Federal Reserve Bank of New York
Staff Reports

Resolving “Too Big to Fail”

Nicola Cetorelli
James Traina

Staff Report No. 859
June 2018

Resolving "Too Big to Fail"
Nicola Cetorelli and James Traina
Federal Reserve Bank of New York Staff Reports, no. 859
June 2018
JEL classification: G21, G28

Abstract

Using a synthetic control research design, we find that "living will" regulation increases a bank's annual cost of capital by 22 basis points, or 10 percent of total funding costs. This effect is stronger in banks that were measured as systemically important before the regulation's announcement. We interpret our findings as a reduction in "too big to fail" subsidies. The size of this effect is large: a back-of-the-envelope calculation implies a subsidy reduction of $42 billion annually. The impact on equity costs drives the main effect. The impact on deposit costs is statistically indistinguishable from zero, representing a good placebo test for our empirical strategy.

Key words: cost of capital, time consistency, too big to fail, resolution plans, Dodd-Frank

---

Cetorelli: Federal Reserve Bank of New York. Traina (corresponding author): University of Chicago (email: traina@uchicago.edu). The authors thank Anya Kleymenova for comments and feedback. They also thank Michael Blank for excellent research assistance, and Maria Ogneva for sharing code. All errors are the authors' own.

To view the authors' disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr859.html.

1 Introduction

Large financial institutions enjoy implicit government guarantees. Numerous studies have documented these "Too Big to Fail" (TBTF) subsidies, often by comparing the cost of capital for large banks against small banks, or large banks against large corporates.[1] Since governments are effectively subsidizing downside risk, the banks that enjoy TBTF status will have artificially lower costs of capital. At best, these subsidies are simply a transfer from taxpayers to bank investors. At worst, these lower financing costs create incentives for banks to either take on more risk or grow beyond their optimal scale.[2]

Following the 2008 financial crisis, the Dodd-Frank Act implemented resolution plans as a policy solution to the TBTF problem. The Act mandates that each large and complex financial institution files a "living will" that details how they could unwind positions when distressed without significant systemic impact. Effective living wills presumably create the conditions for an orderly liquidation, thus reducing or eliminating altogether the need to bail out banks in the first place. Consequently, the reform should lead to lower TBTF subsidies, and therefore higher costs of capital. However, providing empirical evidence is challenging because clean identification requires: (1) a way to measure a bank's cost of capital; (2) a long enough panel after the regulatory changes; and (3) an appropriately balanced control group.

This paper asks the simple question: Are living wills effective at reducing TBTF guarantees? To answer this question, we test whether living will regulations increase bank financing costs, consistent with a reduction in TBTF subsidies. We first develop a method for measuring a bank's cost of capital using regulatory filings, financial statements, and analyst forecasts, adapting the work of Claus and Thomas (2001) and Gebhardt et al. (2001) to banks. This approach has the advantage of allowing us to reliably measure changes in total financing costs at annual frequencies. We combine this cost of capital measure with information on the staggered treatment of living will regulations. The fact that banks are subject to living will regulations at different times helps rule out confounding causes. Finally, because treated banks are typically much larger and more complex than their untreated peers, we use the synthetic control approach pioneered in Abadie and Gardeazabal (2003) and Cavallo et al. (2013) to better balance our comparison. This technique reweights untreated banks to match the cost of capital levels and trends, as well as important covariates such as size, of their treated peers.

We find that living wills reduce TBTF subsidies. The size of the effect is economically significant at around 22 bps on average, or about 10% of total funding costs. This estimate is slightly smaller than the size of TBTF subsidies estimated by earlier studies.[3] Consistent

[1] Recent examples include Gandhi and Lustig (2015) and Acharya et al. (2016).

[2] Penas and Unal (2004), Dam and Koetter (2012), and Duchin and Sosyura (2014) study these moral hazard issues in detail.

[3] Typical estimates range from 22 bps to 100 bps. See Gropp et al. (2010), Ueda and Di Mauro (2013),

with our narrative, the effect is stronger for banks with higher measures of systemic risk before the regulation's announcement. Multiplying each treatment effect by the size of the treated bank, then summing over the treated sample, we estimate the total size of the subsidy to the industry at about \$42B annually. Since these regulations may also decrease systematic risk by changing bank behavior, which would decrease financing costs, we view our estimates as a lower bound. For perspective, this total is similar in magnitude to the estimated one-time transfer from the Emergency Economic Stabilization Act in Veronesi and Zingales (2010). Consistent with recent evidence that bailouts benefit equity holders, we find that most of the action comes from an increase in the cost of equity capital. In fact, the impact on the cost of deposit capital is statistically indistinguishable from zero, representing a good placebo test of our empirical strategy.

We confirm our interpretation of our findings with supplementary evidence in three different ways. First, checking the size of anticipatory bank responses, we find similar but weaker results on a specification that pools the treatment dates to the time of the first announcement. Second, our results are robust to variations in the synthetic control methodology, particularly in the choice of matching variables. Finally, we test whether living wills led to asset mix or capital structure simplification. We find that banks significantly lower their leverage ratio.

The paper unfolds as follows. In the next section, we review the relevant institution details on TBTF and living wills. Section 3 demonstrates our method of measuring a bank's cost of capital. In section 4, we describe the data sources used for the empirical analysis, and develop our synthetic control approach. Section 5 presents our main results, as well as secondary analysis. We conclude with academic and policy implications in the last section.

2 Institutional Background

This section provides relevant institutional context for both TBTF and living wills. First, it reviews the theory and empirical evidence of TBTF. Implicit government guarantees offer banks a subsidy to financing costs since investors do not have to take on part of negative tail risk. A robust conclusion of the literature is that this guarantee creates misaligned incentives, distorting size and risk choices. Second, this section offers the theoretical arguments of why living wills may alleviate the TBTF problem. These wills aim to reduce the need for government guarantees by mitigating negative spillovers when a bank fails.

Gandhi and Lustig (2015), and Acharya et al. (2016).

2.1 The Problem of Too Big to Fail

TBTF is the principle that large financial institutions may expect government support to prevent their failure. These institutions are sufficiently large, complex, and interconnected that their failure may have disastrous effects on the financial system through negative spillovers. Therefore, faced with risking systemic failure, policymakers will likely aid the failing institution. Moreover, since this policy tradeoff occurs at the point of failure, policymakers cannot credibly commit to not assisting failing institutions even if they so desired. Once a bank failure is imminent, policymakers will optimally commit to avoiding short-run downside risks to themselves, despite possible welfare destruction.[4]

The problem is that investors price these banks as TBTF, and consequently demand artificially lower rates of return on capital. Since the government will cover some of investors' losses in the negative tail of returns, banks care less about downside risk. And while the direct effect of these guarantees is to subsidize debt holders, they can also benefit equity holders by either avoiding resolution uncertainty or financial distress costs as the equity value approaches zero.[5] Consequently, this financing subsidy leads to poor incentives for banks. First, banks will try to become bigger just to acquire the subsidy, creating scale distortions.[6] Second, banks will take advantage of the subsidy once they have it by taking on too much risk.[7]

This economic phenomena grows increasingly relevant as the banking industry grows increasingly consolidated. Cetorelli et al. (2014) document this rise in banking conglomeration. In the 1980s, the ten largest banks held 20% of the total banking industry assets; that number is now above 50%. And this consolidation was not just within commercial banking, but throughout the banking industry. Banks grew into complex conglomerates, adding many subsidiaries spanning the entire spectrum of the financial sector. They no longer solely deal in deposits, mortgages, and simple loans – their business activity bounds have expanded to include broker-dealers, insurance brokers and underwriters, asset managers, and many others. These trends have multiplied the number of banks that might threaten the system's health upon failure.

