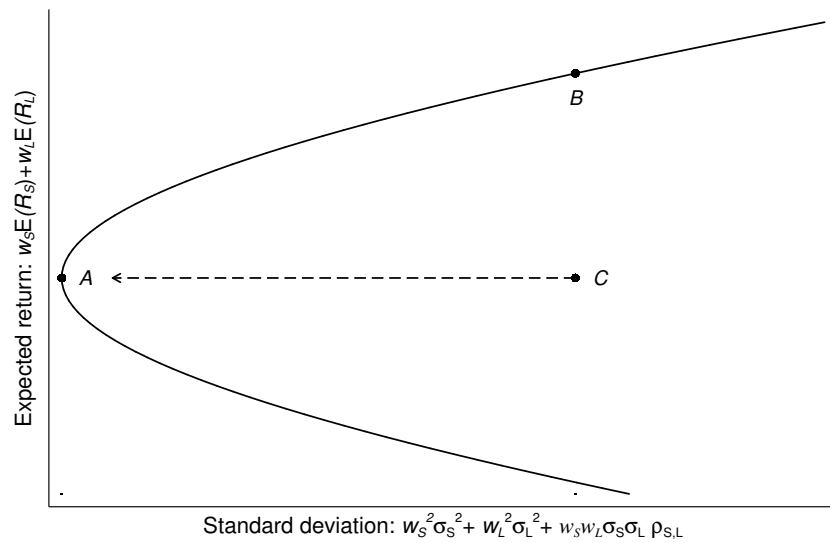



Figure 2: The Federal Reserve districts

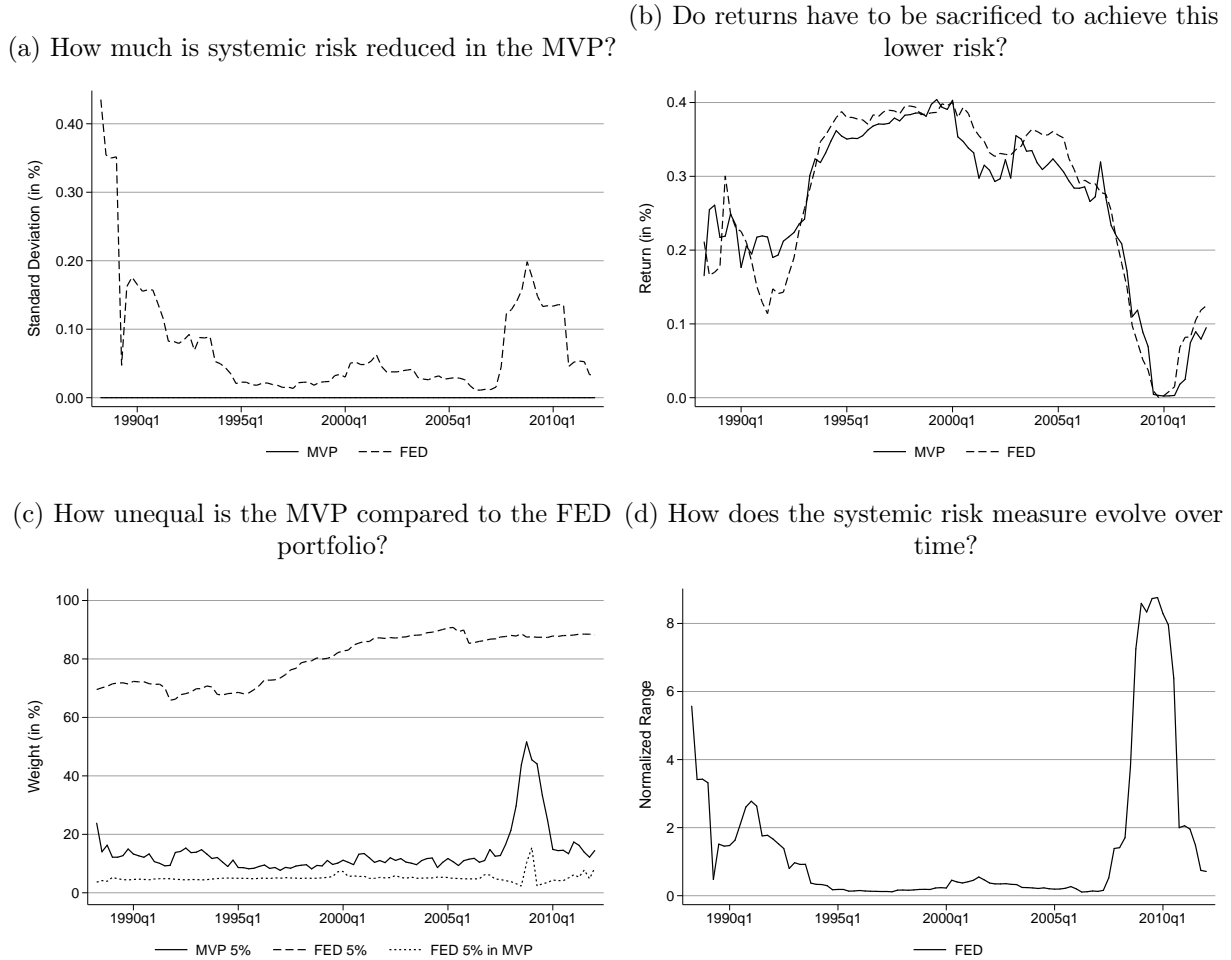


Figure 3: Are large banks more risky?



The figure shows density plots of the average 8 quarter variance of return-on-assets and average pairwise covariances of the smallest 95% and the largest 5% of Bank Holding Companies between 1988Q2 and 2012Q1. Due to long tails for both large and small banks, the variance is truncated at its 90th percentile as it is strictly positive, while the covariance is truncated at its 5st and 95th percentile for graphical purposes.


Figure 4: What role does inequality play in the risk/return trade-off?




The figure shows the comparison between the FED portfolio and the hypothetical MVP: panel 4a and 4b display the difference in the risk and return of each portfolio. Panel 4c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP. Panel 4d shows the demeaned risk measure from Panel 4a to allow for a comparison over time. Instead of using the standard deviation, which is a reflection of the mean, we show the normalized range defined as $(\max - \min) / \text{median}$ and is independent of the mean return.

Figure 5: What are the other features of the Minimum Variance Portfolio?

(a) Is there a large share for low-risk banks?

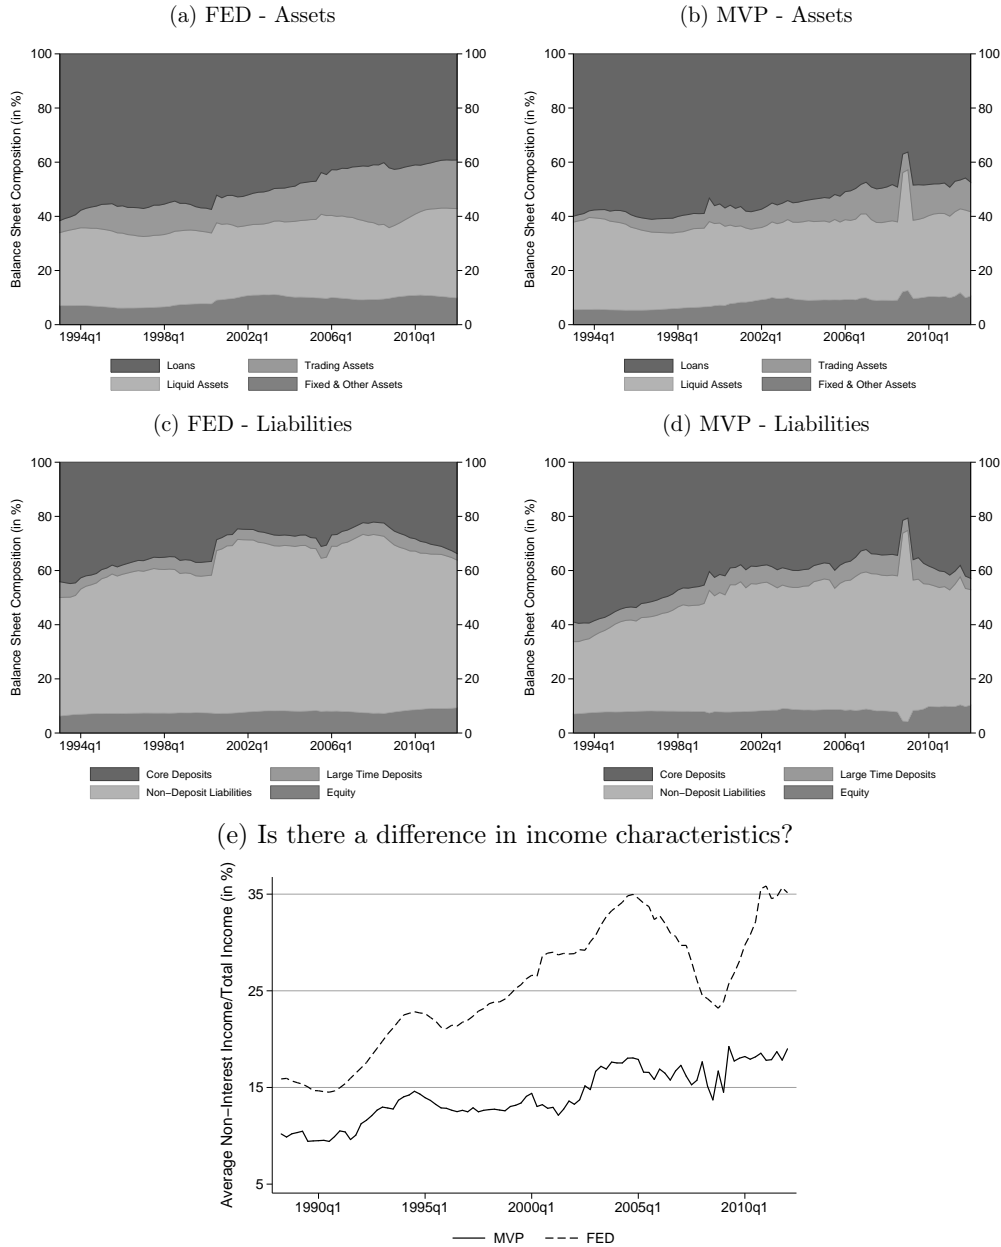


(b) Is there on average less correlation among banks?




The figure shows the average bank level standard deviation of the return-on-assets, as well as the average pairwise correlation between these returns for the largest 5% of banks in both MVP and FED portfolio.

Figure 6: What happens to the intermediary role of banks in a safer banking system?



