DANMARKS
NATIONALBANK

WORKING PAPERS

February 2017 | No. 110

INCORPORATING FUNDING COSTS IN TOP-DOWN STRESS TESTS

Søren Korsgaard
Danmarks Nationalbank

DANMARKS
NATIONALBANK

Please direct any enquiries to
Danmarks Nationalbank,
Communications,
Havnegade 5,
DK-1093 Copenhagen K
Denmark
E-mail: kommunikation@nationalbanken.dk

Henvendelser kan rettes til:
Danmarks Nationalbank,
Kommunikation,
Havnegade 5,
1093 København K.
E-mail: kommunikation@nationalbanken.dk

ISSN (online) 1602-1193

DANMARKS
NATIONALBANK
WORKING PAPERS

INCORPORATING FUNDING COSTS IN TOP-DOWN STRESS TESTS

Contact for this working paper:

Søren Korsgaard
Danmarks Nationalbank
sk@nationalbanken.dk

RESUME

Centralbanker og tilsynsmyndigheder foretager regelmæssigt stresstest af banksektoren. I takt med at tab akkumuleres i stressscenarier, svækkes bankernes solvens, og de skal betale en højere rente på deres gæld. Jeg undersøger en række risikomål, som er udledt fra Merton-lignende modeller, og undersøger sammenhængen mellem disse risikomål og bankers finansieringsomkostninger. Endelig beskriver jeg en metode til at indarbejde stigninger i fundingomkostninger i en "top-down" stresstest.

ABSTRACT

Central banks and supervisory authorities regularly conduct stress tests of banks. As losses accumulate in stress scenarios, banks' equity position worsens, and they must pay higher interest rates to retain funding. I explore how variations of Merton-type models can be used to measure bank risk, and then examine the link between various risk measures and funding costs. Finally, I outline a method for incorporating funding cost increases into top-down stress tests.

KEY WORDS

Stress testing, Funding costs, Risk measurement

JEL CLASSIFICATION

G12, G21, G31, G38

Incorporating Funding Costs in Top-down Stress Tests

Søren Korsgaard*, Danmarks Nationalbank

February, 2017

Abstract

Central banks and supervisory authorities regularly conduct stress tests of banks. As losses accumulate in stress scenarios, banks' equity position worsens, and they must pay higher interest rates to retain funding. I explore how variations of Merton-type models can be used to measure bank risk, and then examine the link between various risk measures and funding costs. Finally, I outline a method for incorporating funding cost increases into top-down stress tests.

Keywords: Stress testing; Funding costs; Risk measurement

JEL-classification: G12; G21; G31; G38

*The author would like to thank Andreas Brøgger, Kristian Loft Kristiansen, Søren Hesel, Thomas Sangill and other colleagues at Danmarks Nationalbank for valuable feedback. E-mail: sk@nationalbanken.dk

1. Introduction

Since the financial crisis, central banks and supervisory authorities have sought to improve their stress tests of financial institutions. One approach has been to conduct highly detailed "bottom-up" stress tests in cooperation with the financial institutions themselves. These are expensive and time-consuming to conduct, and it is difficult to incorporate e.g. interlinkages among institutions in such tests. Bottom-up stress tests are therefore complemented by simpler "top-down" stress tests which mainly use public data and relatively simple econometric tools to forecast the evolution of bank capital in stress scenarios (e.g. Flannery, Hirtle, and Kovner, 2016). One line of work has sought to make these stress tests "more macroprudential" (BIS, 2015), and part of this workstream has been to better capture the linkage between solvency and funding costs: As a bank's equity capital deteriorates under stress, its funding cost presumably increases. The question is: by how much?

At first glance, there seems to be only a weak relationship between solvency and funding costs. Aymanns, Caceres, Daniel, and Schumacher (2016) find that a 1 percentage point reduction in solvency increases average funding costs by 0.02 percentage points. For wholesale funding, the figure is 0.04 percentage points. These are averages, and the effects are found to be stronger during times of crisis, and there is also evidence of nonlinearities. In particular, low solvency banks are more affected by decreases in solvency than are high solvency banks.

There are a number of reasons why identifying the actual effects of solvency on funding costs may be difficult. Weak banks, for example, may not have access to wholesale funding markets and therefore have to rely on deposits. As many deposits are covered by deposit guarantee schemes, this type of funding is relatively insensitive to changes in solvency. Another difficulty is that banks with riskier assets for precautionary reasons choose to hold more equity capital. To the extent that the riskiness of a bank's assets is imperfectly measured by risk weights, it might be the case that higher solvency banks are, in fact, relatively risky. Some stress test results support such a story. Flannery, Hirtle, and Kovner (2016) find that more solvent banks, as measured by regulatory capital ratios, experience greater losses in US stress tests. Higher risk weights, in contrast, do not predict losses.

In this paper, I consider a number of market-based measures of both bank solvency (or risk) and funding cost, and show a strong negative relationship between solvency and funding cost. The market-based measures I consider are all variations of the distance-to-default (DD) measure from the Merton (1974) model. Default distances are known to be important covariates in default prediction (Bharath and Shumway, 2008; Lando, 2004) and are robust to model misspecifications (Jessen and Lando, 2015). When examining the variations of the distance-to-default, I also discuss what the various models can teach us, qualitatively, about the potential pitfalls of using particular risk measures. One feature of many of these models is that they must be solved numerically. That is, unknown parameters such as the value and volatility of a firm's assets must be numerically estimated based on observations of equity values and estimates of equity volatilities. Interestingly, Bharath and Shumway (2008) show that a "naive" and easy-to-compute default distance based on a Merton model performs as well, if not better, in default prediction as the actual default distance. I similarly construct a naive default distance, which draws on the qualitative insights from extensions to the standard Merton model, and show that it too performs as well as default distances derived from more complicated models.

Finally, I discuss how a market-based solvency measure can be used to project funding cost increases in a balance-sheet based stress test. This evidently presents a challenge: In a stress test, one observes the progressive decline in regulatory capital ratios, not in market values. Furthermore, some banks are not traded and therefore have no market data to condition on in the first place. I outline a procedure for dealing with both problems.

The paper takes its outset in the Merton model. The idea of using Merton-type models to project funding costs in stress tests is explored in BIS (2015). I look at both simplifications and extensions to the Merton model. One simplification, suggested by Atkeson, Eisfeldt, and Weill (2014), is to use the inverse equity volatility as a risk measure. In fact, this is equivalent to the distance-to-default when the "delta", the derivative of the equity value with respect to the underlying asset value, is one in the Merton model. As extensions, I consider models where default can occur at any time, similar to Black and Cox (1976), where the default boundary is adapted to better reflect bank solvency regulation, similar to Chan-Lau and Sy (2006), and finally where the nature of underlying stochastic processes are adapted to better reflect the nature of the underlying bank assets, as in Nagel and Purnanandam (2015).

The paper contributes to the literature on the relationship between bank solvency and funding costs (e.g. Aymanns, Caceres, Daniel, and Schumacher, 2016; Anneart, De Ceuster, Van Roy, and Vespro, 2013; Hasan, Liu, and Zhang, 2016). Another closely related paper is Schmitz, Sigmund, and Valderrama (2016) which also discuss applications to stress testing. Incidentally, these authors identify a very strong effect of changes in regulatory capital on funding costs.[1] Relative to these papers, I take a simpler approach to identifying the effects of funding costs. This is partly motivated by theory, partly by practical consideration. The models I study only suggest that funding costs should depend on default probabilities, not on a host of other control variables. From a practical perspective, moreover, parsimonious models can explain a substantial share of the variation in bank funding costs and seem preferable in applied stress testing as it is difficult to condition on a range of variables which are not necessarily being forecast in the stress test itself. Finally, and this is perhaps the main contribution of the paper, I outline a method for incorporating funding stress in a balance-sheet based stress test where the magnitude of the funding stress is a function of market-based risk measures.

The rest of the paper proceeds as follows. Section 2 discusses how Merton models can be used to measure the risk of banks. In section 3, I examine the relationship between these risk measures and funding costs. Section 4 discusses how funding stress can be included in a stress test, with an example application. Section 5 concludes.

[1] They find that a one percentage point decrease in regulatory capital ratios increases wholesale funding costs, measured by CDS spreads, by more than one percentage point. An effect of such magnitude seems unusually large to this author.

2. Measuring bank risk

In this section, I start from the classic Merton model and then look at how both simplifications and extensions can be used to inform us about the default risk of banks.

In the Merton model, a firm's equity is viewed as a call option on the value of the firm's assets, and option pricing can be used to value a bank's equity and debt and to calculate default probabilities. The analogy between equity and an option is imperfect for a number of reasons. Options prices are typically found using replication arguments, but these do not quite apply in the case of a firm: One typically cannot purchase the underlying asset, i.e. the whole firm. Moreover, unlike an option which has a fixed expiry date, a firm's debt generally does not expiry on a singly date. And so forth.