Fiscal authorities intervened and provided substantial support to distressed financial institutions during the 2008 financial crisis, only reinforcing the belief that governments could not credibly commit to avoiding bailouts if significant conditions of distress materialize again. Indeed, Ueda and Di Mauro (2013) quantify the subsidy for TBTF banks, calculating

[4] For empirical support of this political economy rationale, Brown and Dinç (2011) show that after elections, policymakers are less likely to support failing banks.

[5] Gandhi and Lustig (2015) and Kelly et al. (2016) provide empirical evidence that TBTF subsidies substantially lower costs of equity capital, as well.

[6] Penas and Unal (2004) study bond financing cost effects of bank mergers. They find differentially large benefits when medium-sized banks merge to form a large bank.

[7] Dam and Koetter (2012) use political instruments to show that German banks are more likely to enter distress if they expect a bailout, and Duchin and Sosyura (2014) use Troubled Asset Relief Program data to show that US banks originate riskier loans and shift assets toward riskier securities after getting bailed out.

an increase from 60 bps to 80 bps through the crisis. Acharya et al. (2016) found no change in bond credit spreads in an event window around the passage of the Dodd-Frank Act, which they interpret as indicating no change how markets perceive TBTF as a result of the new law. Gandhi and Lustig (2015) find a size factor in bank stock returns that's related to bank tail risk, which suggests that equity holders of large have some expectation of being bailed out during a crisis; they estimate the overall reduction in the cost of equity capital at a bit below 200 bps. Taken together, these studies suggest the TBTF problem is all too relevant in today's financial markets.

2.2 Living Wills as a Potential Solution

Section 165(d) of the Dodd-Frank Act requires that banks with over $50 billion in assets submit annual resolution plans. The requirement also extends to financial institutions designated by the Financial Stability Oversight Council. Such plans, commonly known as living wills, must describe substantive, detailed strategies to achieve a rapid and orderly resolution when experiencing financial distress or failure. Federal Reserve regulators can approve the living will, or reject until the bank develops a sufficiently robust strategy. Rejections may come with stiff penalties, such as higher capital or liquidity requirements, or forced changes in organizational structure. For example, regulators rejected Wells Fargo's plan in December 2016, and consequently prohibited the bank from establishing international bank entities or acquiring non-bank subsidiaries. Regardless of whether they pass regulatory scrutiny, banks must also disclose a substantial part of their living wills to the public. Hence, the living will treatment is arguably economically meaningful, both de jure and de facto.

Living wills offer a theoretical solution to TBTF by both decreasing systemic externalities and increasing regulatory commitment. From the start, regulators issued specific guidelines so that living wills decrease both the likelihood of distress ("enhancing resiliency") and the systemic externalities should distress materialize ("reducing the impact of a firm's failure").[8] These enhanced resiliency measures attempt to reduce the need for bailouts in the first place. Moreover, if these living wills are credible, then they address the ex-post time inconsistency problem faced by the regulator by decreasing potential agency costs to letting a bank fail. This increase in regulatory commitment has two potential mechanisms: living wills can increase a regulator's bargaining power in navigating a bank's financial distress because there's less uncertainty about hidden costs of failure; and they can increase the political costs to bailing out a bank, since the policymaker has publicly signaled they would not do so.

The regulation had a staggered implementation. Banks at $250 billion or more in total assets produced the first set of plans in July 2012, those with assets at or above $100 billion in

[8]See SR 12-17 / CA 12-14: Consolidated Supervision Framework for Large Financial Institutions at https://www.federalreserve.gov/supervisionreg/srletters/sr1217.htm.

July 2013, and everyone else by December 2013. In terms of legal and regulatory timing, the Dodd-Frank Act became law in July 2010, and the Federal Reserve approved the final rule implementing living wills in October 2011. Qualifying banks therefore understood in advance that they would be subject to living will regulations. However, since the regulation is costly, there is little reason to believe banks meaningfully changed their behavior substantially before implementation. Moreover, as suggested by Acharya et al. (2016), studies that have looked at TBTF usually cannot clearly identify the treated group, because usually there is no explicit list of TBTF firms. Instead, earlier work defines treatment by arbitrary size thresholds or other metrics of potential systematic risk. Our analysis instead has the advantage that the firms and their corresponding years specifically subject to the treatment are clearly identified.

3 Measuring Cost of Capital

We develop our method to measure the weighted average cost of capital (WACC) by combining measures of debt and equity costs with capital structure information. For each bank i at time, we define the WACC as:

$$WACC = \frac{E}{V} * CEC + \frac{B}{V} * CBC + \frac{D}{V} * CDC$$

where E is the book value of equity; B and D are the book values of non-deposit and deposit debt; $V = E + B + D$ is the total book value of the bank; and CEC, CBC, and CDC are the costs of equity, non-deposit debt, and deposit capital, respectively. Note that we use book equity instead of market equity. Regulators require banks to frequently mark their books to market, so these measures are highly correlated. However, book equity offers the advantage of lower idiosyncratic volatility in measurement, as it only reflects the market value of the financial assets on its balance sheet and not the present value of future profits.

We derive estimates of the two components of the cost of debt capital from income statement and balance sheet data. Namely, we calculate the cost of deposit capital as the ratio of total deposit interest expenses from the income statement to total deposits from the balance sheet. Correspondingly, we calculate the cost of non-deposit debt capital as the sum of the remaining interest expenses divided by the end-of-year stock of non-deposit debt. Note that these measurements use book values of debt, and measure expected returns using realized returns. These limitations are less important for debt, particularly short-term debt, since it is less sensitive to information shocks.

Measuring the cost of equity capital is difficult. Traditional methods rely on either using realized returns as a proxy for expected returns, or estimating market risk premia. Neither method is possible to precisely measure the cost of equity capital in even typically long-term horizon panels. Ed Elton highlights this problem in his American Finance Association presidential address, stating simply that “realized returns are a very poor measure of expected returns” (Elton (1999)). Or, as Nobel Laureate Gene Fama quipped in a recent interview, “People like to tell stories about short periods of data, but the reality is that you can’t measure the market premium over periods shorter than an investment lifetime. The 5 percent stock market premium over bills takes about 35 years before it becomes two standard errors from zero."[9]

To overcome this problem, earlier work in accounting and finance developed an ex ante approach using equity prices and analyst earnings forecasts (Claus and Thomas (2001), Gebhardt et al. (2001)). These papers derive the estimated cost of equity capital implicitly as the internal rate of return on equity that equates the current stock price to the present value of all future cash flows to common shareholders. More recent work has shown that these ex ante measures of expected returns demonstrate key empirical properties that were previously unavailable in the ex post measures. In asset pricing, Pástor et al. (2008) and Li et al. (2013) show the usefulness of aggregate ex ante cost of equity capital measures in estimating the risk-return tradeoff and predicting market returns, respectively. In corporate finance, Hann et al. (2013) show the usefulness of individual ex ante cost of equity capital measures in demonstrating how corporate diversification can decrease countercyclical deadweight losses through coinsurance.

In theory, the implied cost of equity capital approach requires both the stock price and the infinite series of expected future cash flows for each bank:

$$P = \sum_{t=1}^{\infty} \frac{E[CF_t]}{(1+CEC)^t}$$

where P is the price today, $E[CF_t]$ is the expected future cash flow given today’s information, and CEC is the cost of equity capital given today’s information.