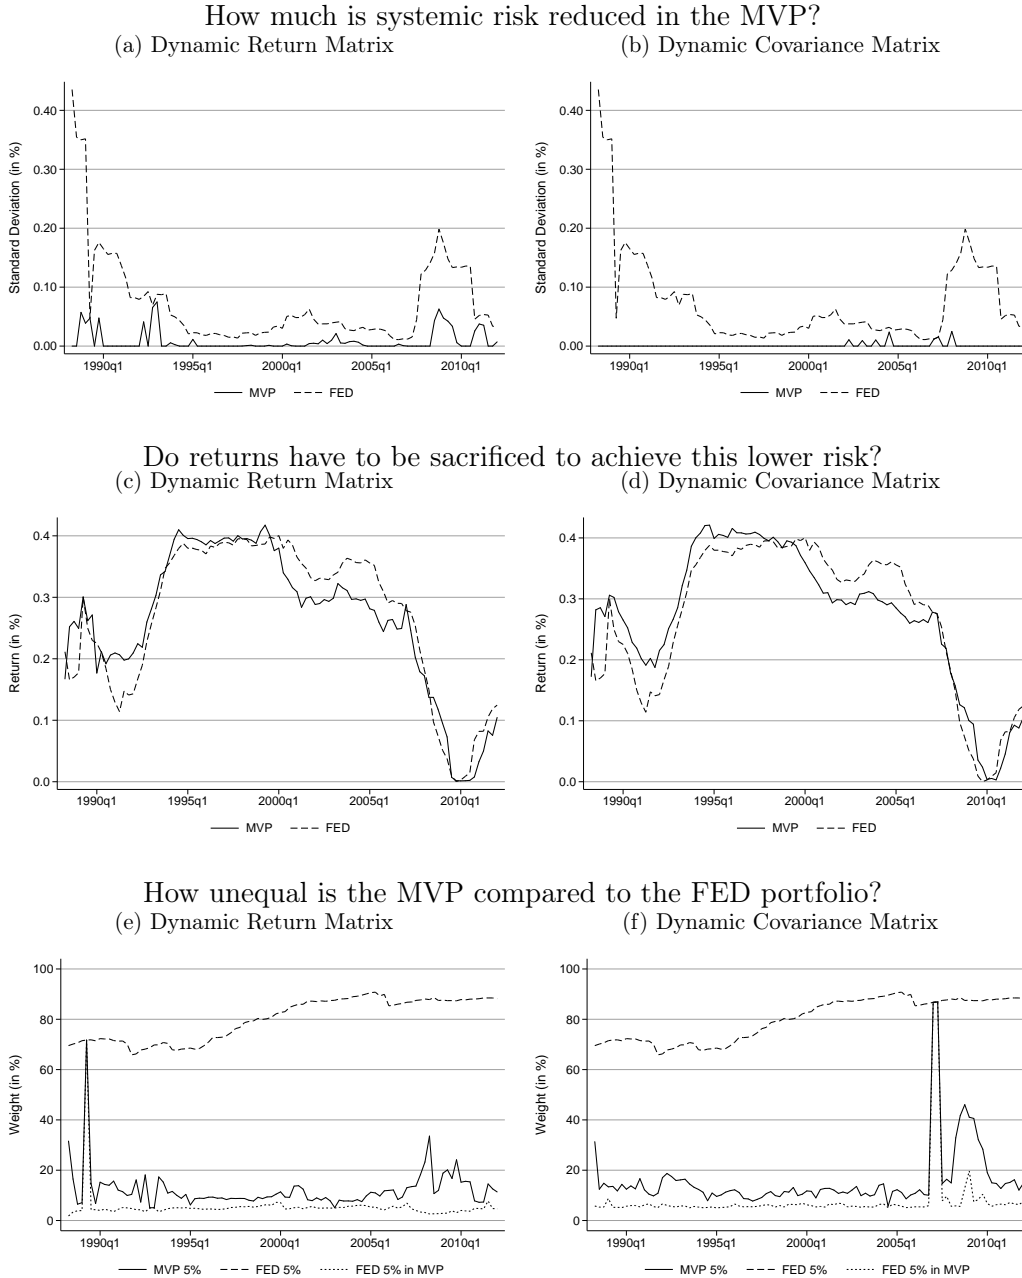
The figure compares the intermediary role of the FED portfolio and MVP, by showing weighted average industry level balance sheets in panels 6a-6d and the different weighted non-interest income/total income ratios in panel 6e. To construct the balance sheet we use the following data series: Loans (BHCK2122); Trading Assets up to 1994 (BHCK2146); Trading Assets after 1994 (BHCK3545); Liquid Assets up to 1994 (BHCK0081 + BHCK0395 + BHCK0397 + BHCK3365 + BHCK0390); Liquid Assets after 1994 (BHCK0081 + BHCK0395 + BHCK0397 + BHCK3365 + BHCK1754 + BHCK1773); Fixed & Other Assets (BHCK2145 + BHCK3163 + BHCK2160); Core Deposits (BHCB3187 + BHCB2389 + BHCB6648 + BHCB2210 + BHOD3187 + BHOD2389 + BHOD6648); Large Time Deposits (BHCB2604 + BHOD2604); Non-Deposit Liabilities (BHCK2948 - Core Deposits - Large Time Deposits); Equity (BHCK2948 + BHCK3210). Non-interest income (BHCK4079); Total income (BHCK4107).

Figure 7: How much more intervention would be required for the MVP?



The figure shows the difference in reweighing between the FED and the MVP by plotting the turnover of each portfolio.


Figure 8: Can portfolio risk be minimized when taking into account a changing covariance matrix due to weight changes?




The figure presents two robustness tests, in which the return and/or covariance matrices are dynamically updated. In the first robustness test (Dynamic Return Matrix), we adjust only the return matrix and then estimate a new sample covariance matrix. In the second test (Dynamic Covariance Matrix), we adjust both the return matrix and the covariance matrix with which the portfolio risk will be minimized. More information on how we perform these tests can be found in the Internet Appendix. Panel 8a and 8a display the difference in the risk each portfolio, while panel 8c and 8d display the difference in return. Panel 8e and 8f show how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Figure 9: What is the risk/return profile of an equally weighted portfolio?

(a) Can we reduce system risk in an equally weighted portfolio?




(b) Is an equally weighted system more profitable?



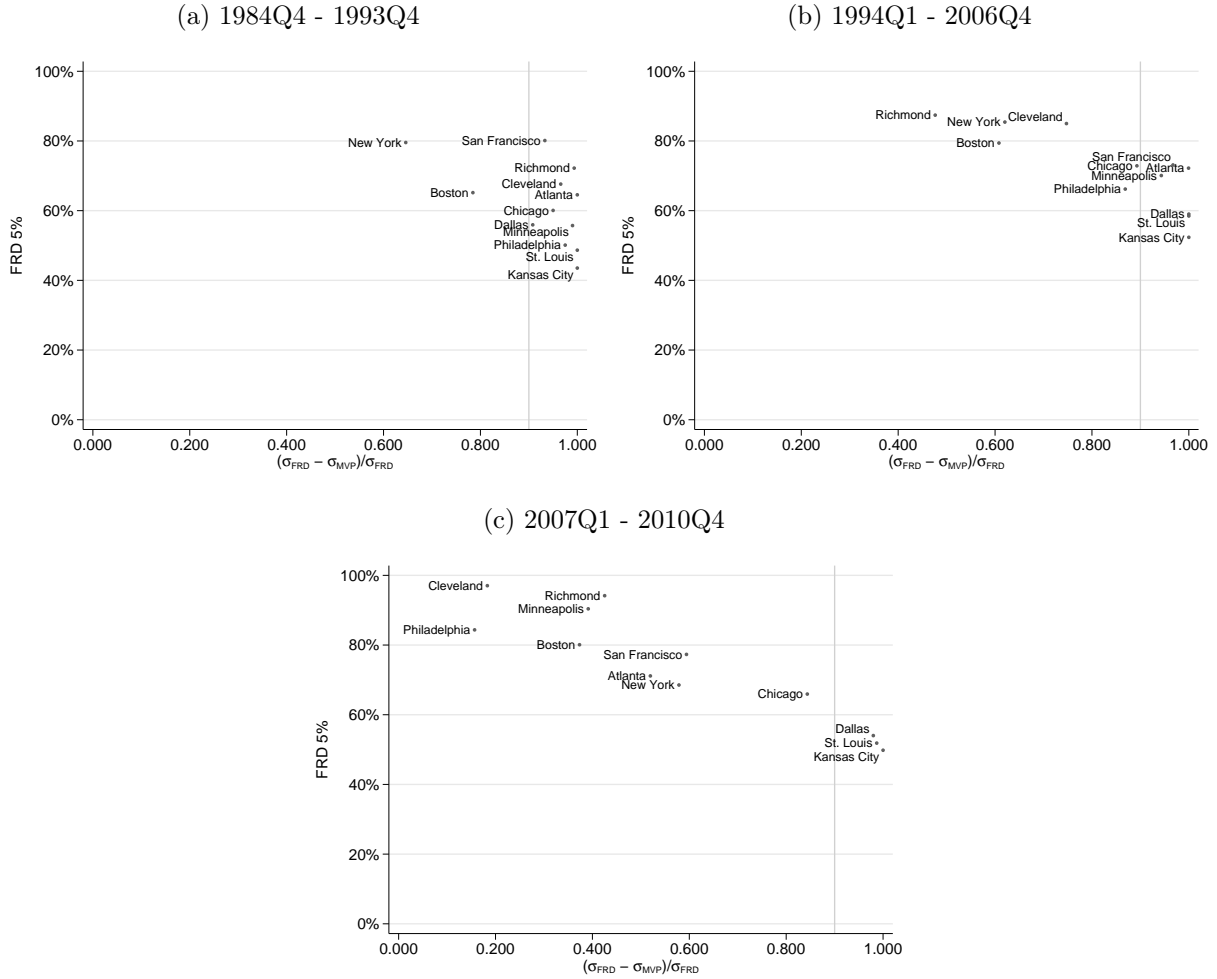
The figure presents results of the risk and return in an equally weighted portfolio, compared to the MVP and FED portfolio.

Figure 10: How much can we lower systemic risk when we keep the largest banks at their actual size?



The figure displays how well the portfolio risk can be minimized while keeping the largest 5% of banks at their actual size and reweighing the remaining 95%.


Figure 11: Can inequality explain why some FRDs can reduce systemic risk while keeping the largest banks at their actual size?



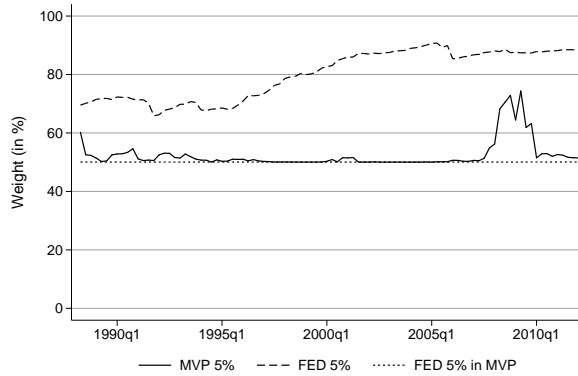
The figure shows the relationship over time between the average weight of the largest 5% of banks and the ability of the FRDs to reduce portfolio risk when the largest banks are allowed to keep their initial weight. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements were possible. The threshold for successfully minimizing portfolio risk is set at 0.9. The weights of the largest 5% of banks are averaged over each time period.

Figure 12: By how much is the risk minimized when the largest banks hold between 50% and 60% of assets?

(a) How much is systemic risk reduced in the MVP? (b) Do returns have to be sacrificed to achieve this lower risk?




(c) How large are the largest banks in the MVP?




The figure shows the risk and return characteristics of an MVP where the largest 5% of banks in the FED portfolio are kept between 50% and 60% of assets, and how the weights in this MVP are distributed.

Figure 13: What role does inequality play in the risk/return trade-off when taking into account the market valuation of assets?

- (a) How much is systemic risk reduced in the MVP? (b) Do returns have to be sacrificed to achieve this lower risk?




- (c) How unequal is the MVP compared to the FED portfolio?



The figure shows the comparison between the FED portfolio and the hypothetical MVP: panel 13a and 13b display the difference in the risk and return of each portfolio. Panel 13c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Figure 14: Can local supervision make the system less risky?



The figure displays the implications of aggregating separate FRD portfolios and their MVPs to countrywide level. Using the weights that the FRDs have in the total FED portfolio, the weighted average FRD/MVP returns are calculated as well as their standard deviations.

Table 1: Summary Statistics

Panel A: Bank Holding Companies

	Percentiles						
	Min.	5 th	25 th	50 th	75 th	95 th	Max.
Net Income (in \$ million)	-15132.20	-0.72	0.64	1.30	3.09	44.0	6414.61
Total Assets (in \$ million)	5.33	156.32	284.62	500.37	1184.04	14456.24	2115728.50
Return on Assets (in %)	-39.29	-0.14	0.18	0.29	0.39	0.58	82.81
	Mean	Std. Dev.		Min.	Max.		
Number of Banks	1502	387		964	2333		

Panel B: Commercial Banks

	Percentiles						
	Min.	5 th	25 th	50 th	75 th	95 th	Max.
Net Income (in million \$)	-11168.98	-0.20	0.08	0.24	0.58	2.94	4682.32
Total Assets (in million \$)	0.14	15.91	39.48	79.39	171.95	896.35	1594746.30
Return on Assets (in %)	-78.23	-0.34	0.19	0.33	0.46	0.72	90.85
	Mean	Std. Dev.		Min.	Max.		
Number of Banks	10115	2655		6477	14474		

The table presents summary statistics for Net Income (bhck/riad4340), Total Assets (bhck/rcfd2170) and the Return-on-Assets of the banks in the analysis. Panel A displays these descriptives for Bank Holding Companies between 1986Q3 and 2012Q1, while Panel B summarizes them for Commercial Banks between 1984Q1 and 2010Q4. Due to the highly skewed distributions, we summarize the data according to their percentiles as well as the minimum and maximum values.

Table 2: Changes in the return and its covariance matrix

		ROA		Variance ROA		Covariance ROA	
		Pre-Jump	Post-Jump	Pre-Jump	Post-Jump	Pre-Jump	Post-Jump
<u>Positive Jump</u>	Value	0.341	0.259	0.023	0.111	0.002	0.005
	Average Difference	-0.082		0.088		0.004	
	Percentage Difference	-23.97%		379.70%		233.70%	
	P-value T-test	0.000		0.007		0.010	
<u>Negative Jump</u>	Value	-0.142	0.213	0.281	3.249	0.007	-0.015
	Average Difference	0.355		2.967		-0.022	
	Percentage Difference	249.52%		1054.64%		-324.52%	
	P-value T-test	0.110		0.359		0.293	

The table shows the average differences in ROA, its variance and its average pairwise covariance pre- and post-jump for both negative jumps and positive jumps using the procedure described in the text and the Internet Appendix.