From a practical perspective, however, the default probabilities - or, equivalently, the distance-to-default (DD) - produced by the Merton model have proven useful for ranking firms in default prediction. To generate default distances, one needs estimates of the value of a firm's assets and their volatility. Given these, the default probability in the Merton model is given by

$$N(-DD) = N\left(-\frac{log(V/D) + (\mu - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}\right) \tag{1}$$

where V is the value of the underlying assets, D is the face value of the debt, μ and σ the drift and volatility parameters, T the time-to-maturity, and N is the cumulative distribution function of the normal distribution.

In practice, the drift term is often ignored, and the distance-to-default is written by practitioners in the more intuitive form (e.g. Lando (2004))

$$DD = \frac{V - D}{\sigma_V V} \tag{2}$$

where σ_V is the volatility of the assets.

That is, the distance-to-default can be thought of as the number of standard deviations by which the value of the underlying assets must decline in value in order for the firm to be insolvent. For practical purposes, D may not necessarily be the actual debt level, but some other value, for example the trigger point at which default occurs. If the actual debt level is used, the distance-to-default is sometimes called a distance-to-insolvency (DI).

2.1. Towards even simpler risk measures

Taking another structural model, the Leland (1994) model, as their outset, Atkeson, Eisfeldt, and Weill (2014) show that

$$DI < \frac{1}{\sigma_E} < DD_{Leland}. \tag{3}$$

As long as creditors force insolvent firms into default before they are deep into insolvency, this bound is relatively tight. This suggests using $\frac{1}{\sigma_E}$ as a risk measure. This measure is much simpler to compute than traditional distance-to-default measures which require the use of numerical procedures to simultaneously estimate both V and σ_V, which are unobserved.

The use of $\frac{1}{\sigma_E}$ as a risk measure can likewise be motivated within the Merton model. Defining DD as in (2), and noting that the equity volatility is

$$\sigma_E = N(d_1)\frac{V}{E}\sigma_V \tag{4}$$

where $d_1 = \frac{log(V/D)+(r+\sigma^2/2)T}{\sigma_V\sqrt{T}}$ and $N(d_1)$ is the derivative of the equity value w.r.t. the asset value. When the asset value far exceeds the debt level, this value is close to one. On the other hand, when V is close to D, or below as is possible in the Merton model, the derivative can be smaller than one due to debt overhang effects.

If the face value of the debt is used as in the formula for the distance-to-default, inserting (4) into (2) gives:

$$DD = N(d_1)\frac{V-D}{E}\frac{1}{\sigma_E} \tag{5}$$

Since $E \geq V - D$ (because of the option value) and $N(d_1) \leq 1$, it follows that $\frac{1}{\sigma_E}$ is an upper bound on the distance-to-default in the Merton model, with the two measures being close when V is large relative to D. One can furthermore argue that the sensitivity of equity value to the value of the underlying assets cannot be too small for banks. For one, as we see later (in figure 2), the relationship between equity and asset values is closer to being linear in models when a default boundary is introduced, as opposed to default only being possible on a particular date as in the Merton model. In addition, the default boundary for banks is higher than the actual debt level due to solvency regulations.

Aside from computational simplicity, the inverse of a firm's equity volatility is intuitive as it combines information about the riskiness (volatility) of a firm's assets with its leverage, both quantities of interest when evaluating the risk of a firm.

Another simple approach is to look at a banks market solvency, i.e. $\frac{E}{V}$. One can think of this as a distance-to-default measure which ignores information about asset volatility, arbitrarily setting it equal to 1 in the expression for DD. Since default distances are typically used to rank firms, such an approach could be justified if one believes that the set of firms being compared, in this case banks, are relatively similar in terms of how risky their assets are. Another justification is that it avoids the problem of having to estimate the asset volatility.

2.2. Towards more complicated risk measures

The use of default distances for banks can be criticized on a number of grounds. In the following, I look at three alterations to the Merton model. The first is inspired by Chan-Lau and Sy (2006) who argue that the traditional distance-to-default measure should be altered for banks as regulators take a number of pre-emptive actions to close or restructure a bank before it is technically insolvent. The distance should therefore reflect capital regulations which require banks to have equity capital in excess of a ratio of risk-weighted assets. A second critique, which can be combined with the first, relates to the assumption that banks can only default at some arbitrary fixed date T. In reality, banks can default whenever their asset values fall below a regulatory threshold. The Black and Cox (1976) model can be used to calculate default probabilities in this case. Finally, it may not make sense to assume that the stochastic process for bank assets is a geometric Brownian motion as in these models. A bank, the argument goes, is in the business of making loans, and loans have limited upside.

At most, the bank is repaid what it is owed. This limited upside features means that the payoff profile of a loan is similar to the payoff profile of a short position in a put option. A bank's assets can therefore be viewed as a portfolio of short positions in put options. Equity is still a call option on the value of the assets and is therefore like an option-on-options. That is the essence of the model in Nagel and Purnanandam (2015) of which I implement a slightly simplified version.

Before proceeding to these models, a brief note on model comparison: As noted earlier, DD is often written in simplified form as $DD = \frac{V-D}{\sigma_V V}$ in the context of the Merton model. It might not be consistent to derive the asset values and volatilities from any model and then define DD similarly. In the standard Merton model, we know how the default distance is related with the default probability through the normal distribution. The more complicated models will generally be evaluated under the risk-neutral measure, and one can think of the risk-neutral default probability ($RNDP$) being related to the default distance through: $RNDP_{model} = N(-DD_{model})$, where $N(\cdot)$ is the cumulative normal distribution. Other models might not produce an exact expression for DD, but they do result in risk-neutral default probabilities. To have comparable distance measures across models, I will therefore define the distance-to-default from a particular model as

$$DD_{model} = N^{-1}(RNDP_{model}). \tag{6}$$

2.2.1. A model with a regulatory default ratio as trigger

Rather than working with a traditional distance-to-default, Chan-Lau and Sy (2006) define a so-called "distance-to-capital", which they define as

$$DC = \frac{V - \lambda D}{\sigma_V V} \tag{7}$$

where $\lambda = \frac{1}{1-CR}$ and CR is a regulatory capital ratio. The traditional DD would result by using $\lambda = 1$.

In the following, I follow a slightly different approach by making the capital ratio a function of the riskiness of the assets, as is the intention when using risk-weights. To motivate this, consider a regulator who might have to take over a bank in case of default and wants a buffer to ensure that the probability of losses is limited conditional on default.

One can think of the regulator as trying to choose a default point V_d such that the probability of losses occuring over a certain period, T (say, a year), after the takeover is limited to a certain probability p. This logic is similar to that used in Value-at-Risk calculations. The problem is then to identify a default point satisfying

$$Pr\left\{V_d e^{\left(\mu-0.5\sigma_V^2\right)T+\sigma_V\sqrt{T}N(0,1)} \leq D\right\} = p. \tag{8}$$

Rearranging, this becomes

$$Pr\left\{N(0,1) \leq \frac{ln(D) - ln(V_d) - \left(\mu - 0.5\sigma_V^2\right)T}{\sigma V\sqrt{T}}\right\} = p, \tag{9}$$

and

$$N^{-1}(p) = \frac{ln(D) - ln(V_d) - \left(\mu - 0.5\sigma_V^2\right)T}{\sigma_V\sqrt{T}}. \tag{10}$$

Solving for V_d:

$$V_d = De^{-\sigma_V\sqrt{T}N^{-1}(p)-(\mu-0.5\sigma_V^2)T}. \tag{11}$$

Using a first-order Taylor expansion and ignoring the drift term gives

$$V_d = D\left(1 - \sigma_V\sqrt{T}N^{-1}(p)\right). \tag{12}$$

This is similar to the formulation in Chan-Lau and Sy (2006), with the difference that the above formulation does not directly use a capital ratio, but instead links the excess solvency of the bank to the riskiness of its assets and how safe $(1-p)$ the regulator wants banks to be (which can be viewed as a proxy for capital ratios).

An alternative way of arriving at a similar formulation is to first note that banks' solvency requirement is of the form:

$$E \geq \sum_i rw_i A_i \tag{13}$$

where rw_i are risk-weights associated with some asset class i and A_i are the asset values in class i.

If we assume that risk weights are set so as to grow linearly in asset volatility, we can think of this as[2]:

$$E = V - D \geq \kappa \sigma_V V \tag{14}$$

where κ reflects the regulators risk aversion, with a greater κ implying that banks must satisfy a tougher capital requirement.

This implies a default boundary of the type:

$$V_d = (1 - \kappa \sigma_A)^{-1} D \approx (1 + \kappa \sigma_V) D \tag{15}$$

with the last approximation holding for small $\kappa \sigma_a$. To estimate the model, I follow the same procedure as for the standard Merton model, only with V_d as defined in (12) as default point at T instead of D. By "estimate", I again mean numerically derive values of V and σ_V so as to match observed stock prices and volatilities and then calculate risk-neutral default probabilities.

2.2.2. Introducing a default barrier

One of the simplifying assumptions of the Merton model is that the value of a firm's assets is evaluated on a particular date, and if the assets are worth less than some default point, for example the face value of the debt, the equity holders get nothing. In reality, of course, firms can default continuously. It is possible that this makes a difference when trying to explain funding costs, not least because banks are constantly being scrutinized to determine whether they satisfy regulatory capital ratios.