Taking these objects to the data, equity prices are observable and analyst earnings forecasts offer a proxy for expected future cash flows. However, analyst forecasts are also narrow in their horizon. In practice, we follow earlier studies in overcoming this problem by specifying a terminal value of 3 years to the forecasts. After this terminal value, we assume that returns on equity ROE linearly revert to their 3-digit NAICS industry median over a 9-year horizon. Beyond that horizon, we use an infinite annuity of residual returns. Given these estimates of ROE, we impose clean-surplus accounting to construct estimates of future book values of equity per share: $BPS_t = BPS_{t-1} + (1-DPR) * ROE_t * BPS_{t-1}$, where BPS_t is the book value of equity per share at time t, and DPR is the dividend payout ratio. Altogether, we numerically solve for the CEC that rationalizes:

$$P = BPS_0 + \sum \frac{ROE_t - CEC}{(1+CEC)^t} BPS_{t-1} + \frac{ROE_T}{CEC(1+CEC)^T} BPS_{T-1}$$

[9]Source: https://www.minneapolisfed.org/publications/the-region/interview-with-eugene-fama.

Empirical studies of the cost of capital outside of banking use market returns on aggregate bond indices, such as the Barclays Capital Aggregate Bond Index, as a proxy for the cost of debt capital (Hann et al. (2013)). One advantage of using accounting data is that it yields bank-specific measures, thus recognizing heterogeneity in the cross-section of institutions in their ability to raise debt liabilities. It also recognizes the starkly different costs of financing for insured deposits relative to other forms of debt finance, which offers a clean placebo test of our empirical strategy discussed in section 5. Finally, accounting data also allow us to avoid negative costs of debt capital, which occur in one year of our sample when using bond indices. We have confirmed that despite expected time-series differences between the aggregate bond index and the accounting metric of cost of debt, the overall $WACC$ for our sample of banks follows very similar time-series trends irrespective of which approach we use.

4 Empirical Strategy

In this section, we describe the data used to measure costs of capital as derived above, and outline our synthetic control approach.

4.1 Data and Sample

There are three key data sources for testing the impact of living wills on banks' cost of capital. Notably, to build our empirical measures of the cost of capital, we need to combine regulatory filings, financial statements, and analyst forecasts for all available publicly-listed US banks. Given the timing of living will regulations and the financial crisis, we select 2009 - 2016 as the sample period for the analysis. To avoid results driven by banks entering or exiting the sample, we require that banks be in the dataset for the full panel. This filter affects less than 10% of our treatment group.

Bank regulatory filings contain information on bank balance sheets and income statements. This information allows us to measure capital structure, debt stocks and interest expenses, and other relevant covariates. We get these fundamentals from the FR Y-9C Consolidated Financial Statements for Bank Holding Companies. This regulatory report collects quarterly financial data from banks with total consolidated assets exceeding $1B. From this form, we use information on a bank's total assets, risk-weighted assets, total deposits (domestic and foreign, interest-bearing and non-interest bearing), interest expense on total deposits, overall interest income and expense, overall noninterest income, and total liabilities (assets minus equity).

For more firm fundamentals, we collect financial statement information from the population of firms in the annual Compustat file provided by Wharton Research Data Services (WRDS). Specifically, we collect information on common dividends paid, income before extraordinary items, common equity, share price, and number of shares outstanding. We restrict our sample to banks with (1) above \$20M in revenue, (2) positive book equity, and (3) non-negative dividends. We also drop observations for which the bank acquired another firm worth over 5% of acquirer assets. Following the ex ante cost of capital literature, we replace income as 6% of total assets in years that have negative or missing income. Given this measure of income, we calculate the dividend payout ratio as dividends over income. If this ratio is below 0 or above 1, we set it to 0 or 1, respectively.

Analyst forecasts allow us to measure the cost of equity capital. We gather this information from the Thomson Reuters I/B/E/S summary file, also provided by WRDS, as of the end of June of each year. We use stock price per share and forecasts of both earnings and long-term earnings growth. We require that each observation has non-missing 1 and 2 year consensus earnings forecasts. We use median consensus forecasts to proxy for the market's future earnings expectations. We use 3 year forecasts for future earnings per share if they are available; otherwise, if the rate is between 2% and 100%, we impute the forecast by applying the long-term growth rate to earlier forecasts.

To merge these datasets, we rely on the Federal Reserve RSSD identifier to PERMCO identifier link provided by the Federal Reserve Bank of New York, as well as the PERMCO-CUSIP link provided by WRDS. The resulting dataset is an unbalanced panel without gaps, containing 1,272 bank-years across 159 unique banks. Of these banks, 23 are treated at some point in the sample. These are the largest public banks and therefore represent the most critical piece of the credit intermediation industry. Moreover, there is a range of bank size, with the banks spanning in total assets from about \$900M to over \$2.5T.

We also collect data on two systemic risk measures for the treated banks in our sample. The first one is SRISK (with an imposed 40% market decline), developed in Brownlees and Engle (2016) and taken from the website of the Volatility Institute at NYU Stern. The second one is 95% CoVaR, developed in Adrian and Brunnermeier (2016) and taken from that paper's online appendix.

Figure 1 shows the time series average WACC for treated and untreated banks from 2001 to 2016. The trend is in keeping with priors, displaying a common decreasing path through the first part of the earlier decade and a common increasing path to the onset of the financial crisis. The lines then display an expected temporary reversal during the 2008 financial crisis, and then a continuation of the downward trend, with levels bottoming out toward the end of the sample period.

Figure 1: Weighted Average Cost of Capital by Treatment

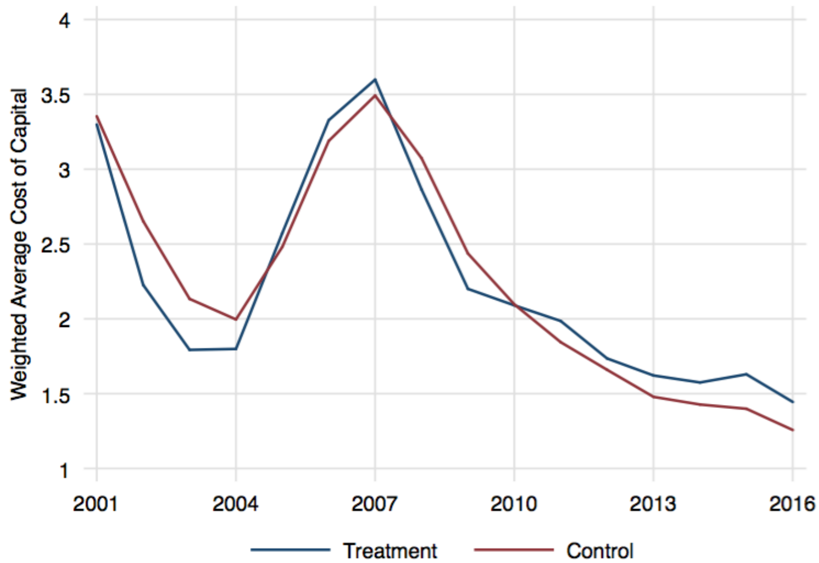

The picture does not allow us to detect significant differences between the two groups of banks, although the WACC for the treated banks declines at a lower pace starting in 2012. The difference begins at the same time as the living will regulation announcement, but the raw data do not lend themselves to a clean check on the impact of the new policy. It does not take into account its staggered introduction (thus commingling temporarily treated and controlled firms), nor that the set of public banks that were not treated may just not offer a proper benchmark – they are significantly different in fundamental characteristics, chiefly size, which correlate with the cost of capital itself.

The trend exhibited in the WACC may also overlook the fact that while in aggregate cost of capital does not seem seems to diverge across banks, the differences in its components may still matter. Figure 2 focuses on the equity component of the cost of capital. The chart still indicates a commonality in trend between treated and untreated banks, although the cost of equity capital for the treated bank increases more during the crisis, and the gap remains roughly constant in the following years although slightly widening starting in 2012.