Table 3: Robustness tests

		Risk minimized without large banks?		
Issue		Robustness test	BHC	COMM
(1)	Independence of covariance matrix	Dynamic return matrix	✓	n.a.
(2)	Independence of covariance matrix	Dynamic covariance matrix	✓	n.a.
(3)	Noise in sample covariance matrix	Shrinkage estimator	✗	0/12
(4)	Multiple minima	Equal starting weights	✓	12/12
(5)	Multiple minima	Randomized starting weights	✓	n.a.
(6)	Length of rolling window	Analysis on 16 quarters	✓	12/12

The table presents a summary of the three robustness tests that were performed, indicating whether portfolio risk was successfully minimized while keeping the levels of concentration of the largest 5% of banks low. To see if a test is successful we define the ratio $(\sigma_{\text{FED}} - \sigma_{\text{MVP}})/\sigma_{\text{FED}}$, which measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements were possible. A test is regarded as successful when this ratio averages 0.9 or higher throughout the sample period, and when the level of concentration of the largest 5% of banks in the MVP is below 50%. The second, fifth and sixth tests are only performed on BHC data due to their computationally intensive nature.

Table 4: Extensions and policy implications tests

Issue	Policy test	Risk minimized?	
		BHC	COMM
(1) High turnover	Equally Weighted Portfolio	✗	0/12
(2) High turnover	Reweighting limited to 10% or 20% of assets ^a	✗	0/12
(3) High turnover	Reweighting limited to mean/std. dev. of growth ^a	✗	0/12
(4) High turnover	Largest 5% keep their weight ^a	✗	5/12
(5) High turnover	Largest 5% are weighted between 50% and 60%	✓	12/12
(6) Small/DeNovo banks	Only BHCs with publicly traded equity – book value	✓	n.a.
(7) Small/DeNovo banks	Only BHCs with publicly traded equity – market value	✓	n.a.
(8) Level of supervision	Aggregate District MVPs to national level	✓	n.a.

The table presents a summary of the seven policy tests that were performed and how well they worked in reducing systemic risk on FED and FRD level. To see if a test is successful we define the ratio $(\sigma_{\text{FED}} - \sigma_{\text{MVP}})/\sigma_{\text{FED}}$, which measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements were possible. A test is regarded as successful when this ratio averages 0.9 or higher throughout the sample period.

^a To avoid the no portfolio loss constraint driving these results, we also performed the tests without the no-loss constraint. The results however do not change.

Internet Appendix to

“Carrying the (Paper) Burden: A Portfolio View of Systemic Risk and Optimal Bank Size”

This Internet Appendix contains results that were omitted from the body of the paper for brevity. In Section A, we report the results of the baseline analysis for each of the separate Federal Reserve Districts (FRDs). Section B shows results for several robustness tests, while Section C shows results for the extensions and policy implications tests.

A Extensive Results: Baseline Specification

Tables A.1 and A.2 show the differences in the risk-return trade-off and distribution of weights for the FRD portfolios and their respective MVPs. We show the results for the entire sample (1984Q4 - 2010Q4), as well as different subsamples (1984Q4 - 1993Q4, 1994Q1 - 2006Q4 and 2007Q1 - 2010Q4). The numbers reported are averages during the relevant time span and Kolmogorov-Smirnov tests are performed to see if the distributions in the FRDs and MVPs are different from each other. In the second panel of Table A.1, we report the relative difference in portfolio risk between the MVPs and FRD portfolios. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements were possible.

The tables show mixed results for the differences in return over the whole sample: some FRDs outperform their MVPs, whereas others exhibit lower returns. One interesting fact is that in the boom period of 1994Q1 - 2006Q4, we find that all FRDs outperform their MVP counterparts in terms of returns. Regarding size disparity, the MVPs constantly have a much lower level of concentration compared to their FRD portfolios, their difference ranging on average between 44% and 78% throughout the sample period.

Table A.1: Portfolio Optimization Baseline – Risk-Return Trade-Off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.032***	0.131***	-0.037**	0.026
New York	0.044***	0.167***	-0.016***	-0.045*
Philadelphia	-0.105***	0.036***	-0.271***	0.107**
Cleveland	0.011***	0.036***	-0.026***	0.073**
Richmond	-0.023***	0.036***	-0.050***	-0.075*
Atlanta	-0.028***	-0.042**	-0.047***	0.066***
Chicago	0.018***	0.042***	-0.017**	0.079***
St. Louis	-0.018***	-0.020***	-0.030***	0.023
Minneapolis	-0.123***	-0.091***	-0.150***	-0.113**
Kansas City	-0.047***	-0.063***	-0.044***	-0.019
Dallas	0.012***	0.070**	-0.019***	-0.019
San Francisco	-0.081***	-0.075**	-0.102***	-0.030*
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	1.000***	1.000***	1.000***	1.000***
New York	1.000***	1.000***	1.000***	1.000***
Philadelphia	1.000***	1.000***	1.000***	1.000***
Cleveland	1.000***	1.000***	1.000***	1.000***
Richmond	1.000***	1.000***	1.000***	1.000***
Atlanta	1.000***	1.000***	1.000***	1.000***
Chicago	1.000***	1.000***	1.000***	1.000***
St. Louis	1.000***	1.000***	1.000***	1.000***
Minneapolis	1.000***	1.000***	1.000***	1.000***
Kansas City	1.000***	1.000***	1.000***	1.000***
Dallas	1.000***	1.000***	1.000***	1.000***
San Francisco	1.000***	1.000***	1.000***	1.000***

This table shows the results of the baseline specification for the analysis on commercial bank level. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table A.2: Portfolio Optimization Baseline – Weights

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<i>FRD 5% in MVP - FRD 5%</i>				
Boston	-68.8***	-59.7***	-74.0***	-73.3***
New York	-76.1***	-75.3***	-80.4***	-64.2***
Philadelphia	-58.0***	-45.2***	-60.6***	-79.3***
Cleveland	-75.6***	-62.3***	-79.8***	-92.6***
Richmond	-78.0***	-66.8***	-82.4***	-89.3***
Atlanta	-64.5***	-59.8***	-67.2***	-66.3***
Chicago	-62.2***	-55.1***	-67.5***	-61.7***
St. Louis	-48.5***	-43.6***	-52.4***	-47.1***
Minneapolis	-62.4***	-50.5***	-63.9***	-85.0***
Kansas City	-43.8***	-38.3***	-47.3***	-45.2***
Dallas	-52.3***	-51.3***	-54.0***	-49.1***
San Francisco	-71.1***	-75.2***	-67.8***	-72.3***
<i>MVP 5% - FRD 5%</i>				
Boston	-54.2***	-42.2***	-63.4***	-52.0***
New York	-65.7***	-61.4***	-73.5***	-49.9***
Philadelphia	-48.5***	-32.5***	-55.1***	-64.1***
Cleveland	-68.1***	-52.6***	-74.9***	-82.1***
Richmond	-68.0***	-53.7***	-74.8***	-79.1***
Atlanta	-55.3***	-48.0***	-62.2***	-49.3***
Chicago	-55.8***	-48.2***	-62.4***	-52.0***
St. Louis	-38.4***	-33.7***	-42.7***	-35.1***
Minneapolis	-52.6***	-40.2***	-55.2***	-73.0***
Kansas City	-35.4***	-31.9***	-37.9***	-35.7***
Dallas	-44.0***	-42.3***	-46.9***	-38.5***
San Francisco	-62.4***	-65.3***	-62.6***	-55.2***

This table shows the results of the baseline specification for the analysis on commercial bank level. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

B Extensive Results: Robustness Tests

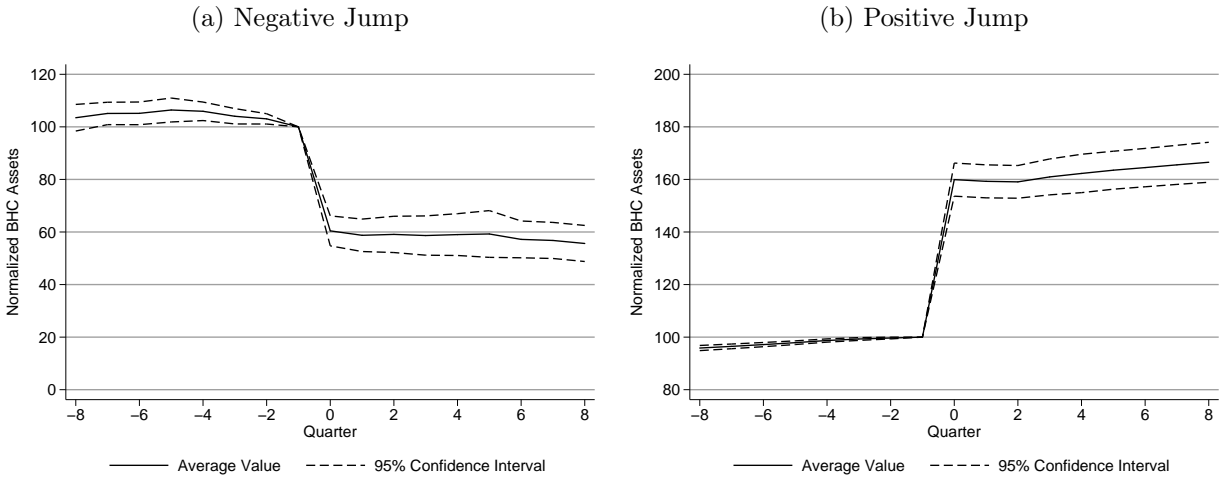
In this section, we revisit several assumptions in the baseline specification and, where possible, relax them. In section B.1, we first explore the assumption that the sample covariance matrix does not change when banks face a large increase/decrease in total assets. Second, in section B.2, we use an alternative estimator of the covariance matrix which reduces the noise in a small T and large N setting employed in our baseline specification. Third, section B.3 examines whether the baseline findings are a result of the starting weights, before confirming in section B.4 that the results also hold for a different length window on which the covariance matrix is estimated.

B.1 Identifying Jumps in Assets of BHCs for Portfolio Optimization under a Dynamic Covariance and Return Matrix

In the baseline specification, the sample covariance matrix was assumed to be constant, such that a large change in the size of the bank would not have changed its return structure. Although we have shown that large and small banks share a common support in this structure, here we identify cases in which banks have seen a large increase/decrease in their total assets and analyze how the elements of its covariance matrix changed.

A jump is defined as an increase/decrease of bank assets of 25% or greater from one quarter to the next, provided that the preceding and following 8 quarters did not show jumps larger than 10% in each of the quarters, nor a cumulative change in the preceding and following 8 quarters of 25%. These last two conditions are imposed to make sure that bank size before and after the jump was relatively stable and that the change in the elements of the covariance matrix can be chiefly attributed to the one-time jump. We find 15 negative jumps and 287 positive jumps during the sample period. Figure B.1.1 depicts the movement in assets before and after the jump, where the total assets are normalized to 100 at the quarter prior to the jump. Banks that experience a negative jump lose on average 40% of their assets, whereas banks experiencing a positive jump gain on average 60%.

Figure B.1.1: Jump in Bank Assets



The figure displays the average jump in total bank assets for the selected bank holding companies, as well as its 95% confidence interval. The negative jump in assets is based on 15 BHCs, whereas the positive jump is based on 287 BHCs. Jumps were defined as a decrease (increase) of bank assets of 25% or greater from one quarter to next, provided that the preceding and following 8 quarters did not show jumps larger than 10% in each of the quarters, nor a cumulative change in the preceding and following 8 quarters of 25%. Assets are normalized to 100 in the quarter before the jump.

Table B.1.1: Changes in the Return and Covariance Matrices due to Jumps

		ROA		Variance ROA		Covariance ROA	
		Pre-Jump	Post-Jump	Pre-Jump	Post-Jump	Pre-Jump	Post-Jump
<u>Positive Jump</u>	Value	0.341	0.259	0.023	0.111	0.002	0.005
	Average Difference		-0.082		0.088		0.004
	Percentage Difference		-23.97%		379.70%		233.70%
	P-value T-test		0.000		0.007		0.010
<u>Negative Jump</u>	Value	-0.142	0.213	0.281	3.249	0.007	-0.015
	Average Difference		0.355		2.967		-0.022
	Percentage Difference		249.52%		1054.64%		-324.52%
	P-value T-test		0.110		0.359		0.293

The table shows the average differences in ROA, its variance and its average pairwise covariance pre- and post-jump for both negative jumps and positive jumps.

The average changes to the return and covariance matrices are displayed in Table B.1.1. T-tests show that banks experiencing a positive jump in assets have a statistically significant lower return-on-assets 8 quarters following the jump and a higher average variance and covariance of these returns. On the other hand, banks experiencing a negative jump do not see statistically significant changes in their average variance and covariance, and only a marginally significant increase in their return-on-assets after the jump.