[2]For example, with two asset classes we get: $V \geq (1 - \kappa(\sigma_1 w - \sigma_2(1-w)))^{-1} D$ where w is the weight of the first class.

To address this, I solve the model numerically with default barrier $De^{-r_f T}$.[3] One can also find closed-form solutions for this type of problem. Black and Cox (1976) show how to calculate risk-neutral default probabilities, bond prices, etc. in a setting with a default boundary of the form $De^{-r_f T}$.[4]

2.2.3. Bank equity as an option-on-options

Nagel and Purnanandam (2015) take a starker departure from the traditional Merton setup. Rather than simply changing the default point or introducing a default barrier, they argue that the assumption of modeling bank assets as a geometric Brownian motion is inappropriate. Instead, they model banks as owning a portfolio of loans. Since loans have limited upside, it is not possible for bank's assets to have unlimited upside as implied by a lognormal distribution. Instead, the payoff structure of a loan resembles a short position in a put position, and the bank can therefore be viewed as having a call option (equity) on a portfolio of put options. Bank equity is therefore an option-on-options and has a payoff structure similar to mezzanine claim.

In the following, I sketch a slightly simplified version of the model in Nagel and Purnanandam (2015). It is simplified in the sense that I maintain the main insight, the option-on-options-like nature of the equity claim, but dispense with the staggered maturity structure which is also part of the original model.

The model involves a portfolio of loans, each with face value F. Underlying a loan is a collateral of value A_t^i. There is a continuum of borrowers, index by $i \in [0,1]$, with unit mass. Under the risk-neutral measure, the value of the collateral is assumed to evolve according to

$$\frac{dA_t^i}{A_t^i} = rdt + \sigma_A + \left(\sqrt{\rho}dW_t + (\sqrt{1-\rho}dZ_t^i\right). \tag{16}$$

[3] I have also tried to estimate the model using $V_d e^{-r_f T}$ as default barrier, but this does not improve the model's ability to e.g. explain funding costs.

[4] My numerical solution differs slightly from the traditional Black-Cox model. Instead of assuming that debt holders become equity holders when the default barrier is hit, I assume that the bank is unwound over a period of one year. The debt holders then get the minimum of the asset value or the face value of the debt at the end of the period.

This formulation entails that collateral values are subject to both common (dW) and idiosyncratic shocks (dZ^i), with ρ determining the relative weight of these shocks. For example, if ρ equals 1, only common shocks matter. The idiosyncratic shocks are independent across borrowers, and all shocks are assumed to be normally distributed.

The payoff to a bank on a loan is given by

$$L^i = min\left\{A^i_T, F\right\}. \tag{17}$$

To find A_T, note that the differential of the log asset value can be written as

$$dln\, A^i_t = \left(r - 0.5\sigma^2\right) dt + \sigma\sqrt{\rho} dW_t + \sigma\sqrt{1-\rho} dZ^i_t. \tag{18}$$

Normalizing the initial collateral value to one, i.e. $A_0 = 1$ gives

$$A^i_T = e^{(r-0.5\sigma^2)T + \sigma\sqrt{\rho}W(T) + \sigma\sqrt{1-\rho}Z^i(T)}. \tag{19}$$

For a particular realization of the common shock, a borrower repays less than the face value whenever $A^i_T < F$, that is when

$$Z^i(T) \leq \frac{ln\, F - \left(r - 0.5\sigma^2\right) T - \sigma\sqrt{\rho}W(T)}{\sigma\sqrt{1-\rho}}. \tag{20}$$

Using this, one can show the aggregate payoff to the equity owners to be

$$L_T = \int_i L^i_T di = A_T N\left(d_1\right) + F_1 N\left(d_2\right) \tag{21}$$

where A_T, the aggregate payoff to the entire assets, is $A_T = e^{\left(r-0.5\rho\sigma^2\right)T + \sigma\sqrt{\rho}W(T)}$, and:

$$d_1 = \frac{ln\, F - \left(r - 0.5\sigma^2\right) T - \sigma\sqrt{\rho}W(T)}{\sigma\sqrt{1-\rho}\sqrt{T}} - \sigma\sqrt{1-\rho}\sqrt{T}, \tag{22}$$

$$d_2 = -\frac{ln\, F - (r - 0.5\sigma^2)\, T - \sigma\sqrt{\rho}W(T)}{\sigma\sqrt{1-\rho}\sqrt{T}}. \tag{23}$$

Using the expression for L_T, the value of the assets can then be found by Monte Carlo simulation, noting that $V_0 = e^{-rT}E_0^Q\,[L_T]$. One can similarly find the value of the debt and the equity by calculating the payoffs to debt and equity holders.

In practice, we need to derive both asset values and volatilities from equity values and volatilities. It would therefore be convenient if we could express the equity volatility in terms of the asset volatility. To do so, it is helpful to write the log of the aggregate payoff from the assets

$$a_T = \int_i ln\, A_T^i di = \left(r - 0.5\sigma_A^2\right) T + \sigma_A\sqrt{\rho}W(T). \tag{24}$$

The payoff of the assets can then be written as

$$L_T = e^{a_T + 0.5(1-\rho)\sigma_A^2 T} N\left(\frac{ln\, F - a_T}{\sigma_A\sqrt{1-\rho}\sqrt{T}} - \sigma_A\sqrt{1-\rho}\sqrt{T}\right) + FN\left(-\frac{ln\, F - a_T}{\sigma_A\sqrt{1-\rho}\sqrt{T}}\right). \tag{25}$$

Now, since the equity value depends on L_T, and L_T depends on a_T, we can write the equity value as a function of the log asset value, ie. $E = E(a_t, t)$. Applying Ito's lemma and focusing on the stochastic part, we find

$$dE_t = (...)dt + \frac{\partial E}{\partial a_t}\sigma_A\sqrt{\rho}dW_t. \tag{26}$$

One possible scheme for linking equity and asset volatilities is therefore to calculate $\frac{\partial E}{\partial a_t}$ numerically and then use

$$\sigma_E = \frac{\frac{\partial E}{\partial a_t}}{E}\sqrt{\rho}\sigma_A. \tag{27}$$

This condition, along with $E = E(a_t, t)$, can then be used to find values of a_t and σ_A to match observed values of E and σ_E.

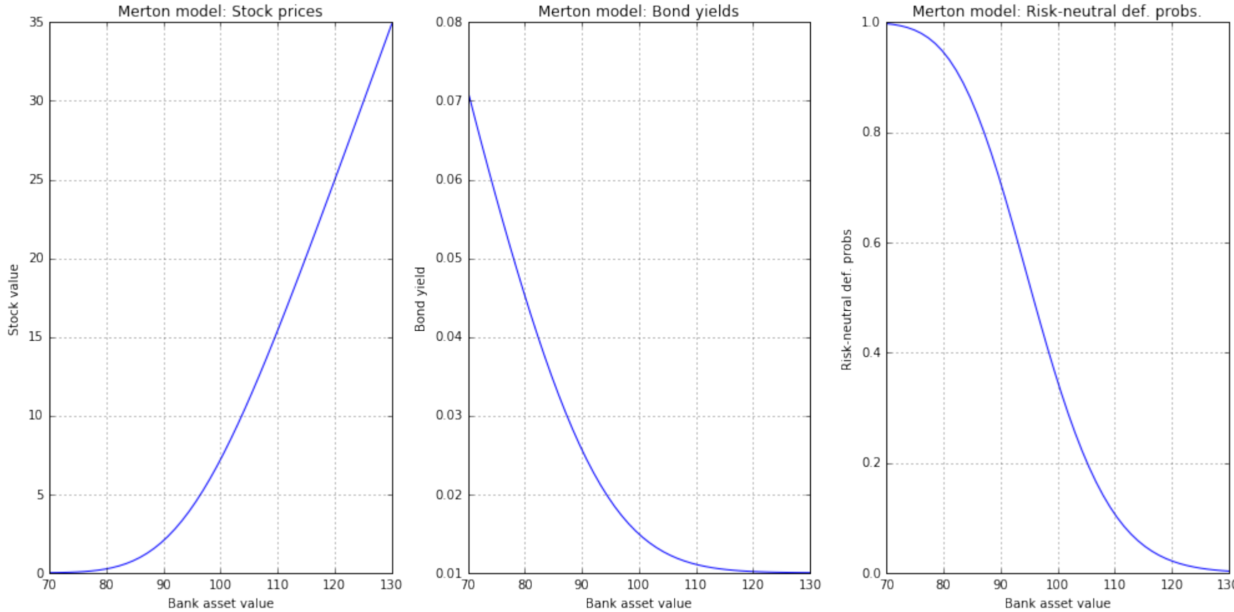

Figure 1: The figure shows the stock price (left), bond yield (center) and risk-neutral default probability (right) as a function of the underlying bank asset value in a Merton model.