Figure 2: Cost of Equity Capital by Treatment

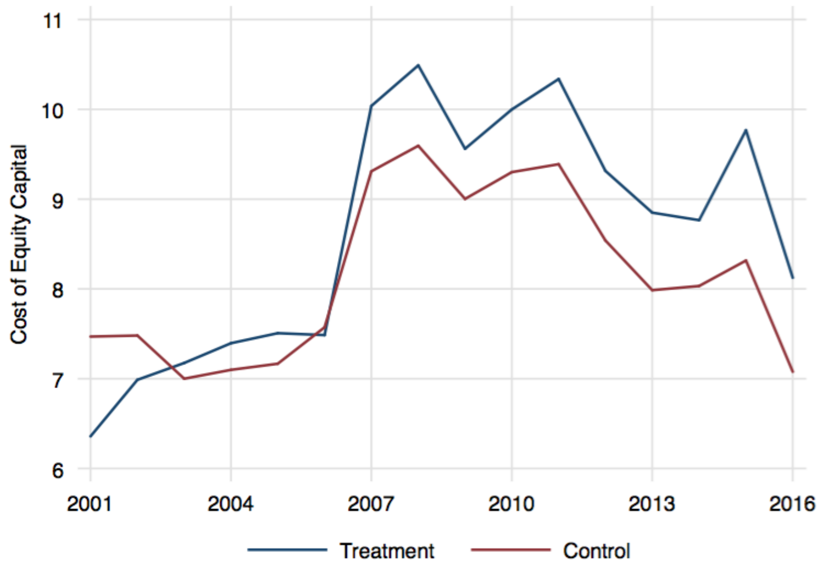

The higher cost overall is what one might have expected for a number of concomitant factors that have likely affected those banks that later become subject to the living will regulation and that should also have an impact on the cost of equity capital. For instance, the difference could be attributed to an overall enhancement in monitoring and regulation of the larger banking institutions, such as stress testing or enhanced liquidity and capital rules. These policies that may have masked changes in the cost of equity capital. Understanding to what extent the living will regulation may have been a significant factor contributing to the differential cost of equity in the ending part of the sample thus requires a more rigorous approach, which we discuss next.

4.2 Synthetic Control Approach

We use a synthetic control approach to mitigate concerns that treated banks are substantially different from control banks. In particular, treated banks tend to be much larger and more complex by qualification, which makes the typical difference-in-differences assumptions unlikely to hold. We alleviate the problem by reweighting the control group banks to match bank characteristics as well as cost of capital trends. The synthetic control methodology chooses covariate weights to minimize the pre-treatment prediction error of the outcome variable, and then rebalances control observations according to these weights (Abadie and

Gardeazabal (2003), Cavallo et al. (2013).[10] Moreover, it does this procedure for each bank, allowing us to identify treatment effects at the individual bank level. In our primary estimates, we target the WACC, CEC, leverage, risk-weighted assets ratio, market capitalization, book-to-market ratio, and operating profitability. We also check the robustness of our results by targeting WACC, CEC, leverage, and market capitalization only.

Inference for synthetic control approaches uses a distribution of placebo estimates. For each bank's treatment effect, we generate placebo treatments for the same period but on all the control banks. The corresponding p-values compare the share of such placebo effects that are at least as large as the actual treatment effect. We also report standardized p-values, which weights the treatment estimates by their prediction error. Intuitively, better matches offer more meaningful evidence for inference, and are therefore weighted higher.

Table 1 shows the covariate balance for the treated and control banks, before and after the synthetic match. We include all covariates used to produce our estimates. The Treatment column displays the pretreatment covariate average across all 23 living will banks from 2009 to 2011; the Unweighted Control column gives this average for all potential control banks in the sample; and the Weighted Control column gives this average for the synthetic control group.

Table 1: Covariate Balance and the Synthetic Control Group

Statistic	Treatment	Unweighted Control	Weighted Control
WACC	2.00	2.04	1.99
CEC	9.98	9.03	8.48
Leverage	0.89	0.89	0.89
RWA Ratio	0.72	0.70	0.60
Market Capitalization	40.56	0.83	8.94
Book-to-Market Ratio	1.17	1.09	1.00
Operating Profitability	0.28	0.27	0.28

As expected, the primary economic change in sample characteristics from unweighted to weighted controls is an increase in the average size. Yet this improvement comes at a price, as the CEC, RWA Ratio, and Book-to-Market ratio all decrease slightly from the treated sample. Other than the outcome variable itself, the synthetic control methodology will not necessarily make the synthetic mean for a given covariate closer to the living will mean. If a given covariate is not important in predicting the pre-treatment variation in the outcome mean, then it will receive a low weight in the synthetic control procedure, and thus it will not necessarily converge to the living will mean via the synthetic match. Hence, the covariates that approach the living will means are especially important for predicting cost of capital over 2009 to 2011.

[10]Note that since the outcome variables change between different tested outcomes, so too do the covariate weights and hence the synthetic control group weights.

5 The Effect of Living Wills on the Cost of Capital

This section collects our results on the cost of capital, as well as supplementary evidence that supports our interpretation of a reduction in the TBTF subsidy. Living wills increase the cost of capital, consistent with reducing the TBTF subsidy. This result is found on average but it is also significant in the cross-section. Moreover, the estimated individual effects display strong correlations with ex ante metrics of systemic risk, corroborating the prior on the expected effects of the policy intervention. This effect is economically large: a back-of-the-envelope calculation implies a total subsidy size of about $42B annually.

5.1 Increases in the Cost of Capital

Figure 3 shows the results of the synthetic control analysis for total cost of capital. In this and all later figures, we apply the treatment on the basis of the staggered implementation of the policy. Hence, the treated and control firm composition varies for each period before the years of the policy implementation ($t = 0$). As expected with the synthetic control methodology, we get a close parallel WACC trend between treated and control groups in the pre-treatment period, as shown in the top left panel. However, the estimation results show an economically significant and persistent divergence in the post-treatment period. The top right panel shows the difference between treatment and control, indicating an overall cost of capital for treated banks of at least 22 bps, or about 10% of total funding costs. The effect size is increasing in the years after the treatment, consistent with a gradual reinforcement of the policy. The bottom panels show that these effects are statistically significant regardless of standardization. Each dot corresponds to an estimate's time-specific p-value, with a horizontal line drawn at the 5% level for reference.