Using the statistically significant changes for positive and negative shocks, we interpolate the effect to the sample covariance matrix and the return matrix, such that these matrices change dynamically with the weights received by the banks. As the change in return-on-assets for the banks experiencing a negative jump is marginally significant at 11%, we also regard this change as significant. During the iterative process, the covariance matrix is updated based on the proposed weights and the actual weights. For example, if a bank has a weight of 5% and the outcome of the iteration is that it should have a weight of 7.5%, i.e. an increase of 50%, its return-on-assets would decrease by $\frac{0.50}{0.60} \times -23.97\% = -19.975\%$, its variance would increase by $\frac{0.50}{0.60} \times 379.70\% = 316.416\%$ and every covariance term would increase by $\frac{0.50}{0.60} \times 233.70\% = 194.75\%$. If, based on the new return and covariance matrices, in the second iteration the proposed weight is 4%, i.e. a 20% decrease from its original size, only its return-on-assets would increase by $\frac{0.20}{0.40} \times 249.52\% = 124.76\%$, while nothing would happen to its original variance and covariance terms as these changes are not statistically significant. Since we do not have data on jumps beyond -40% and +60%, any proposed weight change beyond these thresholds will lead to an adjustment of the return and covariance matrices as if they were -40% or +60%.

In the first robustness test, we adjust only the return matrix and then estimate a new sample covariance matrix. In the second test, we adjust both the return matrix and the covariance matrix with which the portfolio risk will be minimized. The results are shown in the paper in Figure 8.


B.2 Portfolio Optimization using the Ledoit and Wolf (2003, 2004) Covariance Matrix

The sample covariance matrix with which portfolio optimization is applied, relies on a large N and small T setting, which has been shown to introduce a considerable amount of noise in this estimation. Ledoit and Wolf (2003, 2004) proposed a shrinkage based estimator of the covariance matrix to reduce the noise. We use their proposed estimator of the covariance matrix instead of its sample equivalent and re-run the analysis. Figure B.2.1 show the results for the analysis on BHC level, and Tables B.2.1 and B.2.2 do the same for the analysis on commercial bank level.


Overall, we find that the standard deviation of the MVP is now higher compared to the baseline specification, although still lower than that of the FED portfolio. The ratio $(\sigma_{\text{FED}} - \sigma_{\text{MVP}})/\sigma_{\text{FED}}$ is on average only 0.54, indicating that 54% of the portfolio risk was minimized. For the commercial banks, this ratio ranges between 0.40 and 0.88. The portfolio returns, meanwhile, are at a similar level. The levels of concentration in the MVP are basically unchanged, as the largest banks still have a weight of on average 10%, while the largest banks of the FED portfolio have a weight in the MVP of 5%. For the FRD portfolios, on average 60% of portfolio risk is eliminated, and the return of the MVPs is higher. The weight of the largest 5% of banks in the MVP is on average 47% lower than in the FRD portfolio, while the largest banks in the FRD are reduced by 70%. Both of these results are similar to the baseline specification.

Figure B.2.1: Portfolio Optimization using the Ledoit and Wolf (2003, 2004) Covariance Matrix

(a) How much is systemic risk reduced in the MVP? (b) Do returns have to be sacrificed to achieve this lower risk?



(c) How unequal is the MVP compared to the FED portfolio?



The figure shows the comparison between the FED portfolio and the hypothetical MVP using the Ledoit and Wolf (2003, 2004) shrinkage estimator of the covariance matrix: panel B.2.1a and B.2.1b display the difference in the risk and return of each portfolio. Panel B.2.1c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Table B.2.1: Portfolio Optimization using the Ledoit and Wolf (2003, 2004) Covariance Matrix - Risk-return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.024***	0.124***	-0.055***	0.047*
New York	0.089***	0.234***	0.008***	0.016**
Philadelphia	-0.103***	0.032***	-0.288***	0.189**
Cleveland	-0.016***	0.060***	-0.079***	0.017**
Richmond	0.020***	0.090***	-0.028***	0.012*
Atlanta	0.032***	0.078***	-0.036***	0.149***
Chicago	0.035***	0.073***	-0.011*	0.098***
St. Louis	0.010***	0.046***	-0.035***	0.073***
Minneapolis	-0.099***	-0.000***	-0.182***	-0.055*
Kansas City	0.022***	0.083***	-0.023***	0.029*
Dallas	0.100***	0.262***	0.006***	0.031**
San Francisco	-0.003***	0.105***	-0.100***	0.064***
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	0.847***	0.787***	0.885***	0.861***
New York	0.801***	0.737***	0.806***	0.929***
Philadelphia	0.880***	0.810***	0.902***	0.970***
Cleveland	0.729***	0.664***	0.752***	0.808***
Richmond	0.681***	0.525***	0.778***	0.730***
Atlanta	0.402***	0.358***	0.342***	0.699***
Chicago	0.486***	0.495***	0.429***	0.650***
St. Louis	0.394***	0.336***	0.351***	0.666***
Minneapolis	0.587***	0.569***	0.546***	0.762***
Kansas City	0.316***	0.380***	0.224***	0.465***
Dallas	0.501***	0.634***	0.385***	0.572***
San Francisco	0.712***	0.687***	0.726***	0.727***

This table shows the results of the Ledoit and Wolf (2003, 2004) shrinkage based estimator of the covariance matrix for the analysis on commercial bank level. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table B.2.2: Portfolio Optimization using the Ledoit and Wolf (2003, 2004) Covariance Matrix - Weights

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u>FRD 5% in MVP - FRD 5%</u>				
Boston	-63.7***	-53.1***	-66.4***	-79.2***
New York	-74.4***	-70.3***	-79.9***	-65.9***
Philadelphia	-58.9***	-43.9***	-62.3***	-82.6***
Cleveland	-74.8***	-58.6***	-80.4***	-94.0***
Richmond	-76.0***	-60.8***	-82.3***	-90.4***
Atlanta	-61.1***	-56.6***	-62.5***	-67.1***
Chicago	-59.9***	-48.4***	-67.5***	-61.8***
St. Louis	-47.1***	-39.5***	-52.4***	-47.1***
Minneapolis	-61.4***	-47.8***	-64.0***	-84.9***
Kansas City	-42.1***	-36.3***	-45.4***	-45.0***
Dallas	-50.8***	-49.5***	-52.5***	-48.3***
San Francisco	-67.1***	-65.7***	-66.8***	-71.1***
<u>MVP 5% - FRD 5%</u>				
Boston	-26.0***	-9.7***	-36.2***	-30.2***
New York	-47.9***	-29.9***	-65.4***	-32.9***
Philadelphia	-30.7***	-7.4***	-44.3***	-40.3***
Cleveland	-60.2***	-40.5***	-68.3***	-79.5***
Richmond	-62.0***	-44.5***	-73.9***	-63.4***
Atlanta	-53.9***	-47.2***	-59.8***	-49.9***
Chicago	-57.8***	-46.8***	-65.9***	-57.0***
St. Louis	-41.5***	-31.5***	-48.9***	-40.6***
Minneapolis	-53.8***	-38.2***	-58.4***	-75.1***
Kansas City	-37.4***	-29.1***	-42.5***	-40.1***
Dallas	-41.2***	-33.8***	-46.2***	-42.3***
San Francisco	-50.9***	-49.4***	-59.2***	-27.5***

This table shows the results of the Ledoit and Wolf (2003, 2004) shrinkage based estimator of the covariance matrix for the analysis on commercial bank level. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

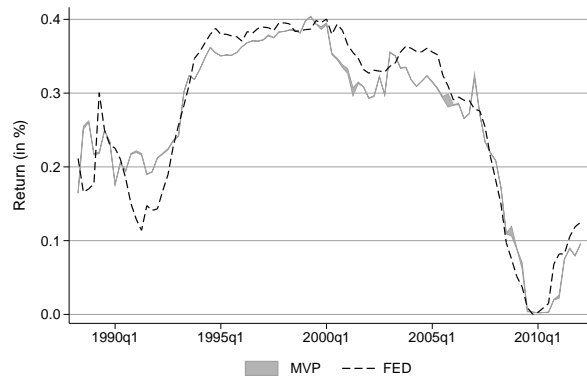
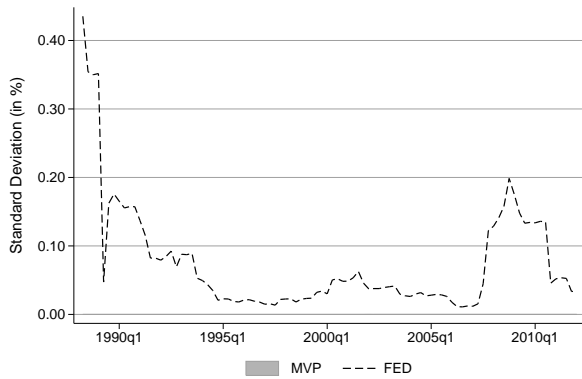
B.3 Equal and Random Starting Weights

Since our baseline specification deals with many banks, the minimization of the portfolio risk is likely to be a complex, highly nonlinear problem comprising multiple minima/solutions. The starting values for the optimization are chosen to be the actual weights in the portfolio, and can have a large impact on whether a global or local minimum is found as well as in what direction the distribution of weights will move. To account for this possible bias, we run two robustness tests. First, we run, for BHC data only, 100 repetitions per quarter using randomized starting values. The random starting values are drawn from a half-normal distribution and then divided by its sum, such that they add up to 1. In the second test, we choose, for the commercial bank data, as starting value an equally weighted portfolio.

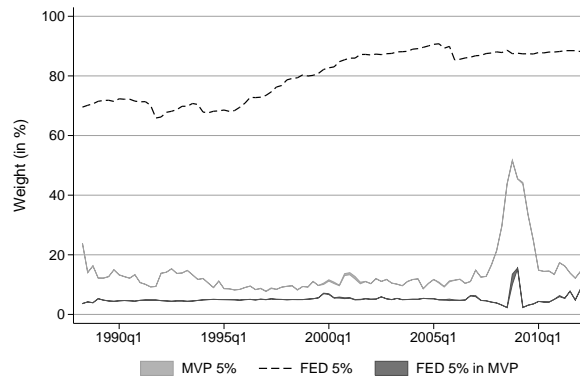
Figure B.3.1 shows the results for the random starting values, while Figure B.3.2 Tables B.3.1 and B.3.2 show the results for the equal starting weights. Figure B.3.1 shows the minimum and maximum portfolio risk, return and concentration found in each period and we can see that there is hardly any variation in the outcomes. The optimization using equal starting weights show almost identical results to those of the baseline specification. Overall, we take this as evidence that we approach a global optimum in the baseline specification.

Figure B.3.1: Do different starting weights values matter in the optimization?

- (a) How much is systemic risk reduced in the MVP?
 (b) Do returns have to be sacrificed to achieve this lower risk?




- (c) How unequal is the MVP compared to the FED portfolio?



The figure presents the robustness test in which 100 random starting values were chosen per quarter, to test whether using the actual weights as starting values drives the results. The lower and upper bound are displayed to show that starting values do not matter.

Figure B.3.2: Do different starting weights values matter in the optimization?



The figure presents the robustness test in which an equally weighted portfolio was chosen as starting weights, instead of the actual weights.

Table B.3.1: Portfolio Optimization using Equal Starting Weights - Risk-Return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.032***	0.131***	-0.037***	0.027
New York	0.044***	0.167***	-0.016***	-0.046*
Philadelphia	-0.105***	0.036***	-0.271***	0.107**
Cleveland	0.011***	0.036***	-0.025***	0.074**
Richmond	-0.023***	0.037***	-0.050***	-0.075*
Atlanta	-0.028***	-0.041*	-0.047***	0.065***
Chicago	0.018***	0.042***	-0.017**	0.079***
St. Louis	-0.019***	-0.020***	-0.032***	0.023
Minneapolis	-0.123***	-0.091***	-0.150***	-0.114**
Kansas City	-0.046***	-0.063***	-0.043***	-0.020
Dallas	0.012***	0.070**	-0.019***	-0.019
San Francisco	-0.081***	-0.075**	-0.100***	-0.030*
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	1.000***	1.000***	1.000***	1.000***
New York	1.000***	1.000***	1.000***	1.000***
Philadelphia	1.000***	1.000***	1.000***	1.000***
Cleveland	1.000***	1.000***	1.000***	1.000***
Richmond	1.000***	1.000***	1.000***	1.000***
Atlanta	1.000***	1.000***	1.000***	1.000***
Chicago	1.000***	1.000***	1.000***	1.000***
St. Louis	1.000***	1.000***	1.000***	1.000***
Minneapolis	1.000***	1.000***	1.000***	1.000***
Kansas City	1.000***	1.000***	1.000***	1.000***
Dallas	1.000***	1.000***	1.000***	1.000***
San Francisco	1.000***	1.000***	1.000***	1.000***

This table shows the results of the analysis on commercial bank level when using equal starting weights. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table B.3.2: Portfolio Optimization using Equal Starting Weights - Weights