2.2.4. Some qualitative features of the models

In this section I briefly look at some qualitative features of the Merton model and its extensions. The standard Merton model and the model with a higher default point are qualitatively quite similar. Figure 1 plots the stock price, bond yield, and risk-neutral default probabilities in a Merton model.

Compare this to the corresponding plots in a model with a default barrier as in figure 2. A noticeable difference is that the stock price is much closer to being a linear function of the underlying asset value when a default barrier is introduced. This is because it is not possible for the asset value to be too low, or “out-of-the-money”, for if it were, default would have occurred.

Now, recall that in the Merton model

$$\sigma_E = N(d_1)\frac{V}{E}\sigma_V. \tag{28}$$

Here $N(d_1)$ represents the “delta” or sensitivity of the stock price with respect to the underlying asset value (with d_1 defined as in the Merton model). In a model with a default

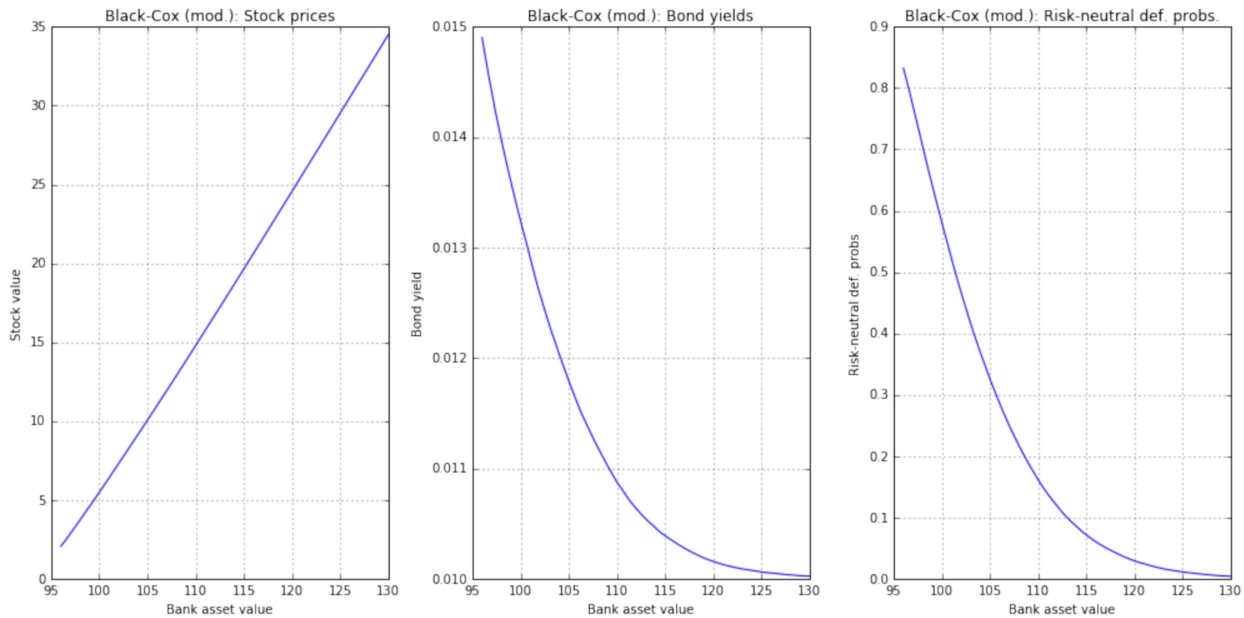

Figure 2: The figure shows the stock price (left), bond yield (center) and risk-neutral default probability (right) as a function of the underlying bank asset value in a Merton model with a default barrier.

barrier this delta is mostly close to one, though it can be somewhat lower when the asset is close to the default boundary. I also showed earlier that the distance-to-default boils down to $\frac{1}{\sigma_E}$ when $N(d_1) = 1$. To the extent that one believes a default barrier to add further realism to a Merton-type model, one can therefore argue that it also adds some support to the use of the inverse equity volatility as a risk measure.

Figure 3 shows similar plots for the options-on-options model, only with a standardized version of the log asset value as a state variable. The plot to the left shows the mezzanine-like payoff structure to equity holders.

An interesting feature of this model, which is discussed at greater length in Nagel and Purnanandam (2015), is that it has implications for the measurement of volatility. To calculate a distance-to-default, one needs an estimate of the asset volatility, and this estimate is typically derived wholly or in part from the equity volatility. What Nagel and Purnanandam (2015) show is that if their model is the “true” model, then the procedures used to derive asset volatilities in other models such as a Merton model may give misleading answers. As a simple illustration of this, figure 4 shows the asset volatilities recovered in the Merton model, the model with a default barrier and the option-options model from the same set of

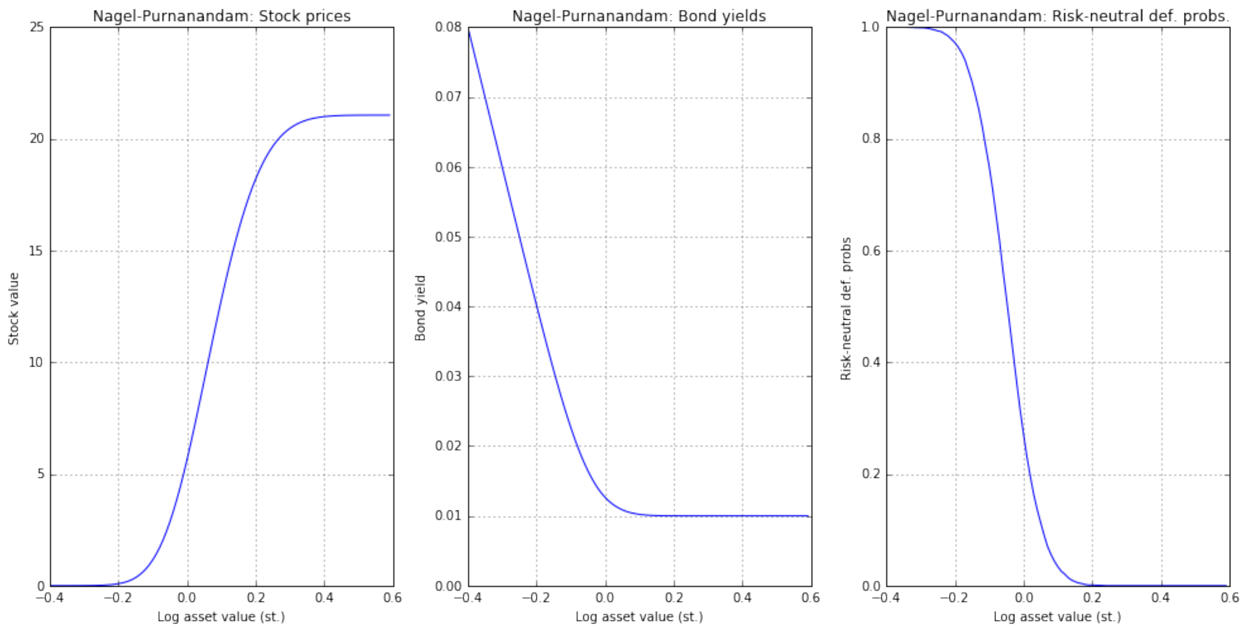

Figure 3: The figure shows the stock price (left), bond yield (center) and risk-neutral default probability (right) as a function of the underlying log asset value in an "option-on-options" model similar to that in Nagel and Purnanandam (2015).

observations of markets solvency (set at 10.0 percent) and equity volatilities. It indicates that the option-on-options model tends to recover higher asset volatilities than do the other models.

One can get an intuitive understanding of why this might be the case simply by looking at the leftmost plot in figure 3. Suppose, for example, that times are good in the sense that bank asset values, or the collateral supporting bank loans, are high relative to debt levels. Banks therefore suffer few losses, and we are situated to the right in figure 3 (left panel). In that area, the equity value is almost a flat function of the underlying asset value. This means that equity volatilities are low. In other models, one would conclude from low equity volatilities that asset volatilities are low. That, in turn, would lead to high levels of default distances, indicating that risk too is low. Hence, to the extent that the option-on-options model is a good description of reality, one might be prone to underestimate risk in "good times" when using other models. The model is also consistent with equity volatility spiking relatively quickly as stock values move from the flat part of the plot to the part where stock prices become sensitive to movements in the underlying asset values.

