

Modeling Municipal Yields with (and without) Bond Insurance

Ethan Namvar, Xiaoxia Ye, and Fan Yu*

This Version: October 10, 2015

*Namvar is from the University of California, Berkeley (namvar@haas.berkeley.edu), Ye is from Stockholm University (xiaoxia.ye@sbs.su.se), and Yu is from Claremont McKenna College (fyu@cmc.edu). We wish to thank Darrell Duffie and Francis Longstaff for helpful discussions. The HPC resources provided by the National Supercomputer Centre in Sweden (NSC) at Linköping University are gratefully acknowledged. The computations in the paper were partially performed with these resources.

Modeling Municipal Yields with (and without) Bond Insurance

Abstract

We develop an intensity-based model of municipal yields with (and without) bond insurance, making simultaneous use of the credit default swap (CDS) premiums of the insurers and both insured and uninsured municipal bond transactions. We estimate the model using 61 issuers with relatively continuous municipal bond trading from July 2007 to June 2008, and decompose the bond yield based on the estimated parameters. The model fits municipal yields as well as Duffee (1999)'s for corporate yields, and the decomposition reveals that the liquidity component plays a dominant role in the municipal yield spread, and is somewhat larger for insured bonds than uninsured bonds. Towards the end of the sample period, our model reproduces the “yield inversion” phenomenon documented by Bergstresser et al. (2010).

1 Introduction

The municipal bond insurance industry, long dominated by the big four of MBIA (Municipal Bond Insurance Association), Ambac (American Municipal Bond Assurance Corp.), FSA (Financial Security Assurance Inc.), and FGIC (Financial Guaranty Insurance Company), suffered a dramatic meltdown during the great recession for insuring structured finance products. While more than half of new municipal bonds were issued with credit enhancement in 2007, the so-called insurance penetration fell to 18 percent in 2008 and only 3.5 percent in 2012. The industry has changed significantly after the financial crisis, with many of the original players either having gone out of business or still sidelined by crisis-related legal problems. The new players, which are still trying to establish a reputation, include AGM (Assured Guaranty Municipal Corp.), BAM (Build America Mutual Assurance Company), and National (National Public Finance Guarantee Corp.). These insurers currently carry AA ratings according to the S&P's and lower ratings like A or Baa according to Moody's.¹

A great deal can be learned from what happened to the municipal bond insurers, also called monolines, during the financial crisis. Intuitively, the default risk of an insured municipal bond should be a function of the credit quality of both the issuer and the monoline, as well as how closely they are correlated with each other. In the pre-crisis years, the big four carried AAA ratings and rating agencies simply let insured bonds assume the top-notch ratings of the monolines; few people cared about quantifying the risk of insured municipal bonds. During the financial crisis and thereafter, a much greater variation in monoline credit quality has become the norm. It is therefore important, for both rating and valuation purposes, to develop a systematic understanding of the risk of municipal bonds subject to bond insurance. Yet, we continue to see the S&P's assigning an insurer's credit rating to every issue that it insures, which suggests that rating agencies still suffer from a lack of rigorous methodology for evaluating the risk of insured municipal bonds.²

¹BAM was created by former FSA employees in July 2012 and has been rated AA by the S&P's since its inception. AGM was formed when Assured Guaranty acquired FSA in 2009 and has been rated AA since March 2014. Finally, National is owned by MBIA and has been rated AA- since March 2014. For further background of the municipal bond insurance industry, see Wells Fargo (2008), Moldogaziev (2013), and Renick and Bonello (2014).

²In principle, it is clearly possible for an insured bond to achieve a higher rating than that of the insurer. For example, assuming that defaults are independent events, a Ba-rated insurer insuring a bond with an

We begin our analysis by regressing the municipal bond yield on a collection of bond characteristics, some known to proxy for bond liquidity, as well as a dummy variable indicating whether the bond is insured or uninsured; in the case of insured bonds, we also include the CDS premium of the corresponding monoline. Because of the difficulty in gathering accounting and financial information at different levels of municipalities, we use issuer-level fixed effects to control for the average issuer credit quality within our sample period. In another version of the regression, we identify a small number of municipal issuers that are traded in the CDS market and use their CDS premiums as a time-varying proxy for their intrinsic credit quality. Results indicate that the municipal bond yield is positively related to the age of the bond and the CDS premiums of the issuer and the insurer, and negatively related to the insurance dummy.