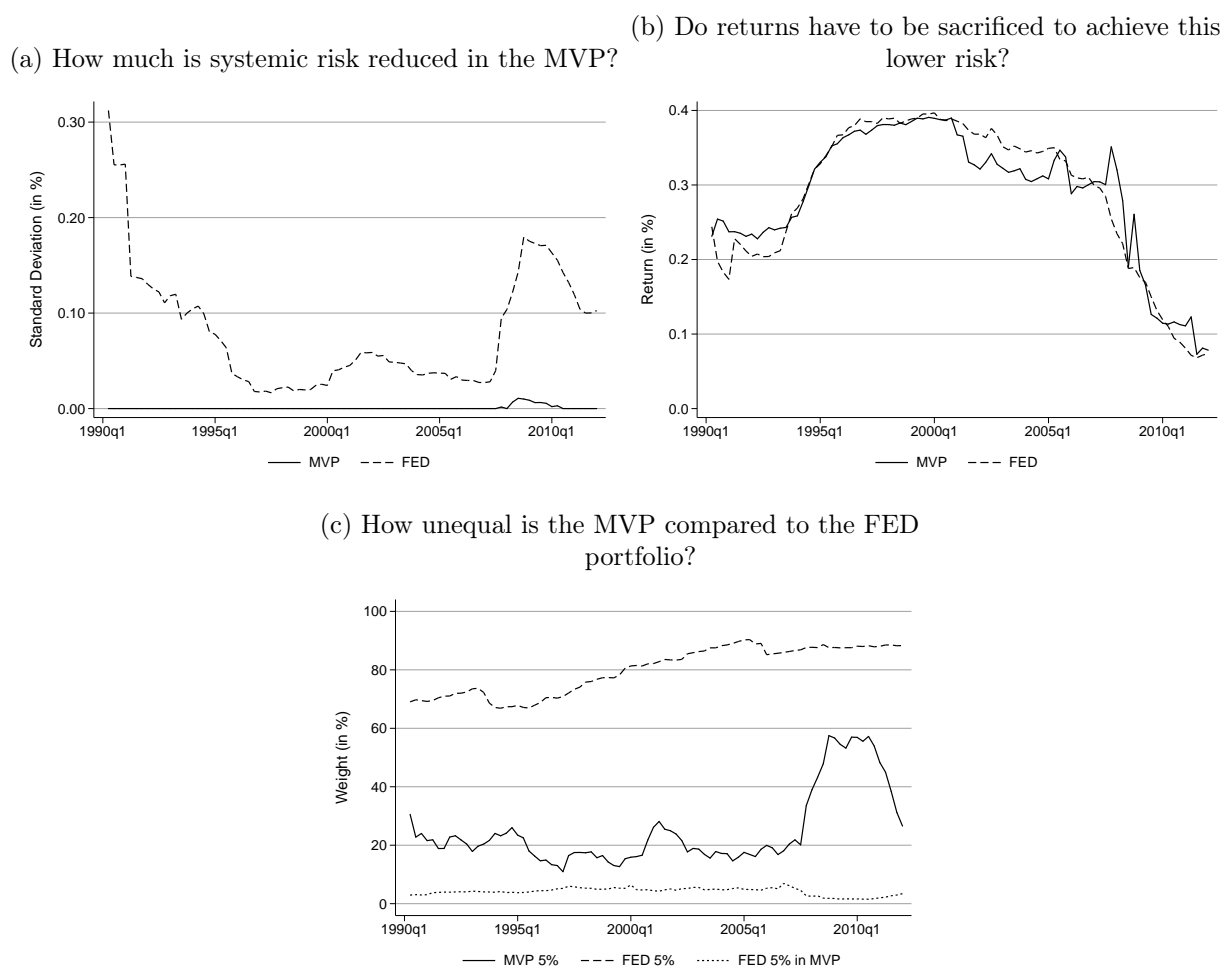
	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u>FRD 5% in MVP - FRD 5%</u>				
Boston	-68.8***	-59.7***	-74.0***	-73.3***
New York	-76.1***	-75.3***	-80.4***	-64.2***
Philadelphia	-58.0***	-45.2***	-60.6***	-79.4***
Cleveland	-75.6***	-62.3***	-79.8***	-92.6***
Richmond	-78.0***	-66.8***	-82.4***	-89.3***
Atlanta	-64.5***	-59.8***	-67.2***	-66.4***
Chicago	-62.2***	-55.1***	-67.4***	-61.7***
St. Louis	-48.5***	-43.6***	-52.4***	-47.1***
Minneapolis	-62.4***	-50.5***	-63.9***	-85.0***
Kansas City	-43.8***	-38.3***	-47.3***	-45.2***
Dallas	-52.3***	-51.3***	-54.0***	-49.1***
San Francisco	-71.1***	-75.2***	-67.8***	-72.3***
<u>MVP 5% - FRD 5%</u>				
Boston	-54.2***	-42.2***	-63.4***	-51.6***
New York	-65.7***	-61.4***	-73.5***	-49.9***
Philadelphia	-48.5***	-32.5***	-55.1***	-64.1***
Cleveland	-68.1***	-52.6***	-74.9***	-82.1***
Richmond	-68.0***	-53.7***	-74.8***	-79.1***
Atlanta	-55.3***	-48.0***	-62.2***	-49.3***
Chicago	-55.8***	-48.2***	-62.4***	-52.0***
St. Louis	-38.3***	-33.7***	-42.6***	-35.1***
Minneapolis	-52.6***	-40.2***	-55.2***	-72.9***
Kansas City	-35.4***	-31.9***	-37.8***	-35.7***
Dallas	-44.0***	-42.3***	-46.9***	-38.5***
San Francisco	-62.4***	-65.3***	-62.6***	-55.2***

This table shows the results of the analysis on commercial bank level when using equal starting weights. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B.4 Portfolio Optimization using a 16 Quarter Window

Finally, we explore if the choice of an 8 quarter window on which the covariance matrices are estimated matters for the optimization. In order to see if our results are robust, we run the analysis using a 16 quarter window. The results for the analysis on BHC level is shown in Figure B.4.1, whereas the results for the analysis on commercial bank level are shown in Tables B.4.1 and B.4.2, with being quantitatively and qualitatively similar to the baseline specification.

Figure B.4.1: Robustness - 16 Quarter Analysis



The figure shows the comparison between the FED portfolio and the hypothetical MVP when using 16 quarters to estimate the sample covariance matrix: panel B.4.1a and B.4.1b display the difference in the risk and return of each portfolio. Panel B.4.1c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Table B.4.1: Portfolio Optimization using a 16 Quarter Window – Risk-Return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.015***	0.140***	-0.045***	-0.018*
New York	0.048***	0.202***	-0.003***	-0.063***
Philadelphia	-0.143***	0.081***	-0.269***	-0.139***
Cleveland	-0.025***	0.048***	-0.066***	-0.026
Richmond	-0.010***	0.050***	-0.025***	-0.073*
Atlanta	-0.036***	-0.060***	-0.044***	0.032
Chicago	0.013***	0.042***	-0.009*	0.036*
St. Louis	-0.026***	-0.023***	-0.037***	0.005
Minneapolis	-0.136***	-0.060***	-0.176***	-0.145***
Kansas City	-0.054***	-0.087***	-0.042***	-0.035**
Dallas	-0.002***	0.056***	-0.024***	-0.035*
San Francisco	-0.068***	-0.039***	-0.094***	-0.037
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	0.993***	0.996***	1.000***	0.963***
New York	1.000***	1.000***	1.000***	1.000***
Philadelphia	1.000***	1.000***	1.000***	1.000***
Cleveland	1.000***	1.000***	0.999***	1.000***
Richmond	1.000***	1.000***	1.000***	0.997***
Atlanta	0.999***	1.000***	1.000***	0.995***
Chicago	1.000***	1.000***	1.000***	1.000***
St. Louis	1.000***	1.000***	1.000***	1.000***
Minneapolis	1.000***	1.000***	1.000***	1.000***
Kansas City	1.000***	1.000***	1.000***	1.000***
Dallas	1.000***	1.000***	1.000***	1.000***
San Francisco	1.000***	1.000***	1.000***	1.000***

This table shows the results of the analysis on commercial bank level when using 16 quarters to estimate the sample covariance matrix. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table B.4.2: Portfolio Optimization using a 16 Quarter Window – Weights

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u>FRD 5% in MVP - FRD 5%</u>				
Boston	-69.5***	-60.5***	-72.7***	-75.6***
New York	-75.8***	-75.1***	-80.1***	-62.9***
Philadelphia	-58.8***	-44.6***	-60.8***	-78.1***
Cleveland	-75.8***	-62.0***	-78.6***	-91.6***
Richmond	-79.0***	-66.4***	-82.8***	-89.3***
Atlanta	-65.7***	-61.7***	-67.7***	-66.4***
Chicago	-63.0***	-55.7***	-67.3***	-62.0***
St. Louis	-49.0***	-44.5***	-51.9***	-47.5***
Minneapolis	-61.7***	-51.7***	-61.2***	-81.8***
Kansas City	-44.1***	-38.2***	-46.9***	-46.0***
Dallas	-51.0***	-48.1***	-54.1***	-46.3***
San Francisco	-66.1***	-74.1***	-67.3***	-47.9***
<u>MVP 5% - FRD 5%</u>				
Boston	-35.5***	-14.5***	-48.6***	-31.4***
New York	-53.3***	-41.5***	-65.3***	-36.0***
Philadelphia	-35.9***	-18.1***	-44.8***	-39.6***
Cleveland	-57.8***	-42.4***	-63.8***	-66.3***
Richmond	-54.8***	-39.8***	-64.3***	-51.0***
Atlanta	-44.8***	-37.9***	-54.7***	-25.4***
Chicago	-49.7***	-43.3***	-56.5***	-39.1***
St. Louis	-27.6***	-25.8***	-30.4***	-22.0***
Minneapolis	-42.8***	-32.8***	-45.1***	-53.0***
Kansas City	-27.6***	-26.0***	-30.0***	-22.5***
Dallas	-33.0***	-28.5***	-38.2***	-24.1***
San Francisco	-49.7***	-54.3***	-57.0***	-17.4***

This table shows the results of the analysis on commercial bank level when using 16 quarters to estimate the sample covariance matrix. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C Extensive Results: Extensions and Policy Implications

In this final section of the Internet Appendix, we examine alternative scenarios which approximate better the current regulatory framework. In section C.1 we examine several scenarios that would limit the amount of reweighing a supervisor could perform when determining its MVP. Finally, in section C.2 we repeat the analysis on a more realistic sample of banks in which the supervisor could invest.

C.1 Portfolio Optimization under Limited Reweighing

Given the large amount of turnover in the baseline specification, a more realistic scenario would be that the supervisor is able to change the size of a bank within certain limitations. We explore several of these scenarios here, starting with a static constraint that the bank supervisor can add or subtract a certain percentage of an individual institution. Figure C.1.1 and Tables C.1.1-C.1.2 show the results for the scenario when supervisors can not change banks' total assets more than 10%, and Figure C.1.2 and Tables C.1.3-C.1.4 do the same for 20%. For both BHCs and commercial banks, any improvement is hardly noticeable for the 10%-scenario, although the portfolio risk is somewhat lower for the 20%-scenario. Neither, however, match up to the baseline scenario and still exhibit the twin peaks in portfolio risk during the two crisis periods.

Another scenario is that the percentage with which the supervisor can reweigh his individual holdings depends on the business cycle. We construct a measure of asset growth in the quarter before the supervisor can set her weights, and allow reweighing equal to at most the average growth and the standard deviation of this growth: when average growth is larger, or more dispersed, this allows for more reweighing by the supervisor in order to reduce systemic risk. The measure, which acts as a time-varying constraint, is shown in Figure C.1.3 using the data on asset growth for the BHCs. Indeed we see that the measures indeed exhibit time-variation, especially the standard deviation of the growth. In fact, the latter is even above 1 in 1997Q4, meaning that based on this constraint alone, it would be possible for the supervisor to short her banks. Since this is not possible, we cut the value off at 1 and thus end up in a situation of unlimited reweighing similar to the baseline specification for this quarter.



Figures C.1.4 and Tables C.1.5-C.1.6 display the results where the maximum amount of reweighing equals the mean of the asset growth in the previous quarter, whereas Figures C.1.5 and Tables C.1.7-C.1.8 do the same for the standard deviation of the growth. Outside of 1997Q4, which equals the baseline specification, there are no real improvements for both BHCs and commercial banks.

So far the supervisor has only been able to reweigh based on the weight that the bank had in the actual portfolio. In the next test we allow for dynamic reweighing where the supervisor can change the size of a bank with a percentage of its size in the previous quarter's MVP. In the first quarter that a bank is eligible for the portfolio, it will enter with its actual size. However, once it has been given an MVP weight (which lies in a range around its original size), the supervisor will keep on minimizing with the constraint on the size the bank had in the MVP rather than the actual portfolio. Since this is computationally more intensive, we perform this test only on BHC level. Figure C.1.6 shows the results for a dynamic reweighing of a maximum of 10%, while Figure C.1.7 shows the results for 20%. We can see no real improvement, although portfolio risk is again reduce more in the 20%-scenario compared to the 10%-scenario.