Figure 3: The Effect of Living Wills on the Weighted Average Cost of Capital

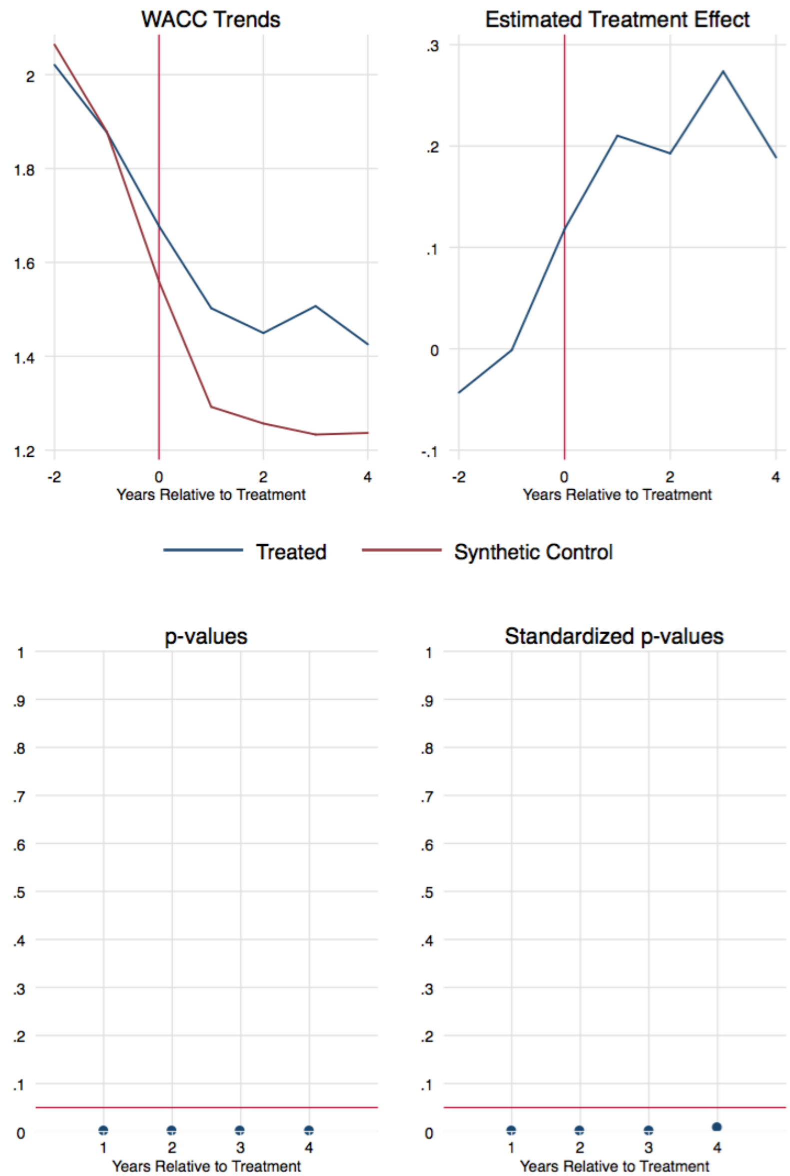

Since living wills attempt to remove TBTF status by increasing bank resiliency while simultaneously removing TBTF guarantees, their net effect on a bank's cost of capital is not straightforward. While the impact on the likelihood of a bailout should lead to an increase in the cost of capital, the complementary impact on the likelihood of default should lead to a decrease in the cost of capital. Hence, our evidence of an increase in the cost of capital post treatment is likely a lower bound of the overall impact on the TBTF subsidy.

5.2 Stronger Effects in Systemically Important Banks

A significant advantage of the synthetic control methodology is that it allows us to estimate a separate treatment effect for each bank. We use this advantage to investigate heterogeneity in our results, and in particular what bank characteristics correlate with our treatment. If the estimated changes in the cost of capital are the result of smaller TBTF subsidies, then the treatment effects will have a positive relationship with ex ante measures of the bank's systemic risk before the policy implementation.

Figure 4 shows binscatters of the estimated effects on WACC against two measures of systemic risk taken in 2011, before the regulation's announcement. The vertical axis displays our estimated measure in percentage points, and the horizontal axis displays the measure of systemic risk. The left panel shows the relationship between our estimated treatment effects and SRISK, developed by Brownlees and Engle (2016), and the right panel shows the relationship with CoVaR, developed by Adrian and Brunnermeier (2016).

Figure 4: The Effect of Living Wills by Systemic Risk

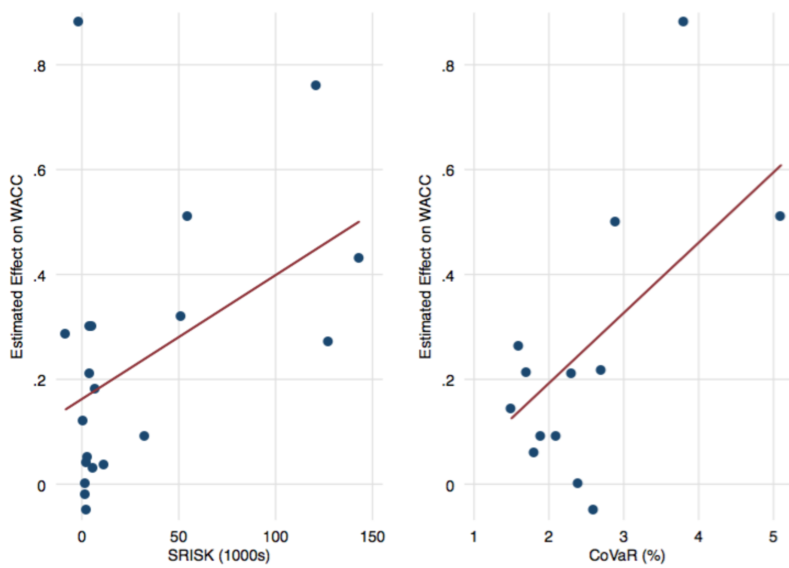

Consistent with our narrative, we find that the cost of capital increases the most for systemically important banks. Both statistical relationships are positive and statistically significant at the 5% level, with R^2 values of about 0.2.

5.3 Large Implied Subsidies

Is the estimated increase in banks' cost of capital economically significant? The 22 bps differential increase is in the ball park of other TBTF estimates provided in the literature. We can however dig a bit deeper. Column 3 of Table 2 shows each bank's individual treatment effect. These estimates are computed by simply taking the difference between the post-treatment mean of the outcome for the living will bank and the mean of the bank's specific synthetic control – their average of 22 bps produces the estimates plotted in the above figures. Banks in the table are sorted by decreasing order of size as of 2011, listed in Column 2.

Table 2: The Effect of Living Wills on the Cost of Capital by Bank

Bank Name	Assets ($B)	Estimate (bps)	Subsidy ($M/yr)
JPMorgan Chase	2,266	27	6,104
Bank of America	2,137	43	9,205
Citigroup	1,874	76	14,227
Wells Fargo	1,314	9	1,151
Goldman Sachs	924	51	4,747
Morgan Stanley	750	32	2,402
U.S. Bancorp	340	7	235
The Bank of New York Mellon	326	1	45
PNC Financial Services	271	21	582
State Street Corporation	216	3	76
Capital One	206	30	620
SunTrust Banks	177	6	109
BB&T	175	30	523
American Express	152	50	759
Regions Financial Corporation	127	18	225
Fifth Third Bank	116	6	64
Northern Trust	100	-5	-47
KeyBank	89	5	44
M&T Bank	78	12	97
Discover Financial	69	88	610
Comerica	61	0	-2
Huntington Bancshares	54	-2	-14
Zions Bancorporation	53	2	13
			41,775

Inspecting the table reveals two important observations about the estimates. First, the range of estimates is consistently of the size of earlier estimated TBTF subsidies. The smallest effects on WACC are roughly 0, and the largest effects are roughly 90 bps. Second, almost all banks subject to living wills experience an estimated increase in the cost of capital after the policy implementation – this average treatment effect is not driven by just a few outliers. As prior perceptions of TBTF were likely present for these living will banks, this widespread effect is consistent with our TBTF hypothesis.

In order to get a better sense of the economic importance of these results, we use the WACC estimates and the corresponding bank asset size to derive implied dollar amount effects. Column 4 in Table 2 display these implied subsidies. We find large economic effects – Wells Fargo alone saw a total increase in cost of capital expenses of \$14B. The effects are less pronounced for the smaller institutitions focused on commercial banking, but stay large for those focused on other financial services. For example, Discover Financial has only \$69B in assets, but enjoyed an annual subsidy of \$610M.

We sum down the column to get a total industry-wide impact of the policy reform of about \$42B per year. We can benchmark this result against the size of the Emergency Economic Stabilization Act, the US government bailout implemented during the 2008 financial crisis. Veronesi and Zingales (2010) estimate that the policy was "…a massive transfer of resources from the taxpayers to the banking sector", with an increase in bank financial claims of \$107B and a corresponding net cost to taxpayers ranging between \$25B and \$47B. The increase in bank value, the authors assert, was in large part driven by a decrease in the perceived risk of bankruptcy. Note that this bailout was a unitary event, whereas our implied subsidy accrues annually.