While these plots highlight differences across models, they also point to a shared feature:

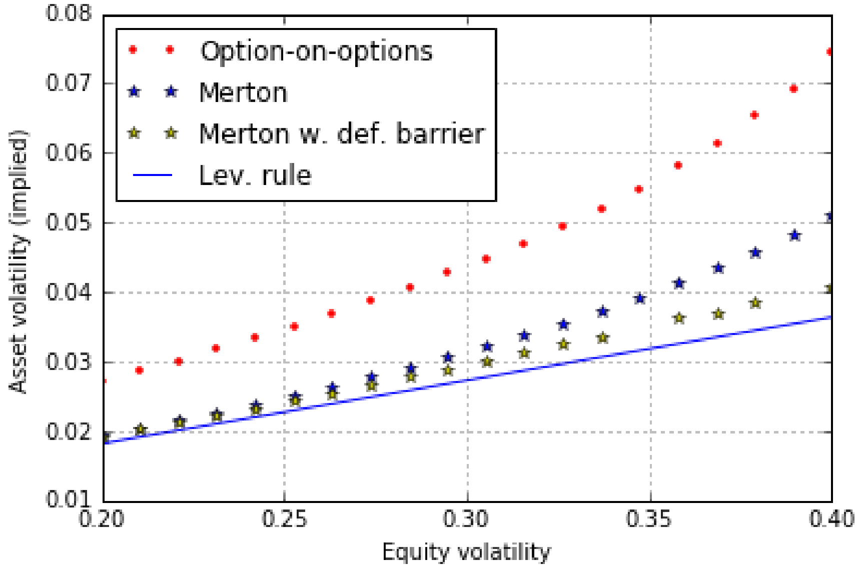

Figure 4: The figure shows the asset volatilities which have been backed out from the Merton model, a model with a default barrier, and the option-on-options model based on a market solvency of 10.0% and equity volatilities in the range 0.2-0.4. Also shown is the asset volatility identified from the following simple leverage rule: $\sigma_V = \frac{E}{V}\sigma_E$.

A non-linear relationship between asset values and funding costs. When asset values are high, funding costs are not particularly responsive to changes in asset values, whereas they are substantially more responsive when asset values are low.

2.2.5. Implementation details

All of the above models can be used to calculate risk-neutral probabilities. This involves numerical estimation of asset values and volatilities from equity values and volatilities.[5] For example, in the Merton model we write the following two equations,

$$E = VN(d_1) - e^{-rT}DN(d_2) \tag{29}$$

where $d_1 = \frac{log\, V_t - log\, D + (r+0.5\sigma^2)T}{\sigma\sqrt{T}}$ and $d_2 = d_1 - \sigma\sqrt{T}$. Furthermore,

$$\sigma_E = N(d_1)\frac{V}{E}\sigma_A. \tag{30}$$

[5]In the case of the option-on-options model, I use the log asset value, as defined in the previous section, as a state variable rather than the asset value itself.

Hence, one needs to solve two nonlinear equations in two unknowns. I solve these numerically by choosing V and σ_A to minimize the sum of the squared differences between actual and model-implied equity values and volatilities. I match the model to realized equity volatilities based on the past 3 months equity returns. For the Merton model, the procedure recovers asset values and volatilities in 95.7% of cases, which is comparable with recovery rates using other methods.[6] The algorithm mainly seems to have difficulties when encountering very high equity volatilities. The parameters are recovered similarly in the other models, with minor adjustments to the procedure. In each model the time to maturity, T, is set to 5 years as the models will be compared, among other things, by their ability to explain 5-year credit default swap (CDS) spreads.

The options-on-options model requires additional parametric assumptions, and here I mainly follow Nagel and Purnanandam (2015). For example, one must specify the parameter ρ, which governs how much idiosyncratic shocks matter relative to common shocks. I set ρ to equal 0.5. I also normalize the bank's debt level to equal one, and then specify the equity value relative to D. Another issue is how to specify F, the face value of the debt. Nagel and Purnanandam (2015) set $F = e^{rT}$, but such an assumption makes it difficult to recover asset values for a larger number of observed equity values. This is because the size of F determines the maximal return on loans, and it is difficult to explain observations of high equity values without allowing for more upside (especially when risk-free rates are low). I therefore set $F = e^{(r+0.06)T}$ since this allows me to recover more values, but I am nonetheless only able to recover values in about 65% of cases. This, presumably, is related to both the limited upside feature of the model, which means that there are a range of observations of equity values and volatilities with which the model simply is not consistent, and to the fact that the equity value is a virtually flat function of the log asset value in some regions, which makes convergence difficult.

[6] As an example, Atkeson, Eisfeldt, and Weill (2014) report a convergence rate of 95% when using an iterative algorithm

3. Bank risk measures and funding costs

In this section I examine how well the risk measures derived from the variations - simplifications and extensions - on the Merton model help explain actual funding costs. Following Bharath and Shumway (2008) I will also define a "naive" risk measure which attempts to integrate some of the qualitative lessons from the various risk measures, but which can be easily calculated without the need for numerical estimation.

3.1. Data sources

The subsequent analysis draws on data from three sources. (1) All of the bank balance sheet and market data is from SNL Financial and covers the period from beginning 2007 in the case of price data and beginning 2008 in the case of balance sheet data to end-June 2016. (The coverage for non-US banks seems rather limited pre-2008). 947 banks from the US, Canada and Europe are included in the sample. (2) Data on 5-year CDS premia and 5-year government bond rates (used as proxies for risk-free rates) is from Bloomberg. CDS premia are only available for a small subset of banks (57 banks) in the full sample. (3) Later, I consider an application to the top-down stress test conducted by Danmarks Nationalbank. The balance sheet data used in that application is from the Danish FSA whereas the market data is from Bloomberg.

3.2. A first glance at the risk measures

To get a initial sense of whether the distance measures actually capture the riskiness of banks, I look at how the measures compare to a perhaps more familiar measure of riskiness, ratings (concretely, Moody's long-term ratings). Figure 5 compares values of some of the risk measures across rating classes. The risk measures are clearly able to distinguish between banks, at least as long as the banks are somewhat risky (have a rating of B or less). On the other hand, the measures seemingly do not discriminate between banks as long as banks are highly rated, though this could be due to timing issues. As an example, some banks

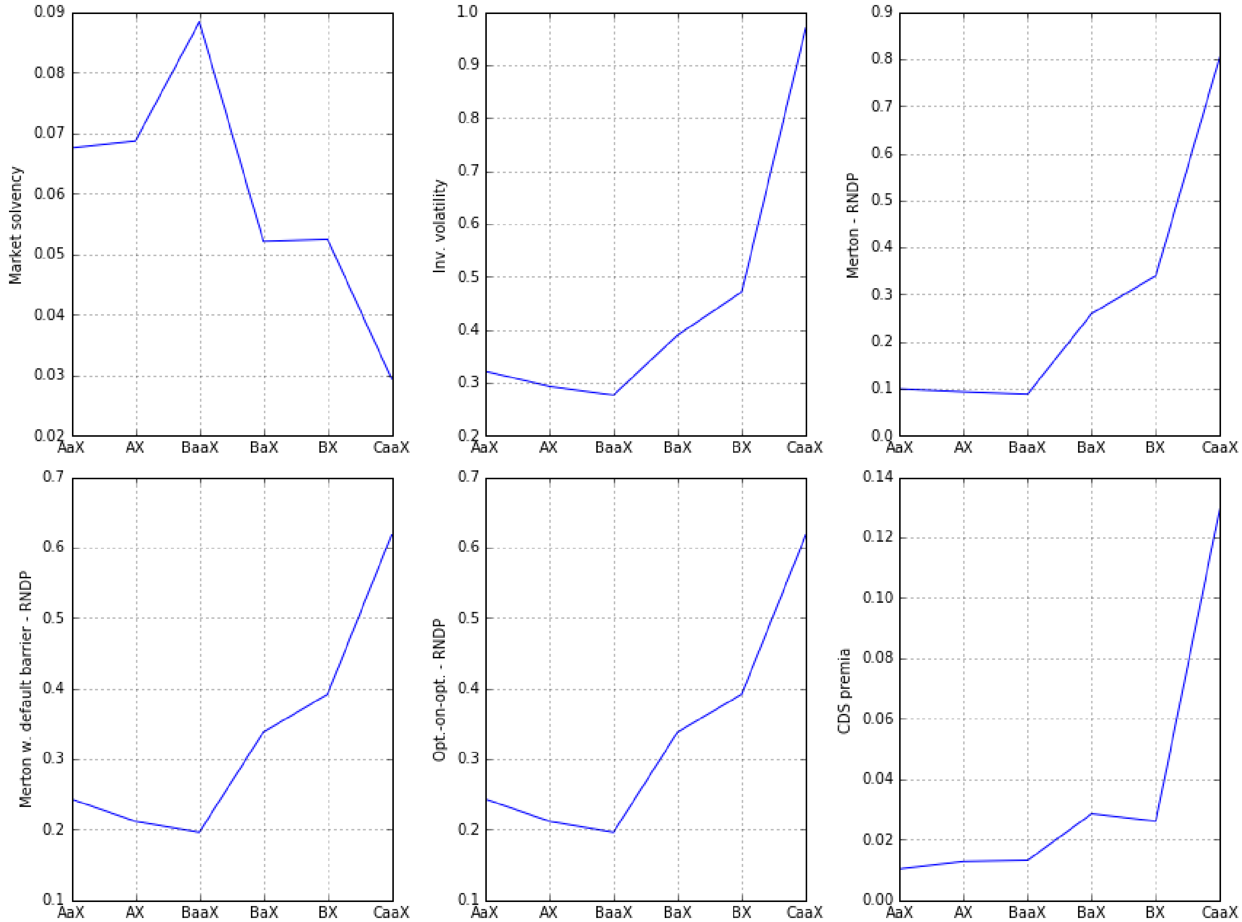

Figure 5: The figure shows how market solvency, the inverse of volatility, risk-neutral default probabilities from the Merton model with and without default barrier, risk-neutral default probabilities from the option-on-options model, and also CDS premia compare across rating classes. Ratings are Moody's corporate long-term rating. A ratings class such as "AaX" covers ratings "Aaa", "Aa2" and "Aa3". The ratings data is based on quarterly from 2008Q1-2016Q2 from SNL Financial. The risk measures are calculated for the same period, also based on data from SNL Financial. Data on CDS premia is from Bloomberg.

may still have had relatively high ratings at the beginning of the sample, in 2008, when market-based risk measures began to look worse, and ratings methods themselves may have changed within the period.