Given the inability to reduce portfolio risk when constraining the reweighing of individual banks, we finally explore some scenarios in which we constrain reweighing of the largest 5% of banks. Specifically, we add a constraint that the largest 5% of banks need to retain their current individual size, and allow unlimited reweighing of the remaining banks while still adhering to the no-shorting and no-loss constraint. The results for the commercial bank level are shown in Tables C.1.9 and C.1.10. There are 3 districts which were able to reduce systemic risk to a minimum in each of the three sub-samples. Interestingly, a common feature they share is that they have the lowest levels of concentration among all districts. Since the weight of these largest banks in the three successful districts never exceeds 60%, this leads us to a final test using limited reweighing: is it possible to reduce portfolio risk while keeping the cumulative weight of the largest 5% of banks between 50% and 60%? The results for the commercial banks are reported in Tables C.1.11 and C.1.12. We observe that under this limited reweighing scheme, it is possible to effectively eliminate risk while maintaining similar returns.

Figure C.1.1: Portfolio Optimization using Limited Reweighing - 10%

- (a) How much is systemic risk reduced in the MVP? (b) Do returns have to be sacrificed to achieve this lower risk?



- (c) How unequal is the MVP compared to the FED portfolio?



The graph shows the comparison between the FED portfolio and the hypothetical MVP when constraining individual reweighing to 10% of the original bank size: panel C.1.1a and C.1.1b display the difference in the risk and return of each portfolio. Panel C.1.1c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Table C.1.1: Portfolio Optimization using Limited Reweighing of 10% – Risk-Return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.004***	0.008	0.002	-0.000
New York	0.004***	0.006	0.004	0.001
Philadelphia	0.000***	0.005	-0.008	0.018
Cleveland	0.003***	0.006	0.002	0.002
Richmond	0.003***	0.005	0.001	0.002
Atlanta	0.002***	0.004	-0.001	0.010
Chicago	0.004***	0.006	0.001	0.009
St. Louis	0.002***	0.002	0.000	0.006
Minneapolis	-0.004***	0.001	-0.008	-0.003
Kansas City	0.002***	0.005	-0.000	0.004
Dallas	0.005***	0.012	0.000	0.003
San Francisco	0.002***	0.007	-0.003	0.004
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	0.070***	0.070	0.074	0.059
New York	0.077***	0.060	0.081	0.103
Philadelphia	0.098***	0.101	0.100	0.084
Cleveland	0.089***	0.093	0.092	0.070
Richmond	0.066***	0.086	0.059	0.038
Atlanta	0.087***	0.076	0.105	0.052
Chicago	0.097***	0.093	0.110	0.067
St. Louis	0.081***	0.076	0.087*	0.070
Minneapolis	0.076***	0.080	0.085**	0.038
Kansas City	0.084***	0.085	0.088*	0.070
Dallas	0.082***	0.062	0.097	0.083
San Francisco	0.092***	0.086	0.108	0.057

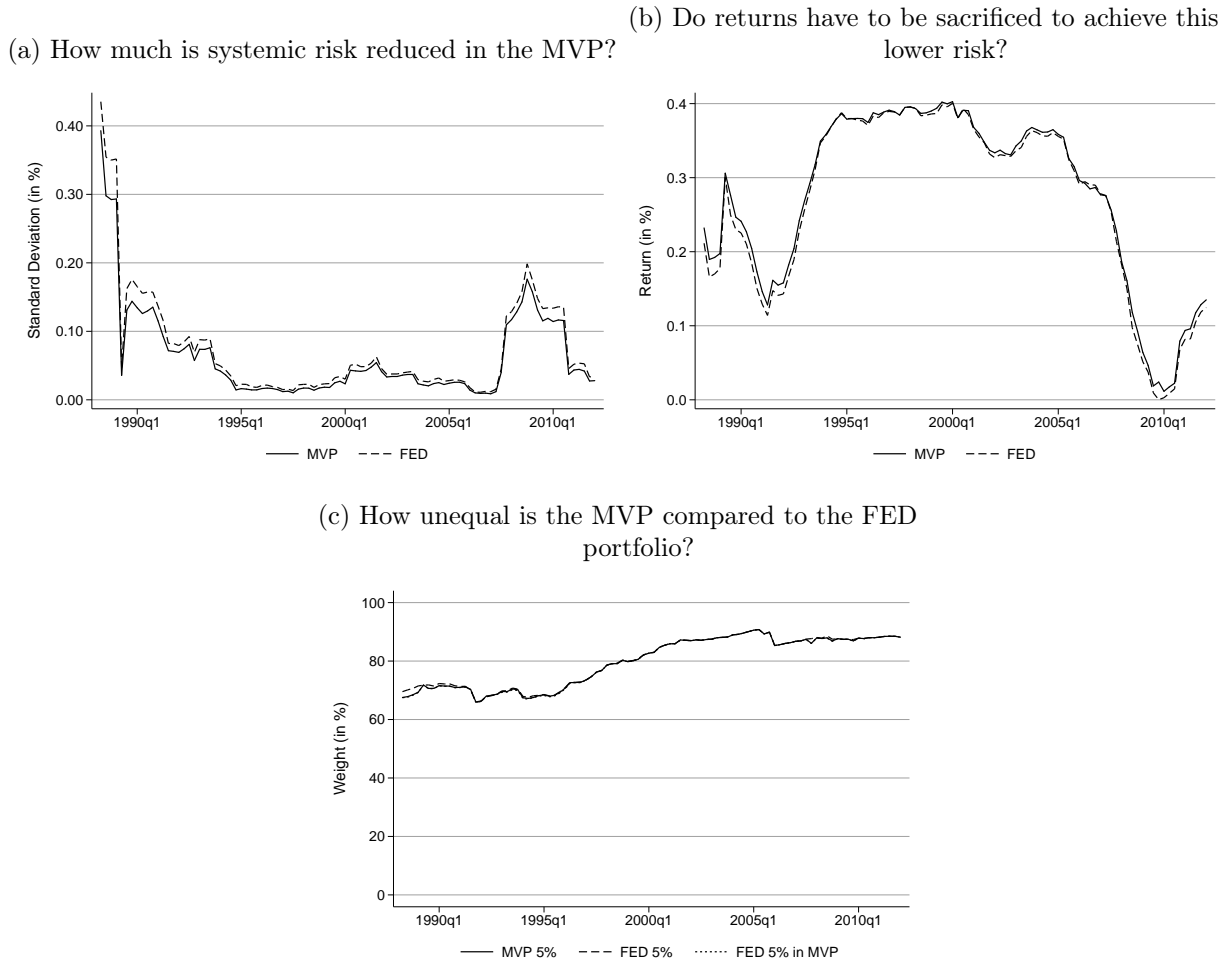
This table shows the results of the analysis on commercial bank level when constraining individual reweighing to 10% of the original bank size. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table C.1.2: Portfolio Optimization using Limited Reweighting of 10% – Weight

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u>FRD 5% in MVP - FRD 5%</u>				
Boston	-0.6***	-0.6	-0.7***	-0.4
New York	-0.4***	-0.6	-0.2	-0.5
Philadelphia	-0.4***	-0.2	-0.4	-0.6
Cleveland	-0.2***	-0.5**	-0.1	-0.0
Richmond	-0.1***	-0.1	-0.1	-0.0
Atlanta	-0.0***	0.2	-0.0	-0.3
Chicago	-0.2***	-0.2	-0.1	-0.5
St. Louis	0.0***	0.3	-0.1	-0.3
Minneapolis	-0.2***	-0.2	-0.2	-0.1
Kansas City	0.1***	0.2*	0.0	0.0
Dallas	-0.2***	-0.3	-0.1	-0.0
San Francisco	-0.2***	-0.2	-0.3	-0.1
<u>MVP 5% - FRD 5%</u>				
Boston	-0.6***	-0.6	-0.6***	-0.4
New York	-0.4***	-0.6	-0.2	-0.5
Philadelphia	-0.4***	-0.2	-0.4	-0.6
Cleveland	-0.2***	-0.5**	-0.1	-0.0
Richmond	-0.1***	-0.1	-0.1	-0.0
Atlanta	0.0***	0.2	-0.0	-0.3
Chicago	-0.2***	-0.2	-0.1	-0.4
St. Louis	0.0***	0.3	-0.1	-0.3
Minneapolis	-0.2***	-0.2	-0.2	-0.1
Kansas City	0.1***	0.2*	0.1	0.0
Dallas	-0.2***	-0.3	-0.1	0.0
San Francisco	-0.2***	-0.2	-0.3	-0.0

This table shows the results of the analysis on commercial bank level when constraining individual reweighting to 10% of the original bank size. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Figure C.1.2: Portfolio Optimization using Limited Reweighing - 20%



The figure shows the comparison between the FED portfolio and the hypothetical MVP when constraining individual reweighing to 20% of the original bank size: panel C.1.2a and C.1.2b display the difference in the risk and return of each portfolio. Panel C.1.2c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Table C.1.3: Portfolio Optimization using Limited Reweighing of 20% – Risk-Return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.008***	0.017	0.003	-0.001
New York	0.009***	0.012	0.008	0.004
Philadelphia	-0.001***	0.010*	-0.019*	0.035*
Cleveland	0.007***	0.012**	0.004	0.004
Richmond	0.005***	0.010	0.002	0.004
Atlanta	0.005***	0.008	-0.002	0.019
Chicago	0.008***	0.012	0.002	0.018
St. Louis	0.004***	0.005**	0.001	0.012
Minneapolis	-0.009***	0.002	-0.017*	-0.007
Kansas City	0.004***	0.010*	-0.000	0.007
Dallas	0.009***	0.022	0.001	0.005
San Francisco	0.003***	0.014**	-0.006	0.007
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	0.142***	0.139	0.152	0.116*
New York	0.154***	0.121	0.161***	0.205*
Philadelphia	0.202***	0.209***	0.207*	0.169
Cleveland	0.175***	0.178	0.184***	0.140
Richmond	0.125***	0.162	0.115	0.074
Atlanta	0.166***	0.153**	0.194**	0.105
Chicago	0.188***	0.188	0.203***	0.140
St. Louis	0.163***	0.150**	0.176***	0.149
Minneapolis	0.156***	0.163	0.175***	0.076
Kansas City	0.176***	0.181*	0.181***	0.148
Dallas	0.167***	0.127*	0.196*	0.164
San Francisco	0.185***	0.180***	0.210*	0.113


This table shows the results of the analysis on commercial bank level when constraining individual reweighing to 20% of the original bank size. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table C.1.4: Portfolio Optimization using Limited Reweighting of 20% – Weight

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u>FRD 5% in MVP - FRD 5%</u>				
Boston	-1.4***	-1.2**	-1.6***	-0.9
New York	-1.0***	-1.4**	-0.6	-1.3
Philadelphia	-1.0***	-0.5	-1.3**	-1.3
Cleveland	-0.5***	-1.1**	-0.2	-0.0
Richmond	-0.1***	-0.2	-0.2	-0.0
Atlanta	0.0***	0.4	-0.0	-0.6
Chicago	-0.3***	-0.4	-0.1	-1.0
St. Louis	-0.1***	0.5	-0.2	-0.7
Minneapolis	-0.4***	-0.4	-0.4	-0.2
Kansas City	0.2***	0.4**	0.2	0.0
Dallas	-0.3***	-0.5	-0.2	-0.1
San Francisco	-0.5***	-0.4*	-0.7	-0.1
<u>MVP 5% - FRD 5%</u>				
Boston	-1.3***	-1.2**	-1.5***	-0.9
New York	-0.9***	-1.4**	-0.6	-1.1
Philadelphia	-0.9***	-0.2	-1.2**	-1.3
Cleveland	-0.5***	-1.0**	-0.2	-0.0
Richmond	-0.1***	-0.1	-0.2	-0.0
Atlanta	0.1***	0.4	-0.0	-0.5
Chicago	-0.3***	-0.4	-0.1	-0.9
St. Louis	0.0***	0.6**	-0.2	-0.6
Minneapolis	-0.4***	-0.4	-0.4	-0.2
Kansas City	0.3***	0.5***	0.2	0.1
Dallas	-0.2***	-0.4	-0.1	-0.1
San Francisco	-0.5***	-0.4*	-0.7	-0.0


This table shows the results of the analysis on commercial bank level when constraining individual reweighting to 20% of the original bank size. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Figure C.1.3: Limited Reweighting - Mean and Standard Deviation of Growth Variable



The figure shows the mean and standard deviation of the asset growth in the previous quarter, which are used as a time-varying constraint determining how much the supervisor can alter the assets of individual banks.