Based on these intuitive findings, we develop an intensity-based model for municipal yields, taking into account that insured bonds will not suffer any loss unless the issuer and the insurer both default before bond maturity. We model both defaults as doubly stochastic Poisson processes (Lando, 1998) that are independent given their intensities, allowing for the intensities to be correlated. Furthermore, our model can accommodate two features that are crucial in fitting municipal yields. First, the municipal bond market is known to be plagued by illiquidity in the form of high transaction costs and dealer markups, and the lack of transparency and dealers' substantial market power have been suggested as the underlying cause (see Harris and Piwowar 2006 and Green et al. 2007). To address this concern, we estimate an aggregate liquidity factor using all municipal bond transactions following the procedure of Driessen (2005), allowing for a direct compensation for the lack of liquidity in municipal yields as well as letting the default intensity load on the liquidity factor. In terms of modeling flexibility, while insured and uninsured bonds with the same issuer share a common default intensity, they can have different exposures to the aggregate liquidity factor. Second, we note that the state of the world in which an insured municipal bond suffers a loss of capital can look quite different from the one in which an issuer is defaulting on its

intrinsic rating of Ba would lower the probability of bond default to 0.01 percent per year, assuming a default rate of one percent per year for Ba-rated firms (see Exhibit 27 in Ou, 2011). This would allow the insured bond to easily clear the threshold for an Aa rating. While the independent default assumption needs to be relaxed, the basic message of the above example comes through even if defaults are correlated.

uninsured bond. While the latter is likely to be caused by idiosyncratic risk, the former requires the default of both the issuer and the insurer, which makes it much more likely to be systematic in nature. Therefore, it seems sensible to assume that insured bonds have a greater loss-given-default (LGD), or a lower recovery rate, than uninsured bonds.

To facilitate the estimation of our model, we specify the default intensities of the issuer and the monoline, as well as the liquidity factor, as squared-root diffusions or sums of independent squared-root diffusions. This allows for closed-form formulas for monoline CDS premiums and both insured and uninsured municipal bond prices, which produce the measurement equations in an Unscented Kalman Filter (UKF) procedure. Our estimation proceeds in three steps. In Step 1, we use the entire panel of municipal bond transaction prices to construct short-term and long-term “liquidity spreads,” exploiting bond age as an empirical proxy of bond liquidity, and extracting a single liquidity factor from these spreads. In Step 2, we estimate the default intensity for each of the four major monolines using the time-series of their respective CDS premiums. Each of the monoline default intensities is allowed to load on the liquidity factor, which is important because liquidity in the municipal bond market began to dry up around the same time when default concerns arose for the monolines. In Step 3, we select one municipal bond issuer and use the time-series of all of its bond transactions (both insured and uninsured bonds must be present) to estimate its default intensity, along with the differential liquidity exposures and recovery rates of insured and uninsured bonds. We implement this procedure using data from July 2007 to June 2008 for 61 municipal bond issuers, which include 17 states, 17 cities, and 17 local school districts, among others.³

Our results show that the model fits municipal bond prices quite well. The median RMSE in terms of bond yields is around 12 bps for insured bonds and 9 bps for uninsured bonds, similar to that of Duffee (1999) from fitting the corporate bond yields of 161 firms. For the median issuer, the issuer-specific part of the default intensity is mean-reverting

³These issuers are selected because their bonds have relatively continuous trading, thus offering good coverage of pricing data. We focus on 2007-08 because there is very little fluctuation in monoline CDS premiums prior to this period, and these monolines were essentially out of business with extremely high CDS premiums thereafter. A timeline of rating agencies’ monoline-related rating actions during 2007-08 can be found at http://afgi.org/resources/Subprime_Crisis_Timeline.pdf.

under the physical measure with a long-run mean of 55 bps, and is explosive under the risk-neutral measure, implying an upward-sloping default-related yield component, similar to Duffee’s findings for investment-grade corporate issuers. As expected, we find much lower recovery rates on insured bonds than those on uninsured bonds. The average recovery rate for insured bonds is 0.142 while that of uninsured bonds is 0.526, the latter being consistent with Moody’s historical muni recovery rates. Using the estimated model parameters for the 61 issuers, we decompose the municipal yield spread of a hypothetical bond into default and liquidity components. We find that the liquidity component plays a dominant role in the decomposition, accounting for nearly 80 percent of the yield spread for uninsured bonds and even higher percentages for insured bonds.⁴ For the median issuer, the liquidity component of insured bonds is about 15 bps larger than that of uninsured bonds. For the selected few issuers with good CDS coverage, we find that their bond-implied default intensities resemble the CDS-implied counterparts, but there are also substantial differences that might tempt a fixed income arbitrageur.