5.4 Treatment by Types of Capital

Finally, we assess the relative impact on the individual components of the cost of capital: equity, non-deposit debt, and deposit debt. To the extent that the policy is effective in either decreasing resolution uncertainty or the costs of financial distress, we should see a higher required compensation for equity holders. A similar argument holds for non-deposit debt, which also has the mechanism of being the direct target of past bailouts. However, we note that for our particular measurement, this channel is not as clear-cut: much of the non-deposit debt on bank balance sheets are safe, short-term, deposit-like instruments. At the largest banks, these are largely repurchase agreements and collateralized trading liabilities. Consequently, they are less sensitive to a change in perceived tail risk. The prior on the effect on holders of deposits is small or non-existent – these are already protected by deposit insurance. The cost of deposit capital should therefore act as a natural placebo test for the overall analysis, since we should not expect to find any signficant difference of the policy between treated and control banks.

We show the results of decomposed by type of capital in the next set of figures. Figure 5 displays the estimated impact of the policy on the cost of equity capital. The cost of equity capital appears much higher for the treated banks in the post-treatment years. The difference is economically quite large, increasing over the post-treatment period and averaging about 80 bps. It is also statistically significant as seen in the bottom panels.

Figure 5: The Effect of Living Wills on the Cost of Equity Capital

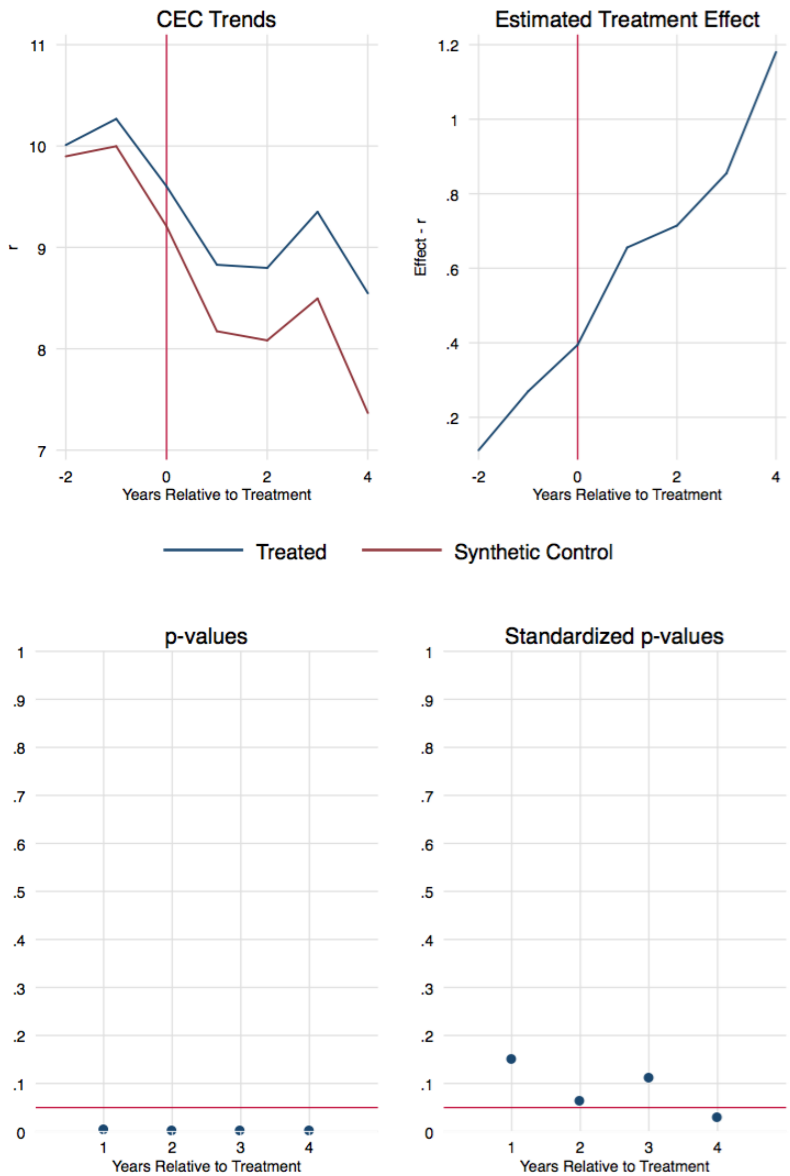

Figures 6 and 7 presents the corresponding analysis applied to the debt components of the total cost of capital. Cost of non-deposit debt is higher post-treatment for the treated banks. However, the effect is not estimated precisely, and the effect sizes are quite small, thus limiting the overall statistical inference. Cost of deposit debt is indistinguishable between treated and control in the post-treated period. Any detectable difference is also very small.