Figure C.1.4: Portfolio Optimization using Limited Reweighing - Mean of Growth



The figure shows the comparison between the FED portfolio and the hypothetical MVP when using the mean asset growth as a time-varying constraint for individual reweighing of banks: panel C.1.4a and C.1.4b display the difference in the risk and return of each portfolio. Panel C.1.4c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Table C.1.5: Portfolio Optimization using Limited Reweighting Mean of Growth – Risk-Return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.003***	0.009	0.000	-0.000
New York	0.002***	0.002	0.001	0.001
Philadelphia	0.002***	0.002	-0.003	0.018
Cleveland	0.001***	0.001	0.000	0.000
Richmond	0.001***	0.001	0.000	0.000
Atlanta	0.001***	0.001	-0.000	0.006
Chicago	0.001***	0.001	0.000	0.003
St. Louis	0.000***	0.000	0.000	0.001
Minneapolis	-0.003***	-0.000	-0.002	-0.012
Kansas City	0.000***	0.001	-0.000	0.001
Dallas	0.002***	0.005	0.000	0.001
San Francisco	0.001***	0.003	-0.002	0.002
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	0.041***	0.076	0.024	0.013
New York	0.031***	0.024	0.033	0.037
Philadelphia	0.040***	0.043	0.039	0.038
Cleveland	0.019***	0.022	0.018	0.017
Richmond	0.021***	0.027	0.021	0.010
Atlanta	0.028***	0.020	0.035	0.024
Chicago	0.019***	0.015	0.023	0.015
St. Louis	0.013***	0.011	0.015	0.012
Minneapolis	0.027***	0.013	0.024	0.069
Kansas City	0.011***	0.011	0.012	0.011
Dallas	0.018***	0.017	0.017	0.022
San Francisco	0.046***	0.038	0.059	0.024

This table shows the results of the analysis on commercial bank level when using the mean asset growth as a time-varying constraint for individual reweighting of banks. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

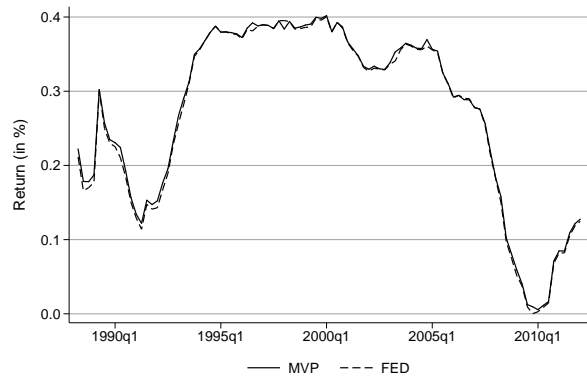

Table C.1.6: Portfolio Optimization using Limited Reweighting Mean of Growth – Weight

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u><i>FRD 5% in MVP - FRD 5%</i></u>				
Boston	-1.0***	-2.4	-0.2	-0.1
New York	-0.2***	-0.2	-0.1	-0.2
Philadelphia	-0.2***	-0.1	-0.2	-0.2
Cleveland	-0.0***	-0.1	-0.0	-0.0
Richmond	-0.0***	-0.0	-0.0	0.0
Atlanta	-0.0***	0.0	-0.0	-0.1
Chicago	-0.0***	-0.0	-0.0	-0.1
St. Louis	0.0***	0.1	-0.0	-0.0
Minneapolis	-0.8***	-0.0	-0.0	-5.2
Kansas City	0.0***	0.0	0.0	-0.0
Dallas	-0.0***	-0.1	-0.0	-0.0
San Francisco	-0.1***	-0.0	-0.2	-0.0
<u><i>MVP 5% - FRD 5%</i></u>				
Boston	-0.7***	-1.6	-0.2	-0.1
New York	-0.1***	-0.2	-0.1	-0.1
Philadelphia	-0.1***	-0.1	-0.2	-0.1
Cleveland	-0.0***	-0.1	-0.0	-0.0
Richmond	-0.0***	-0.0	-0.0	0.0
Atlanta	-0.0***	0.0	-0.0	-0.1
Chicago	-0.0***	-0.0	-0.0	-0.1
St. Louis	0.0***	0.1	-0.0	-0.0
Minneapolis	-0.7***	-0.0	-0.0	-4.4
Kansas City	0.0***	0.0	0.0	-0.0
Dallas	-0.0***	-0.1	-0.0	-0.0
San Francisco	-0.1***	-0.0	-0.2	-0.0


This table shows the results of the analysis on commercial bank level when using the mean asset growth as a time-varying constraint for individual reweighting of banks. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Figure C.1.5: Portfolio Optimization using Limited Reweighting - Std. Dev. of Growth

(a) How much is systemic risk reduced in the MVP?
 (b) Do returns have to be sacrificed to achieve this lower risk?



(c) How unequal is the MVP compared to the FED portfolio?



The figure shows the comparison between the FED portfolio and the hypothetical MVP level when using the standard deviation of asset growth as a time-varying constraint for individual reweighting of banks: panel C.1.5a and C.1.5b display the difference in the risk and return of each portfolio. Panel C.1.5c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Table C.1.7: Portfolio Optimization using Limited Reweighting Std. Dev. of Growth – Risk-Return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.008***	0.021	0.002	-0.001
New York	0.008***	0.011	0.008	0.002
Philadelphia	0.005***	0.015*	-0.024	0.079
Cleveland	0.005***	0.010	0.002	0.001
Richmond	0.004***	0.010	0.001	0.002
Atlanta	0.005***	0.008	-0.002	0.022
Chicago	0.005***	0.008	0.001	0.010
St. Louis	0.001***	0.000	0.001	0.003
Minneapolis	-0.009***	-0.000	-0.014	-0.015
Kansas City	0.002***	0.006	-0.000	0.003
Dallas	0.010***	0.028	0.001	0.002
San Francisco	-0.002***	0.008	-0.012	0.009
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	0.126***	0.169	0.110	0.080
New York	0.142***	0.110	0.163*	0.147
Philadelphia	0.215***	0.267***	0.178	0.215
Cleveland	0.112***	0.157	0.093	0.065
Richmond	0.094***	0.138	0.080	0.040
Atlanta	0.154***	0.134***	0.180**	0.119
Chicago	0.128***	0.113	0.156**	0.070
St. Louis	0.092***	0.092	0.104**	0.053
Minneapolis	0.120***	0.088	0.153***	0.087
Kansas City	0.109***	0.109	0.124*	0.061
Dallas	0.095***	0.115	0.088	0.074
San Francisco	0.223***	0.230**	0.246**	0.134

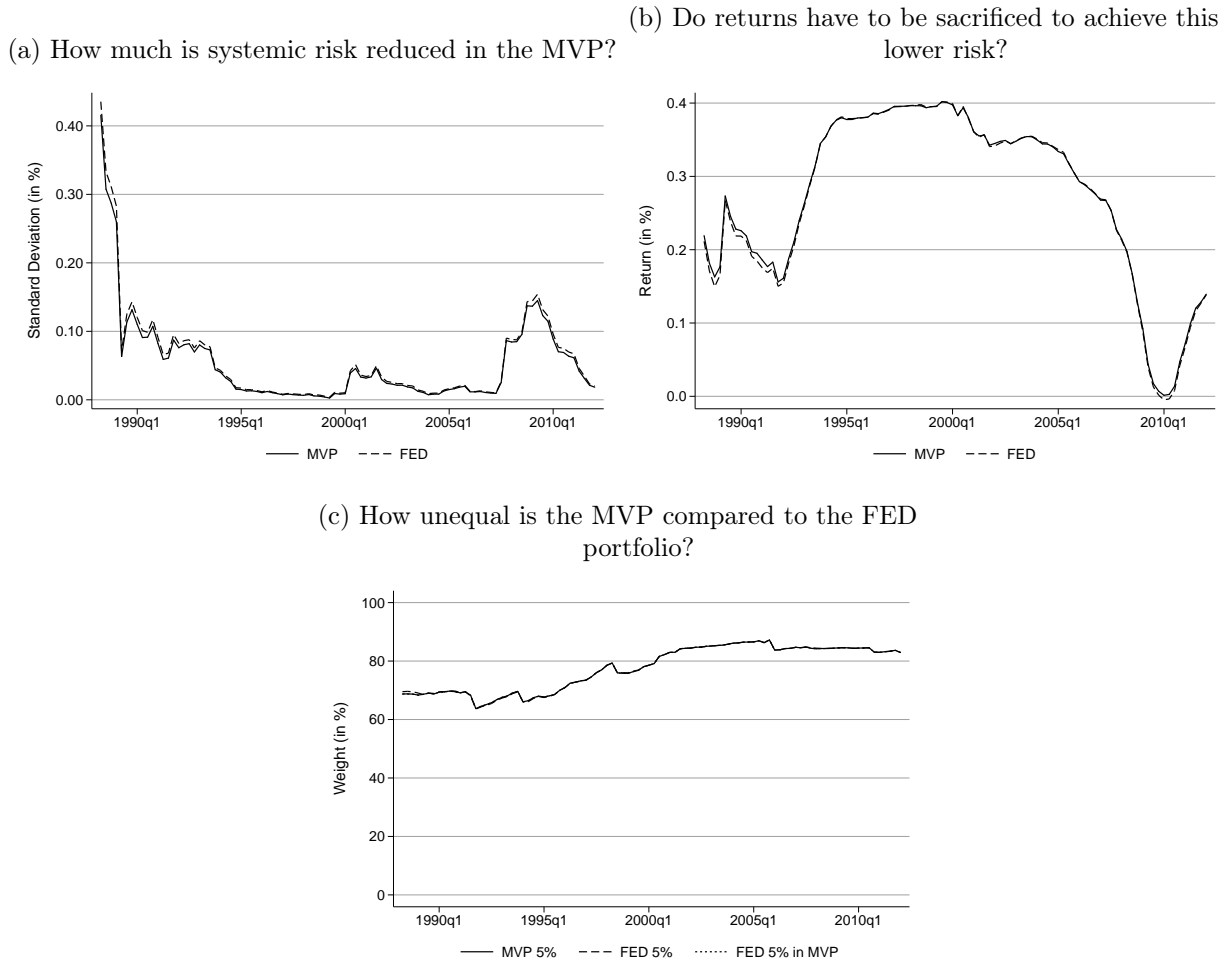
This table shows the results of the analysis on commercial bank level when using the standard deviation of asset growth as a time-varying constraint for individual reweighting of banks. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table C.1.8: Portfolio Optimization using Limited Reweighting Std. Dev. of Growth – Weight

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u>FRD 5% in MVP - FRD 5%</u>				
Boston	-2.3***	-4.5**	-1.2***	-0.6
New York	-0.9***	-1.3**	-0.7	-0.8
Philadelphia	-4.1***	-4.1	-3.2**	-6.6
Cleveland	-0.9***	-2.5**	-0.1	-0.0
Richmond	-0.7***	-2.0	-0.1	-0.0
Atlanta	-1.2***	0.3	-1.3	-4.4
Chicago	-0.8***	-0.2	-1.3	-0.5
St. Louis	-0.4***	-0.9	-0.1	-0.2
Minneapolis	-2.2***	-0.2	-2.6	-5.3
Kansas City	0.1***	0.2	0.1	0.0
Dallas	-0.7***	-1.8	-0.1	-0.0
San Francisco	-5.4***	-4.6*	-6.0	-5.2
<u>MVP 5% - FRD 5%</u>				
Boston	-1.9***	-3.4**	-1.2***	-0.6
New York	-0.8***	-1.3**	-0.5	-0.7
Philadelphia	-3.2***	-2.2	-3.0**	-6.0
Cleveland	-0.8***	-2.2**	-0.1	-0.0
Richmond	-0.6***	-1.6	-0.1	-0.0
Atlanta	-1.0***	0.4	-1.3	-3.2
Chicago	-0.7***	-0.2	-1.2	-0.4
St. Louis	-0.3***	-0.6	-0.1	-0.2
Minneapolis	-1.9***	-0.2	-2.4	-4.5
Kansas City	0.1***	0.2*	0.1	0.0
Dallas	-0.6***	-1.5	-0.1	0.0
San Francisco	-4.9***	-4.0*	-5.7	-3.9

This table shows the results of the analysis on commercial bank level when using the standard deviation of asset growth as a time-varying constraint for individual reweighing of banks. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.



Figure C.1.6: Portfolio Optimization using Dynamic Limited Reweighing of 10%




The figure shows the comparison between the FED portfolio and the hypothetical MVP when constraining individual reweighing to 10% of the bank size in the previous quarter's MVP: panel C.1.6a and C.1.6b display the difference in the risk and return of each portfolio. Panel C.1.6c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Figure C.1.7: Portfolio Optimization using Dynamic Limited Reweighing of 20%

(a) How much is systemic risk reduced in the MVP? (b) Do returns have to be sacrificed to achieve this lower risk?



(c) How unequal is the MVP compared to the FED portfolio?