The extant literature on municipal bonds has had to contend with several interesting puzzles. First, the long-end of the municipal yield curve appears too high relative to the Treasury yield curve using the marginal tax rate implied at the short-end. Green (1993) resolves this tension by showing that investors can structure a bond portfolio to replicate the pre-tax cash flows of a taxable bond, while at the same time lowering the tax burden on its coupon payments. Because long-term bonds generate most of their cash flows from coupons, this strategy is more effective at long maturities, pulling down taxable yields relative to tax-exempt yields at the long end.⁵ We take as given the relation between after-tax and pre-tax discount factors in Green’s model, while leaving the marginal tax rate as one of the parameters to be estimated.

Recently, Bergstresser et al. (2010) document a second puzzle in municipal bonds. Specifi-

⁴This stands in contrast with the yield spread decomposition for corporate bonds—Longstaff et al. (2005) find liquidity to account for up to 50 percent of the yield spread for bonds rated AAA/AA, while the percentages are lower on bonds rated below AA.

⁵Chalmers (1998) examines the relative yields of Treasury bonds and pre-funded municipal bonds (those backed by Treasury bonds). He finds that these municipal yields display the same tendency to be too high relative to Treasury yields, concluding that default risk does not explain the muni puzzle. Wang et al. (2008) introduce default and liquidity components into Green’s model and show that the estimated tax rate is stable across maturity and credit rating.

cally, they use cross-sectional regressions to show that bonds insured by “troubled monolines” have higher yields than uninsured bonds with the same S&P’s issuer credit rating, both during and after the great financial crisis. Moreover, they use Green et al. (2007)’s method to show that municipal bond dealers charge a higher markup on sales to customers for these insured bonds than uninsured bonds. Consistent with their finding, our results from the top 61 municipal bond issuers demonstrate that insured bonds command a somewhat larger liquidity discount than uninsured bonds, and this difference results in yield inversion at the end of our sample period, when some of the monolines became “troubled.”

Our study is similar to Longstaff (2011) in the sense that both exploit term structure modeling techniques to extract information from municipal bond yields. While Longstaff focuses on the very short-end through municipal swaps and infers a time-varying marginal tax rate (as well as its risk premium) embedded in the one-week MSI (Municipal Swap Index) rate, our model is estimated using municipal bonds of all maturities and considers the effect of tax-exemption, default, liquidity, and insurance on municipal bond yields. Though our tax rate is treated as a constant, our median estimate of 50 percent during our sample period of 2007-08 is close to Longstaff’s time-varying estimate for the same period.

Another study that is similar to ours is Wang et al. (2008), who decompose municipal yields into liquidity and non-liquidity components by estimating an extension of Green (1993)’s model that accounts for sensitivities of the municipal yield to an aggregate liquidity factor. They do not model the time variation of default risk or the effect of bond insurance. Moreover, they estimate a much smaller liquidity component than ours because their estimates only include the part of the liquidity discount that covaries with the aggregate liquidity factor. In contrast, our liquidity discount incorporates a constant term as well as sensitivity to the liquidity factor. The constant part of the liquidity discount cannot be separated from the issuer’s credit risk if not for the presence of insured bonds in our estimation—as the insurer’s credit risk approaches zero, the insured bond becomes virtually risk-free, and its yield above and beyond the risk-free level must be attributed to a liquidity discount. Once we pin down the liquidity discount for insured bonds, the issuer default risk can be estimated. Then, since uninsured bonds share the same issuer default intensity with insured bonds, their liquidity discount can also be identified. Thus, the necessity for the

joint estimation of insured and uninsured bonds is a unique feature of our approach.