Figure 6: The Effect of Living Wills on the Cost of Non-Deposit Debt Capital

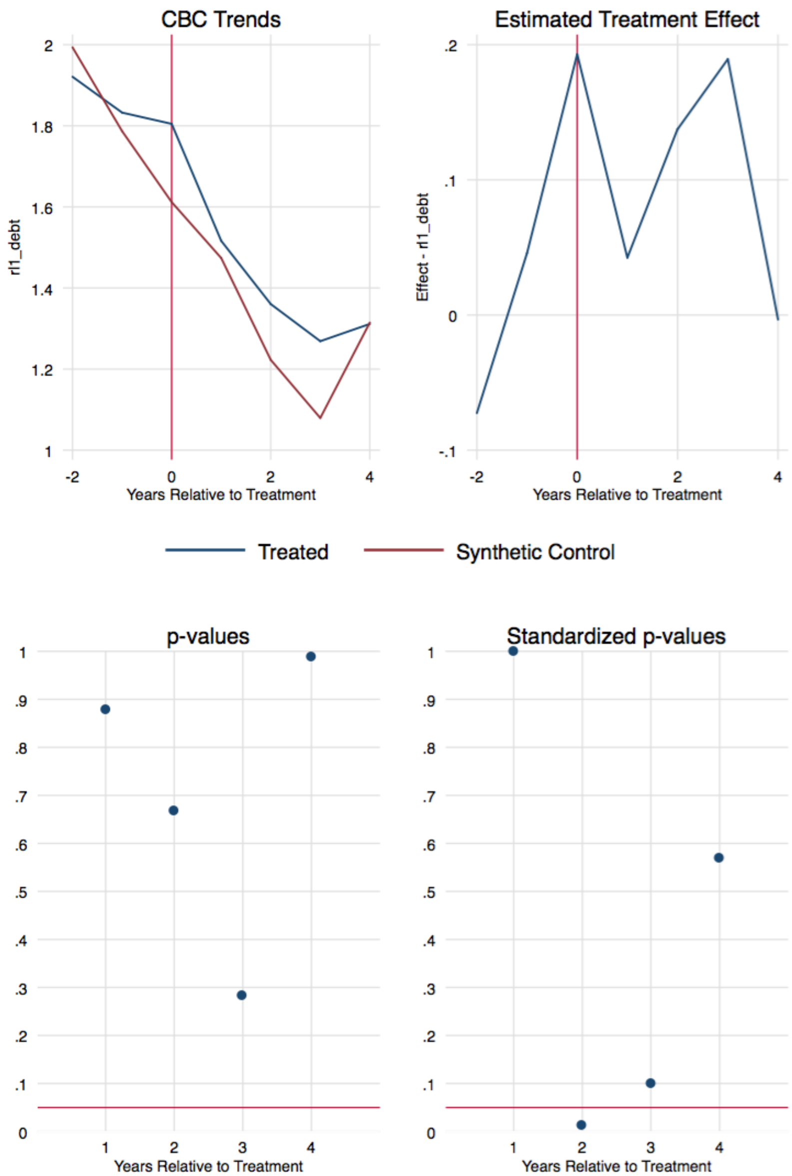

Figure 7: The Effect of Living Wills on the Cost of Deposit Capital

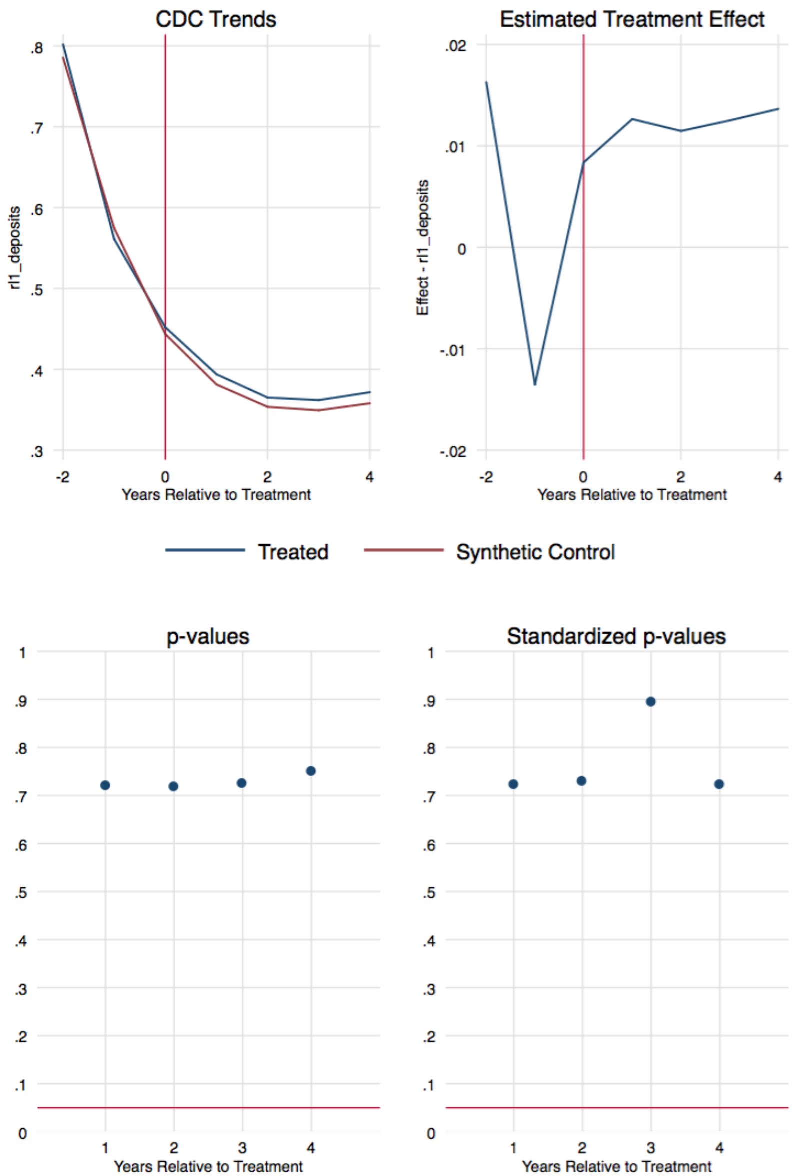

5.5 Robustness Tests

For robustness, we confirm our results hold under two alternative specifications. Figure 8 in the Appendix shows the effect of living wills on the cost of capital under an alternative belief in treatment timing – that banks and investors fully anticipate and react to living will implementation at the time of the announcement. Figure 9 in the Appendix shows the original specification but with only cost of capital, size, and leverage targets in the matching covariates. Both variations do not substantially alter the effect, both in terms of statistical and economic size.

We also test whether banks adapt to the new policy with significant capital structure changes. Figure 10 in the Appendix extends the analysis to estimated the effect on leverage ratios. Treated banks show a lower leverage level in the post-treatment period. The difference is statistically significant and economically relevant, reducing leverage by about 1.5% toward the end of the post-treatment period. To the extent that excessive leverage is a proxy for risk-taking behavior, its reduction can be interpreted as consistent with the effectiveness of a policy that has an impact not only on the likelihood of being bailed out given default, but also on the likelihood of experiencing distress to begin with. The result pointing at a reduced leverage is in line with the mixed findings on the debt component of the cost of capital.

6 Implications

Our evidence suggests living wills are an effective policy tools for banks. Moreover, it also suggests similar resolution plans are a good policy tool for large companies generally, particularly when systemic risk is a concern.

Fiscal authorities across multiple jurisdictions implemented massive intervention policies during the 2008 financial crisis, aimed at preventing the default of the largest and most complex financial institutions. The rationale for such policies was the recognized inability to otherwise go ahead with standard bankruptcy and the justified concerns for the systemic externalities that would have likely stemmed from such distress events. While arguably successful in their immediate goal, these policies may have reinforced the market belief that the government will never let large and complex banks fail, thus strengthening the value of any implicit subsidy and the associated distortion in firm and market conduct. A proposed solution to this problem, forcefully voiced in the policy arena, has been to just do away with large and complex banking institutions altogether. Recent commentators have reasoned that curbing size and imposing severe regulatory constraints to prevent complexity buildups as the only practical solution, given the lack of a credible commitment by government authorities to implement ex post bail out interventions. Congress has sought a compromising approach by creating mandatory, annual resolution plans for the largest banks. The law's intent and the associated regulations is to set up the conditions through which an ex post policy of not bailing out banks could be credibly implemented.

Our study provides empirical support to the policy approach followed by the US authorities, or at least evidence that markets have internalized an economically significant reduction in the cost of capital subsidy. We find that living will regulations increase a bank's weight average cost of capital by about 22 bps, consistent with a reduction in TBTF guarantees. This decrease is differentially concentrated in changes to the cost of equity. This estimate exploits both the staggered treatment of the regulation's implementation and a reweighting of untreated banks to match their treated peers. Alternative mechanisms such as reduced resolution uncertainty or financial distress would only bias our estimate down, suggesting our effect is a lower bound. Indeed, we do see treated banks cut their leverage ratios despite the higher costs of equity capital. Moreover, banks also seem to simplify their organizational structure as evidence by a decrease in subsidiary count. More generally, since the very mechanism by which living wills may reduce TBTF subsidies is by decreasing systemic risk directly, our results shed light on the more general problem of negative spillovers from firm failure.

References

Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the basque country. *American Economic Review 93*(1), 113–132.

Acharya, V. V., D. Anginer, and A. J. Warburton (2016). The end of market discipline? investor expectations of implicit government guarantees.

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

Brown, C. O. and I. S. Dinç (2011). Too many to fail? evidence of regulatory forbearance when the banking sector is weak. *The Review of Financial Studies 24*(4), 1378–1405.

Brownlees, C. and R. F. Engle (2016). Srisk: A conditional capital shortfall measure of systemic risk. *The Review of Financial Studies 30*(1), 48–79.

Cavallo, E., S. Galiani, I. Noy, and J. Pantano (2013). Catastrophic natural disasters and economic growth. *Review of Economics and Statistics 95*(5), 1549–1561.

Cetorelli, N., J. McAndrews, and J. Traina (2014). Evolution in bank complexity. *Economic Policy Review* (Dec), 85–106.

Claus, J. and J. Thomas (2001). Equity premia as low as three percent? evidence from analysts' earnings forecasts for domestic and international stock markets. *The Journal of Finance 56*(5), 1629–1666.

Dam, L. and M. Koetter (2012). Bank bailouts and moral hazard: Evidence from germany. *The Review of Financial Studies 25*(8), 2343–2380.

Duchin, R. and D. Sosyura (2014). Safer ratios, riskier portfolios: Banks' response to government aid. *Journal of Financial Economics 113*(1), 1–28.

Elton, E. J. (1999). Presidential address: expected return, realized return, and asset pricing tests. *The Journal of Finance 54*(4), 1199–1220.

Gandhi, P. and H. Lustig (2015). Size anomalies in us bank stock returns. *The Journal of Finance 70*(2), 733–768.

Gebhardt, W. R., C. Lee, and B. Swaminathan (2001). Toward an implied cost of capital. *Journal of Accounting Research 39*(1), 135–176.