The figure shows the comparison between the FED portfolio and the hypothetical MVP when constraining individual reweighing to 20% of the bank size in the previous quarter's MVP: panel C.1.7a and C.1.7b display the difference in the risk and return of each portfolio. Panel C.1.7c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

Table C.1.9: Portfolio Optimization Keeping the Largest 5% at Their Actual Weight – Risk-Return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.060***	0.094***	0.046***	0.025
New York	-0.001***	0.041	-0.009*	-0.076*
Philadelphia	0.005***	0.063***	-0.032*	-0.010
Cleveland	0.005***	-0.011	0.015**	0.008
Richmond	-0.028***	-0.032	-0.026***	-0.028
Atlanta	-0.014***	-0.044	-0.016*	0.061**
Chicago	-0.029***	-0.061***	-0.017	0.006
St. Louis	-0.035***	-0.036***	-0.033***	-0.038
Minneapolis	-0.037***	-0.114***	-0.002*	0.028
Kansas City	-0.052***	-0.072***	-0.035***	-0.059**
Dallas	-0.010***	0.047**	-0.035***	-0.062**
San Francisco	-0.018***	-0.090***	0.027*	0.001*
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	0.634***	0.784***	0.608***	0.373*
New York	0.623***	0.646***	0.620***	0.579***
Philadelphia	0.798***	0.975***	0.869***	0.157
Cleveland	0.738***	0.966***	0.747***	0.183
Richmond	0.651***	0.994***	0.476***	0.425
Atlanta	0.927***	1.000***	0.999***	0.520*
Chicago	0.905***	0.950***	0.893***	0.844***
St. Louis	0.998***	1.000***	1.000***	0.987***
Minneapolis	0.876***	0.990***	0.943***	0.391***
Kansas City	1.000***	1.000***	1.000***	1.000***
Dallas	0.965***	0.908***	1.000***	0.980***
San Francisco	0.898***	0.933***	0.967***	0.594***

This table shows the results of the analysis on commercial bank level when constraining the largest 5% of banks to their individual current weigh, while allowing unlimited reweighing for the remaining banks. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table C.1.10: Portfolio Optimization Keeping the Largest 5% at Their Actual Weight – Weight

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u>FRD 5% in MVP - FRD 5%</u>				
Boston	0.0***	0.0	0.0	0.0
New York	0.0***	0.0	0.0	0.0
Philadelphia	0.0***	0.0	0.0	0.0
Cleveland	0.0***	0.0	0.0	0.0
Richmond	0.0***	0.0	0.0	0.0
Atlanta	0.0***	0.0	0.0	0.0
Chicago	0.0***	0.0	0.0	0.0
St. Louis	0.0***	0.0	0.0	0.0
Minneapolis	0.0***	0.0	0.0	0.0
Kansas City	0.0***	0.0	0.0	0.0
Dallas	0.0***	0.0	0.0	0.0
San Francisco	0.0***	0.0	0.0	0.0
<u>MVP 5% - FRD 5%</u>				
Boston	7.8***	13.3***	5.7***	1.6*
New York	8.3***	10.4***	5.3***	13.2***
Philadelphia	13.8***	16.3***	14.0***	7.2***
Cleveland	7.9***	13.1***	6.1***	1.5***
Richmond	7.0***	7.5***	7.6***	4.1***
Atlanta	6.9***	5.8***	3.7***	19.9***
Chicago	9.8***	13.4***	4.9***	17.6***
St. Louis	8.6***	5.1***	7.4***	20.5***
Minneapolis	12.4***	18.0***	10.0***	7.3***
Kansas City	5.6***	3.9***	5.1***	11.1***
Dallas	11.6***	17.9***	5.3***	17.7***
San Francisco	7.3***	8.4***	6.6***	7.2***

This table shows the results of the analysis on commercial bank level when constraining the largest 5% of banks to their individual current weigh, while allowing unlimited reweighing for the remaining banks. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table C.1.11: Portfolio Optimization Keeping the Largest 5% Between 50% and 60% – Risk-Return Trade-off

	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
$r_{MVP} - r_{FRD}$				
Boston	0.060***	0.140***	0.003	0.058
New York	0.051***	0.128***	0.021***	-0.028*
Philadelphia	-0.049***	0.057***	-0.178***	0.123**
Cleveland	0.032***	0.067***	-0.011***	0.093***
Richmond	-0.006***	0.034***	-0.026***	-0.037
Atlanta	0.006***	0.001	-0.017***	0.091***
Chicago	0.044***	0.085***	0.005	0.076***
St. Louis	0.002***	0.010***	-0.022***	0.059***
Minneapolis	-0.074***	-0.033**	-0.108***	-0.060
Kansas City	-0.017***	-0.029**	-0.013**	-0.003
Dallas	0.031***	0.102**	-0.013***	0.005
San Francisco	-0.013***	0.015*	-0.035***	-0.004*
$(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$				
Boston	0.994***	1.000***	1.000***	0.964***
New York	1.000***	1.000***	1.000***	1.000***
Philadelphia	1.000***	1.000***	1.000***	1.000***
Cleveland	1.000***	1.000***	1.000***	1.000***
Richmond	1.000***	1.000***	1.000***	1.000***
Atlanta	1.000***	1.000***	1.000***	1.000***
Chicago	1.000***	1.000***	1.000***	1.000***
St. Louis	1.000***	1.000***	1.000***	1.000***
Minneapolis	1.000***	1.000***	1.000***	1.000***
Kansas City	1.000***	1.000***	1.000***	1.000***
Dallas	1.000***	1.000***	1.000***	1.000***
San Francisco	1.000***	1.000***	1.000***	1.000***

This table shows the results of the analysis on commercial bank level when constraining the largest 5% of banks to have a weight between 50% and 60%. $r_{MVP} - r_{FRD}$ measures the difference between returns of the actual portfolio and the returns of the minimum variance portfolio. $(\sigma_{FRD} - \sigma_{MVP})/\sigma_{FRD}$ measures the relative difference in portfolio risk between the actual and minimum variance portfolio. A score of 1 indicates that the risk has been effectively eliminated, while a score of 0 indicates that no improvements are possible. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

Table C.1.12: Portfolio Optimization Keeping the Largest 5% Between 50% and 60% – Weight

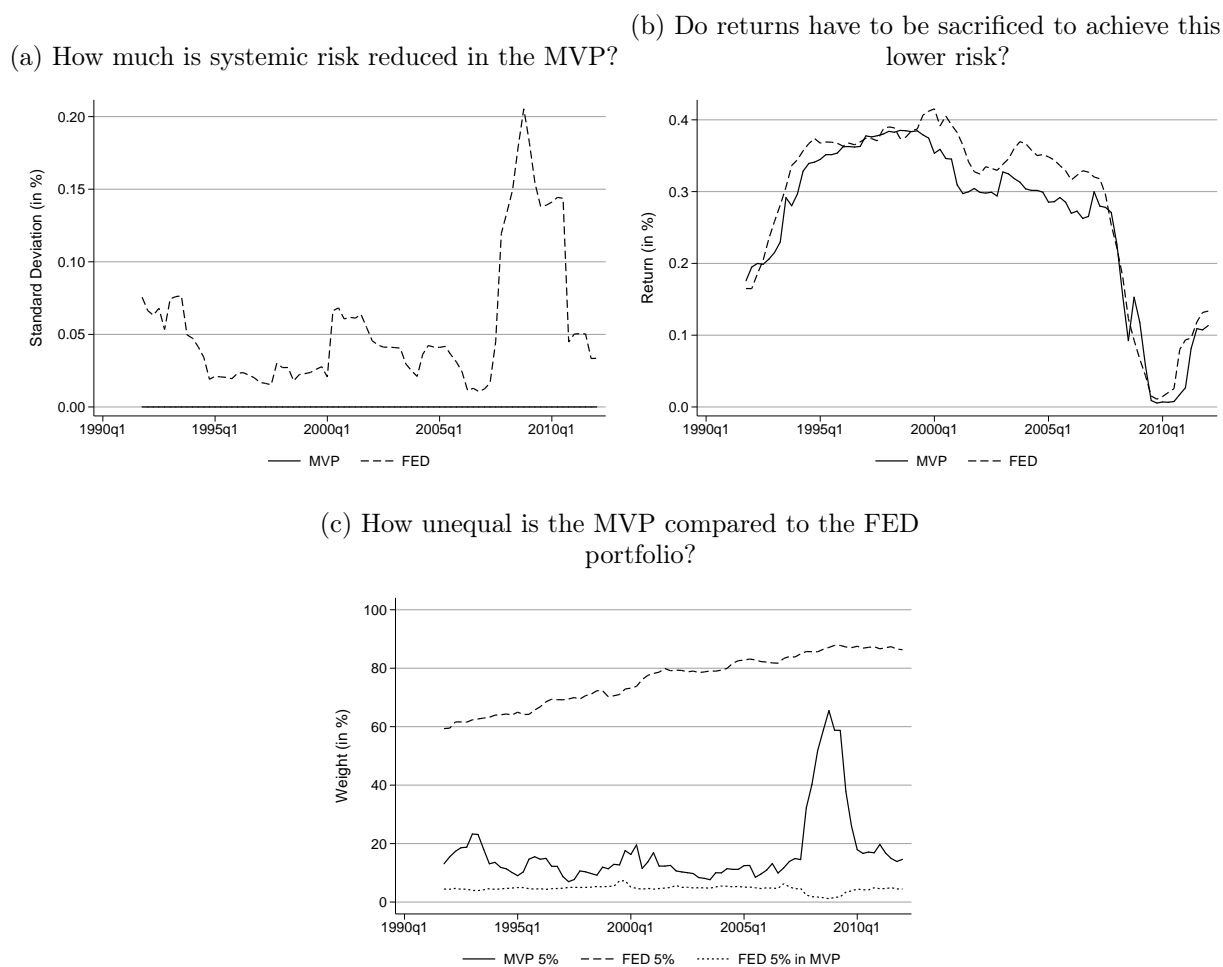
	1984Q4 - 2010Q4	1984Q4 - 1993Q4	1994Q1 - 2006Q4	2007Q1 - 2010Q4
<u>FRD 5% in MVP - FRD 5%</u>				
Boston	-24.2***	-14.9***	-29.0***	-29.9***
New York	-30.6***	-29.4***	-35.2***	-18.3***
Philadelphia	-13.1***	0.1***	-15.9***	-34.1***
Cleveland	-30.6***	-17.5***	-34.8***	-46.8***
Richmond	-33.0***	-22.1***	-37.3***	-44.0***
Atlanta	-19.3***	-14.5***	-22.1***	-21.1***
Chicago	-17.3***	-10.0***	-22.8***	-15.9***
St. Louis	-3.9***	1.4***	-8.4***	-1.8***
Minneapolis	-18.1***	-5.7***	-20.0***	-40.3***
Kansas City	1.2***	6.5***	-2.3***	0.2***
Dallas	-7.1***	-5.9***	-8.9***	-3.9***
San Francisco	-26.1***	-30.0***	-23.0***	-27.2***
<u>MVP 5% - FRD 5%</u>				
Boston	-16.9***	-8.3***	-24.5***	-12.1***
New York	-27.1***	-24.8***	-34.1***	-9.5***
Philadelphia	-9.8***	2.2***	-14.7***	-21.2***
Cleveland	-29.5***	-16.6***	-34.6***	-43.2***
Richmond	-30.7***	-19.0***	-36.1***	-40.0***
Atlanta	-17.5***	-12.2***	-21.9***	-15.7***
Chicago	-16.6***	-9.6***	-22.8***	-12.8***
St. Louis	-2.8***	2.4***	-7.7***	0.9
Minneapolis	-17.1***	-4.8***	-19.5***	-37.7***
Kansas City	2.0***	7.1***	-1.7***	2.2***
Dallas	-5.9***	-4.1***	-8.3***	-1.9
San Francisco	-25.0***	-28.4***	-22.7***	-24.6***

This table shows the results of the analysis on commercial bank level when constraining the largest 5% of banks to have a weight between 50% and 60%. *FRD 5% in MVP - FRD 5%* measures the average difference between weight of the top-5% in the actual portfolio with the weight of the actual top-5% in the minimum variance portfolio. *MVP 5% - FRD 5%* measures the average difference between the weight of the top-5% in the actual portfolio with the weight of the new top-5% in the minimum variance portfolio. Kolmogorov-Smirnov tests are performed to see if the distributions of the actual and minimum variance portfolio are different from each other. *** p<0.01, ** p<0.05, * p<0.1.

C.2 Portfolio Optimization using Listed BHCs

Lastly, we check whether the results hold if we only consider BHCs that had publicly traded equity, using the database of the Federal Reserve Bank of New York (2013). Banks without a listing are removed and only those that do have one are considered for the portfolio the investor can manage. Figure C.2.1 shows the results, which are again similar to the baseline scenario.

Figure C.2.1: Portfolio Optimization using Listed BHCs



The figure shows the comparison between the FED portfolio and the hypothetical MVP when using only BHCs that have publicly traded equity: panel C.2.1a and C.2.1b display the difference in the risk and return of each portfolio. Panel C.2.1c shows how the weights are distributed in each portfolio by plotting their concentration ratios, as well as the weights that the current largest banks have in the MVP.

References

Federal Reserve Bank of New York (2013). CRSP-FRB Link.

Ledoit, O. and M. Wolf (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance* 10(5), pp. 603–621.

Ledoit, O. and M. Wolf (2004). Honey, I shrunk the sample covariance matrix. *The Journal of Portfolio Management* 30(4), pp. 110–119.