As a second point, we might consider how different these risk measures really are. Table 1 displays the sample correlations between the default distances for each model. The upper panel shows the correlations for the full sample while the bottom panel shows the correlations for weaker banks, here defined as banks with a market solvency of less than 3 percent.

The risk measures are generally highly correlated, with the market solvency a bit of an outlier. Moreover, comparing the two panels, we observe that the correlation decreases for

Table 1: Correlations between default distances / risk measures from models

Full sample					
Model	1	2	3	4	5
1. Merton					
2. - w. solvency adjustment	0.928				
3. - w. default barrier	0.940	0.957			
4. Option-on-options	0.977	0.957	0.952		
5. Inverse volatility	0.802	0.810	0.822	0.772	
6. Market solvency	0.346	0.479	0.418	0.540	0.314

Weaker banks (market solvency $\leq 3\%$)					
Model	1	2	3	4	5
1. Merton					
2. - w. solvency adjustment	0.678				
3. - w. default barrier	0.822	0.711			
4. Option-on-options	0.977	0.732	0.896		
5. Inverse volatility	0.785	0.641	0.798	0.787	
6. Market solvency	0.269	0.310	0.384	0.381	0.216

more troubled banks, a sign that the models are not entirely identical when it matters most.

Interestingly, the default distance implied by option-on-options model is highly correlated with the distance implied by the traditional Merton model. This may be because of the difficulty in recovering parameters from the option-on-options model in the regions where the equity value is (virtually) a flat function of the asset value: We therefore do not observe the values implied by the option-on-options model in the cases where they would have been different.

3.3. A naive measure: Combining information from more complicated risk measures to construct a naive measure

The discussion of the qualitative features of the various models offers a number of suggestions as for how to integrate some features of more complicated models into simpler models. For example, introducing a default barrier generally means that the derivative of the equity value to the asset value is closer to one than in a model without such a barrier. This by itself suggests using the approximation $\sigma_E \approx \frac{V}{E}\sigma_V$, which motivates the use of $\frac{1}{\sigma_E}$ as a risk-measure.

We can incorporate further insights from the models, for example from the options-on-options model. That model in particular suggests that Merton-type models will sometimes tend to underestimate true asset volatilities. Now, from the above discussion we know that using a simple leverage rule of the type $\sigma_V = \frac{E}{E+D_{Book}}\sigma_E$ might be reasonable if the true model is close to a Merton model with a default barrier. Perhaps the simplest way of bridging these observations is to take a simple average of a moderately high (for a bank) asset volatility, say 0.04, and the asset volatility calculated using the simple leverage rule.

Combining these findings suggests the use of the following simpler measure

$$DD_{simple} = \frac{E}{\sigma_V \left(E + D_{Book}\right)}, \tag{31}$$

with the asset volatility defined by

$$\sigma_V = 0.5\left(0.04 + \frac{E}{E + D_{Book}}\sigma_E\right). \tag{32}$$

This admittedly is arbitrary, and it is easy to criticize the procedure. However, as Bharath and Shumway (2008) note in their construction of a naive default distance, the whole point is to use a predictor which is easy to calculate and may have predictive power. If more complicated models are really better, they should have higher predictive power, but - as will be shown in the following - they generally do not.

One could also adapt the numerator to take into account a higher default point due to solvency regulations. This could be accomplished by replacing the numerator with

$$E - \kappa\sigma_V D. \tag{33}$$

In the following, I will not include this final step in the naive default distance since it does not seem to improve the measure's ability to explain funding costs.

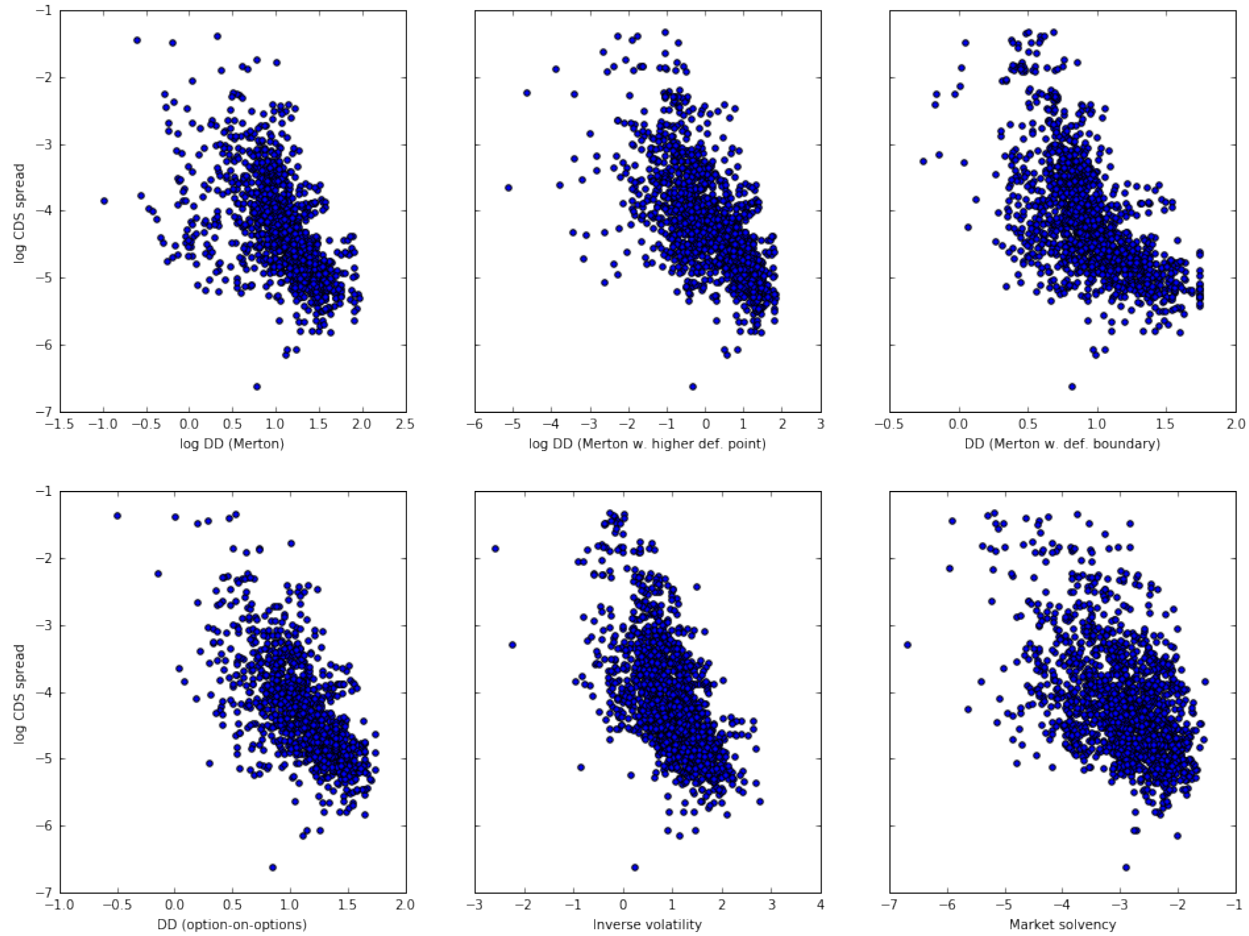

Figure 6: The figure shows scatter plots of log CDS spreads and log default distances from the Merton model and the variations of it considered in this paper. Data on CDS spreads is from Bloomberg, and the data used to calculate the default distances is from SNL Financial for the period 2008Q1-2016Q2.

3.4. Risk measures and funding costs

When attempting to measure funding costs, it is common to look at CDS spreads. Figure 6 shows the relationship between default distances and 5-year CDS spreads for all of the banks in the data set for which spreads are available. If one were to plot the relationship between default distances and CDS spreads, it would be highly non-linear, with CDS spreads increasing strongly for low distances. When CDS spreads are plotted against log distances, the relationship looks roughly linear for most of the risk measures as the figure shows.

All of the risk measures seem quite strongly correlated with CDS spreads, but the relationship looks clearer for some, for example the inverse volatility, than for others, for example market solvency. To get a better sense of which measure best explains funding costs, table 2 shows the results of regressions of log CDS spreads on the log of the default distances. This table also includes the naive default distance. Interestingly, both the naive measure and other simple measures such as the inverse volatility do a rather good job of explaining funding costs.