Gropp, R., H. Hakenes, and I. Schnabel (2010). Competition, risk-shifting, and public bail-out policies. *The Review of Financial Studies 24*(6), 2084–2120.

Hann, R. N., M. Ogneva, and O. Ozbas (2013). Corporate diversification and the cost of capital. *The Journal of Finance 68*(5), 1961–1999.

Kelly, B., H. Lustig, and S. Van Nieuwerburgh (2016). Too-systemic-to-fail: What option markets imply about sector-wide government guarantees. *American Economic Review 106*(6), 1278–1319.

Li, Y., D. T. Ng, and B. Swaminathan (2013). Predicting market returns using aggregate implied cost of capital. *Journal of Financial Economics 110*(2), 419–436.

Pástor, L., M. Sinha, and B. Swaminathan (2008). Estimating the intertemporal risk–return tradeoff using the implied cost of capital. *The Journal of Finance 63*(6), 2859–2897.

Penas, M. F. and H. Unal (2004). Gains in bank mergers: Evidence from the bond markets. *Journal of Financial Economics 74*(1), 149–179.

Ueda, K. and B. W. Di Mauro (2013). Quantifying structural subsidy values for systemically important financial institutions. *Journal of Banking & Finance 37*(10), 3830–3842.

Veronesi, P. and L. Zingales (2010). Paulson's gift. *Journal of Financial Economics 97*(3), 339–368.

Appendix

Figure 8: The Effect of Living Wills on the WACC, Announcement Year

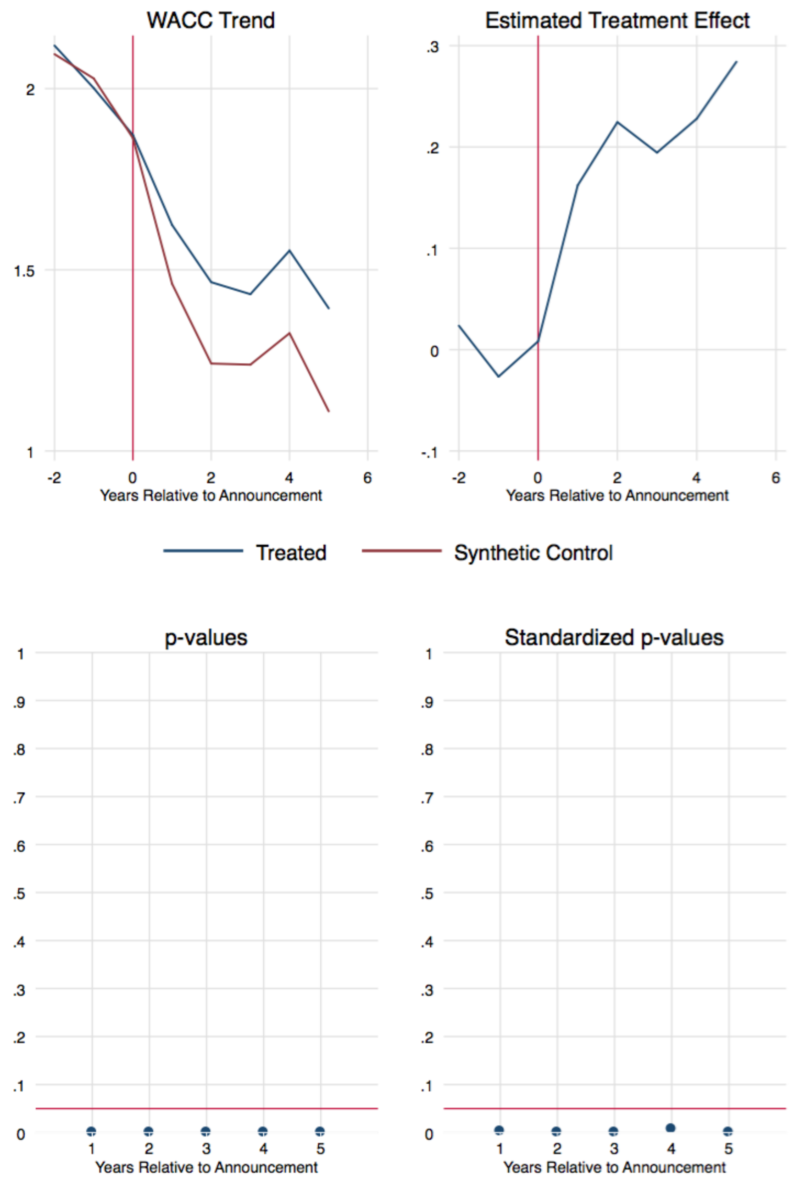

Figure 9: The Effect of Living Wills on the WACC, Alternative Covariates

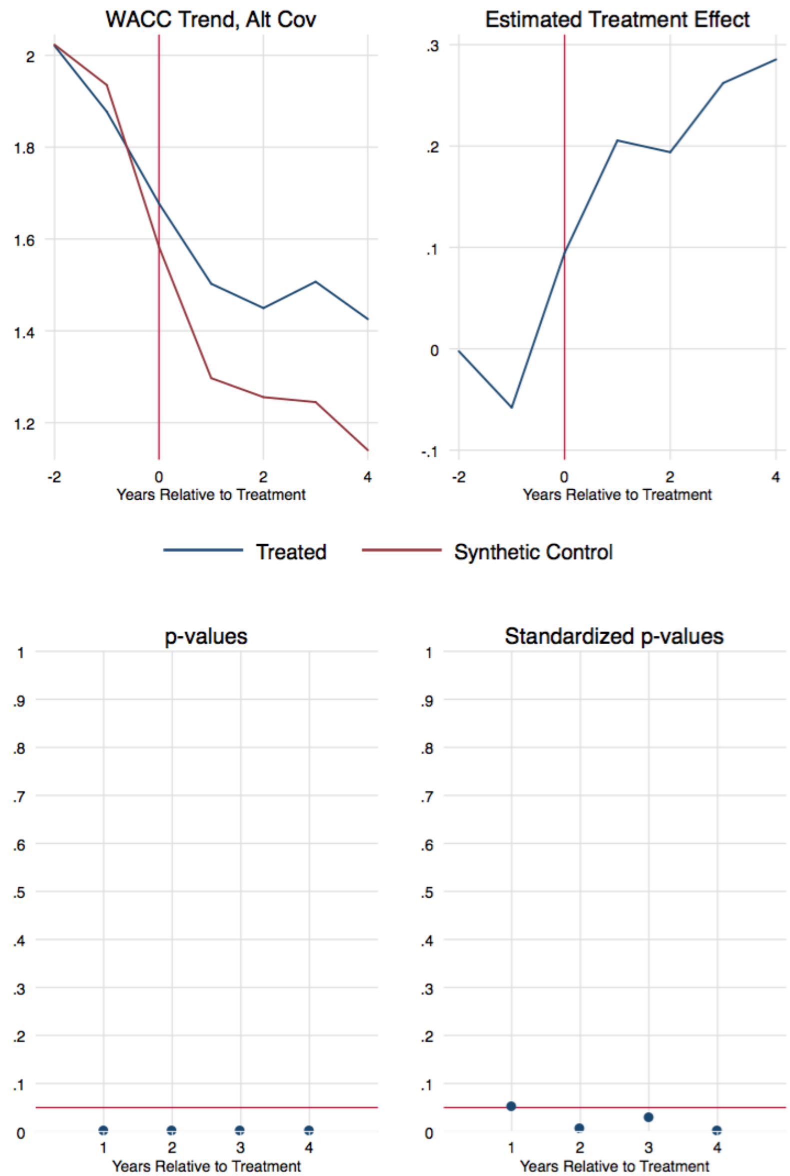

Figure 10: The Effect of Living Wills on Leverage