Table 2: The table shows the beta-coefficients from regressions of the form: $log\,CDS_i = \alpha + \beta log\,DD_i + \epsilon_i$, where DD is a default distance measure. All coefficients are significant at $p < 0.01$. Data on 5-year CDS premia is from Bloomberg, and the distance measures are calculated based on data from SNL Financial over the period 2008Q1-2016Q2.

Distance measure	β	R^2	β	R^2
1. Merton	-0.39	0.29	-0.30	0.47
2. - w. solvency adjustment	-0.41	0.29	-0.34	0.52
3. - w. default barrier	-0.36	0.34	-0.28	0.50
4. Option-on-options	-0.39	0.26	-0.30	0.45
5. Inverse volatility	-0.84	0.38	-0.59	0.59
6. Market solvency	-0.55	0.24	-0.72	0.66
7. Naive measure	-0.76	0.35	-0.83	0.68
Country fixed effects	No		Yes	

The log-log formulation means that percentage changes have the same percentage effects. As an example, if a reduction in default distances from 4 to 2 makes spreads double, say from 0.5 to 1, a further halving from 2 to 1 would again lead to a doubling of spreads from 1 to 2. The implication is that already weak banks are more susceptible to increased funding costs if they suffer further losses.

One could also, as e.g. Bharath and Shumway (2008) do, regress log CDS spreads on the logs of risk-neutral default probabilities, but this approach does not seem to hold greater explanatory power. As an example, for the classic Merton model the R-squared is 0.29 when regressing on the log of default distances versus 0.27 when regression on the log of risk-neutral default probabilities.

In practice, working with CDS spreads is less than ideal in a stress test setting. A regulator conducting a stress test will presumably want to identify specific parameter estimates for banks in its jurisdiction, but few banks have traded CDS contracts to their names. In addition, the set of banks for which CDS data is available may not be representative. As an example, Acharya, Anginer, and Warburton (2016) find that bond credit spreads are more risk-sensitive for smaller financial institutions than for the largest ones. If that is the case, one might underestimate the actual relationship between risk and funding costs since CDS premia are mainly available for the largest financial institutions. It is also unclear how to apply the spread change in a stress test. A bank has multiple sources of funding, some secured and some unsecured, and the different sources mature at different times as well. It is therefore unclear how to translate an increase in CDS spreads into an increase in funding cost for the bank as a whole.

Rather than using CDS spreads, we can look at average bank funding costs. Since average funding costs encompasses a number of funding sources which may not be particularly responsive to changes in risk, e.g. insured deposits and secured loans, we obviously expect the relationship between risk measures and funding costs to be weaker. However, as table 3 shows, it is still present. Riskier banks face higher funding costs on average.

Table 3: The table shows the beta-coefficients from regressions of the forms (1) $r_i = \alpha + \beta DD_i + \epsilon_i$, (2) $r_i = \alpha + \beta log\, DD_i + \epsilon_i$, and (3) $log\, r_i = \alpha + \beta log\, DD_i + \epsilon_i$, where DD is a default distance measure. The interest rate, r_i, is defined as the quarterly interest expenses (times 400 to turn it into an anual interest rate) of a bank divided by total liabilities less equity. Country fixed effects are included in all of the regressions. All coefficients are statistically significant at $p < 0.01$. Data on interest expenses and the data used to calculate the distance measures are from SNL Financial over the period 2008Q1-2016Q2. N = 13,042.

Distance measure	β (lin-lin)	R^2	β (lin-log)	R^2	β (log-log)	R^2
1. Merton	-0.18	0.33	-0.21	0.37	-0.11	0.22
2. - w. solvency adjustment	-0.18	0.34	-0.10	0.44	-0.03	0.27
3. - w. default barrier	-0.24	0.33	-0.16	0.34	-0.08	0.20
4. Option-on-options	-0.25	0.39	-0.20	0.37	-0.10	0.20
5. Inverse volatility	-0.13	0.34	-0.48	0.37	-0.43	0.21
6. Market solvency	-2.95	0.31	-0.34	0.33	-0.17	0.18
7. Naive measure	-0.22	0.36	-0.52	0.37	-0.35	0.20

When the dependent variable is average funding costs, the relationship is not quite as nonlinear. The log-log formulation, which implies constant percentage effects, shows a poorer fit than a linear or lin-log specification. The lin-log formulation seems to do marginally better than the linear specification. This suggests a somewhat non-linear relationship between funding costs and the distance measures. Moving from, say, a default distance of 4 to 2 increases funding costs by the same amount, in percentage points, as moving from 2 to

1. The absolute change in funding costs depends on the percentage change in DD, i.e. $\Delta r = \beta(log\, DD' - log\, DD)$. Again, this implies that already weak banks are more at risk of facing increased funding costs. When incorporated in a stress test, stress of funding costs will therefore tend to have an amplifying effect as the weakest banks, i.e. either as those with low default distances to start with or those who suffer large losses, will also face the greatest increases in funding costs.

Note also that the naive default distance appears to perform about as well as the more complicated alternatives, which cannot always be computed due to problems of numerical convergence.

Regressions such as those in table 3 can be extended to take into account other features which affect funding costs. As an example, in the case of Danish banks it improves the results if one includes an interaction between the default distances and the share of deposits to total debt. In that case, the parameter estimates (see the following section) show that banks with a high degree of deposit financing are less susceptible to increased funding costs due to increases in risk.

4. Incorporating funding cost increases in top-down stress tests

A stress test takes as its outset a stress scenario which describes the evolution of a number of key macroeconomic variables over an extended period such as 3 to 5 years. The development of the macroeconomy in turn affects many of the quantities which determine a bank's income: net interest margins, loan impairment charges, trading book losses, and so forth. As banks' income deteriorate due to stress, this in turn affects the balance sheet and the key capital ratios of banks. These observations, however, are not particularly helpful if we have trouble identifying the link between balance sheet based capital ratios and funding costs in the first place. The question, then, is how to incorporate market data into a stress test which in its forecasts produces balance sheet data. I outline a procedure in the following.

The starting point is to calculate one's preferred measure of default distance at the beginning of the stress period. Note that this default distance is a function of the market value of equity.

Now, a stress test is typically run over multiple scenarios. Aside from one or more stress scenarios, these commonly include a baseline scenario which is a forecast of how the economy is expected to evolve over the horizon in question. A natural approach is to exploit differences relative to this baseline scenario. In particular, the discounted difference in profits in a stress scenario relative to the profits in the baselines scenario can be viewed as a measure of the loss in market value to a bank. By first running a stress test without funding stress, one can therefore get an initial estimate of the loss in market value to a given bank.

This estimate of loss in market value can then be used to calculate a new default distance by inserting the updated equity value, $E' = E - MV_{loss}$, which is the original equity value less the estimated loss in market value. One can then calculate the change in default distance and use this to calculate a funding cost add-on. As an example, suppose one has estimated a funding cost relationship of the form

$$r_{bank} - r_f = \alpha + \beta ln\left(DD_{bank}\right). \tag{34}$$

The increase in bank funding costs is then

$$\Delta r_{bank} = r'_{bank} - r_{bank} = \beta\left(ln\left(DD'_{bank}\right) - ln\left(DD_{bank}\right)\right). \tag{35}$$

This can then be added to one's existing estimate of a bank's funding cost. The process can be repeated if one wants to incorporate feedback effects. The increased funding cost will lead to a further loss in market value, which again affects the default distance and the increase in funding cost. One can keep on going until some convergence criterion is reached. A graphical overview of the process is given in figure 7.

Figure 7: The figure shows how an initial loss, due to a stress scenario, is transmitted through to increased funding costs for a bank.

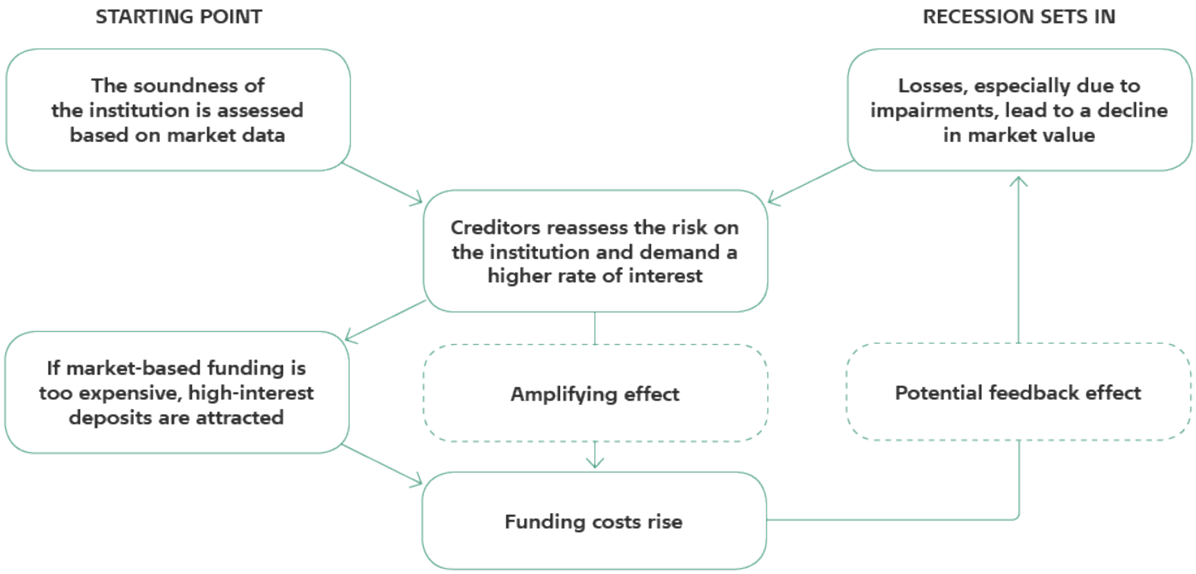

4.1. Some implementation details

In practice, there are a number of implementation details to consider. For example, when does the increased funding cost take effect? The most conservative option, that is the one which induces the maximal stress, is to let the funding cost stress kick-in immediately. This can be justified by thinking of bank creditors as rational and forward-looking. If they understand that they are in a stress scenario, they will immediately demand higher compensation to keep lending to the bank. This assumption, of course, may be too strict as not all of a bank's funding matures straightaway, and one could consider adding a delay in the introduction of the funding stress depending on the maturity structure of the bank's debt.

Another issue is how to deal with non-traded banks for which no default distance can be calculated. One possible approach is to interpolate their funding cost increases from those of other banks. This could be done by first running the stress test without the added funding in both the baseline and stress scenarios. This provides observations of the estimated market value losses for all banks. For the traded banks, it also results in estimates of the log changes in default distances (I am assuming the use of a lin-log specification of funding costs as a function of default distances). One can then regress the log changes in default distances on the loss in market values relative to banks' total assets. This regression relationship can then be used to calculate implied log changes in default distances for the non-traded banks.

Finally, at least if one uses a log-specification, one needs to consider how to deal with cumulative losses that exceed the value of a bank's equity. One cannot work with negative default distances, at least if using a logarithmic specification. One might simply conclude that the bank cannot refinance itself at any rate and therefore suffers a liquidity run, or maybe it repays maturing debt by selling off assets. However, in most solvency stress tests, banks are by assumption not permitted to simply be liquidated or engage in massive deleveraging. A pragmatic alternative then is to introduce a ceiling on how much funding costs can increase. Such a ceiling can be motivated by the presence of insured deposits. Presumably, a bank can attract sufficient deposits if only it offers an adequately high interest rate. The ceiling can then be made a function of how much market financing a bank has and how high an interest rate is required to attract deposits.

4.2. A practical application

A version of the above method has been implemented in the top-down stress test conducted semiannually by Danmarks Nationalbank. This stress test covers the 16 largest banks in Denmark and assesses their solvency over a 3-year horizon. Introducing funding stress into the model has, on average, reduced banks' common equity tier 1 ratio by 0.9 percentage points at the end of the stress period relative to the model without funding stress.

The stress test uses a modified version of the regressions in table 3. It also happens that the naive default distance does a better job than the alternatives of explaining average funding costs. The regressions are modified to include the deposit share as a covariate as well. One experience from the recent financial crisis was that banks, which relied heavily on market funding, were more vulnerable to funding shocks. To capture this, the deposit share, defined as deposits to total debt, is interacted with the distance-to-default. Table 4 shows the effect of including the deposit share.

Table 4: The table shows three regressions of funding costs on (naive) default distances and, in models 2 and 3, interactions between default distances and deposit shares. The interest rate, r_i, is defined as the quarterly interest expenses (times 400 to turn it into an anual interest rate) of a bank divided by total liabilities less equity. All coefficients are statistically significant at $p < 0.01$. Data on interest expenses is from the FSA and the distance measures are calculated based on data from Bloomberg. The sample covers traded Danish banks over the period 2004Q1-2016Q2. N = 657.

Model	β_1	β_2	R^2
1: $\beta_0 + \beta_1 DD$	-0.37		0.29
2: $\beta_0 + \beta_1 DD + \beta_2 (DD \times share_{deposits})$	-0.84	0.71	0.41
3: $\beta_0 + \beta_1 log\, DD + \beta_2 (log\, DD \times share_{deposits})$	-2.34	2.04	0.43

The first regression indicates that funding costs increase, on average, by 37 basis points for unit decrease in the naive default distance. This figure is roughly in line with the figures found for the sample of international banks in table 3. Once the deposit share is accounted for, as in the second regression, the picture changes somewhat. A bank all of whose debt is in the form of deposits will face a 13 ($= (0.84 - 0.71) \times 100$) basis point increase in funding costs per unit decrease in the default distance, whereas the effect is much larger for a bank which relies more on wholesale funding. As was the case for the international sample, the regression fit is marginally improved by modeling funding costs as a function of default distances in logs. The third specification is therefore used in the stress test of Danmarks Nationalbank.

As noted in the previous section, it might be necessary to include a ceiling on how much funding costs can increase. In the case of Danmarks Nationalbank's stress test, this is accomplished using a simple rule-of-thumb. In the wake of the financial crisis, many Danish banks with funding shortfalls advertised time deposits with interest rates roughly 2 percentage points above the standard deposit rates offered by other banks. Since existing depositors could also take advantage of such offers, it also seems reasonable to assume that some of these will take advantage of higher rates. For now, we have somewhat arbitrarily assumed that 20 percent of existing depositors must also be offered a higher rate. Letting d denote the deposit share, the funding cost ceiling is therefore currently given by $[(1 - d) + 0.2d]\, 2\%$, but the calculation may be revisited in the future.

5. Conclusion

This paper has outlined a simple method for incorporating increased funding costs in top-down stress tests. It involves estimating the initial default distance of a bank, identifying the change in default distance as a function of stress test losses, updating the default distance, and finally calculating the increase in funding costs based on the change in default distance.

The main innovation of the method is that it offers a way to include market information into a stress test which otherwise is based on balance sheet information. The method is flexible in that it can be adapted to other risk measures than the default distances in this paper, and one can relatively easily extend the method to include further covariates (such as the share of deposits to total debt) which may affect funding costs.

Finally, the method as described here is adapted for a solvency stress test. It implicitly assumes that banks can attract funding by offering a sufficiently high interest rate, but in reality a bank may not be able to attract funding sufficiently quickly if it suffers a liquidity crisis, so the model abstracts from important liquidity considerations.

References

Acharya, Viral V., Deniz Anginer, and A. Joseph Warburton, 2016, The end of market discipline? Investor expectations of implicit government guarantees, *Working Paper*.

Anneart, Jan, Marc De Ceuster, Patrick Van Roy, and Christina Vespro, 2013, What determines Euro area bank CDS spreads?, *Journal of International Money and Finance* 32, 444–461.

Atkeson, Andrew G., Andrea L. Eisfeldt, and Pierre-Olivier Weill, 2014, Measuring the financial soundness of U.S. firms, 1926-2012, *Working Paper*.

Aymanns, Christoph, Carlos Caceres, Christina Daniel, and Liliana Schumacher, 2016, Bank solvency and funding cost, *IMF Working Paper* WP/16/64.

Bharath, Sreedhar T., and Tyler Shumway, 2008, Forecasting default with the Merton distance to default model, *The Review of Financial Studies* 21, 1339–1369.

BIS, 2015, Making supervisory stress tests more macroprudential: Considering liquidity and solvency interactions and systemic risk, *BIS Working Paper* No. 29.

Black, Fischer, and John Cox, 1976, Valuing corporate securities: some effects of bond indenture provisions, *Journal of Finance* 31, 351–367.

Chan-Lau, Jorge A., and Amadou N. R. Sy, 2006, Distance-to-default in banking: A bridge too far?, *IMF Working Paper* WP/06/215.

Flannery, Mark, Beverly Hirtle, and Anna Kovner, 2016, Evaluating the information in the Federal Reserve stress tests, *Journal of Financial Intermediation* (forthcoming).

Hasan, Iftekhar, Liuling Liu, and Gaiyan Zhang, 2016, The determinants of global bank credit-default-swap spreads, *Journal of Financial Services Research* (forthcoming).

Jessen, Cathrine, and David Lando, 2015, Robustness of distance-to-default, *Journal of Banking and Finance* (50, 493–505.

Lando, David, 2004, Credit risk modelling: Theory and applications, *Princeton University Press.*

Leland, Hayne E., 1994, Corporate debt value, bond covenants, and optimal capital structure, *Journal of Finance* 49, 157–196.

Merton, Robert C., 1974, On the pricing of corporate debt: the risk structure of interest rates, *Journal of Finance* 29, 449–470.

Nagel, Stefan, and Amiyatosh Purnanandam, 2015, Bank risk dynamics and distance to default, *Working paper.*

Schmitz, Stefan W., Michael Sigmund, and Laura Valderrama, 2016, Bank solvency and funding cost: New data and new results, *Working paper.*