Last but not least, Ang et al. (2014) extract a daily zero-coupon yield curve from all pre-funded (hence default-free) municipal bonds and adjust the curve downward for the liquidity difference between pre-funded and regular bonds. They then use this adjusted zero curve to compute a default-free yield for any regular municipal bond. They also use the highest federal income tax rate to infer an after-tax yield for a hypothetical Treasury bond with the same cash flows as the municipal bond under consideration. These two yield benchmarks allow for a decomposition of any municipal yield into default, liquidity, and tax components. One difference between our methodology and Ang et al. (2014)'s is that our yield decomposition derives from an intensity-based model exploiting the pricing of a single issuer's insured and uninsured bonds, without relying on zero curves constructed from groups of municipal bonds that could potentially differ in their credit risk. Another difference is that the tax rate in our model is a parameter to be calibrated at the issuer level, which allows for clientele effects across municipal issuers. Looking beyond these differences, however, the dominant proportion of the liquidity component in the municipal yield spread is a finding shared by both papers.

The remainder of our paper is organized as follows. Section 2 explains the sources of our data and offers a preliminary analysis of municipal bond yields using panel regressions. Section 3 then presents an intensity-based model of municipal bond pricing with the aforementioned elements of default, liquidity, and insurance. Section 4 contains the results of model estimation and the ensuing municipal yield decomposition. We conclude with Section 5.

2 Preliminary Analysis

In this section, we use panel regressions to gain some insights into the various determinants of municipal bond yields, which will be useful for the development of a formal model.

2.1 Data

In Figure 1, we plot the five-year CDS premiums of Ambac, FGIC, FSA, and MBIA, the four major municipal bond insurers, from January 2007 to December 2008. Evidently, the premiums were extremely low and stable prior to July 2007 for all four monolines, but reached dizzying heights in 2008. In particular, FGIC's five-year premium stayed above 5,000 bps for most of the second half of 2008, and Ambac and MBIA's were over 2,000 bps in extended stretches as well. The only exception is FSA, presumably because it was the only major monoline to have reduced its exposure to structured credit products by half between 1995 and 2010 (Moldogaziev, 2013). While we intend to exploit changes in monoline credit quality in our estimation of the effect of bond insurance, variations of this magnitude could necessitate more sophisticated econometric specifications, such as those with regime switches and jumps, that would overly complicate our simple setup. Because of this concern, we limit our sample period to July 2007 to June 2008, where monoline CDS premium variations are substantial but not excessively large.

To select the municipal bonds used in our analysis, we apply several filters to those included in the Mergent Municipal Bond Securities Database. Specifically, we select tax-exempt investment-grade general obligation bonds with semi-annual fixed coupons and no embedded options or sinking fund provisions. In addition, these bonds must be either uninsured or insured by Ambac, FGIC, FSA, or MBIA.⁶ We then match the characteristics of these bonds with corresponding trades in the MSRB (Municipal Securities Rulemaking Board) historical transactions data, which has been used in virtually all muni-related research since mandatory trade reporting by municipal bond dealers began in 1997. Three types of trades are reported in the MSRB data: dealer sales to customers, dealer purchases from customers, and inter-dealer transactions. We focus exclusively on dealer sales to customers because 1) these are the majority among the three types of trades, and 2) we want to avoid modeling the bid-ask bounce. This procedure results in a total of 441,103 dealer-to-customer sales on 64,771 bonds from 7,425 unique municipal bond issuers during our one-year sample

⁶In the 2008 version of the Mergent Municipal Bond Securities Database, there are 2.28 million unique securities, about half of which are insured. Among the insured bonds, approximately 90 percent are insured by the four major monolines. The percentages are: 26 percent (MBIA), 25 percent (FSA), 22 percent (Ambac), and 17 percent (FGIC).

period.

2.2 Summary Statistics

Panel (a) of Table 1 shows that the median bond in our sample is 2.91 years old and has a yield-to-maturity of 3.42 percent, a coupon rate of four percent, a maturity of 4.58 years, and an issue size of \$2.62 million. The median trade size is \$50,000. While both the trade size and issue size seem small, they are highly positively skewed with the means much larger than the medians, and the standard deviations in turn much larger than the means. For more detailed summary statistics and sample composition, the reader is referred to existing studies such as Downing and Zhang (2004), Harris and Piwowar (2006), Green et al. (2007), Bergstresser et al. (2010), and Chung et al. (2015), since all of these studies make use of the same MSRB transactions data from which our sample is derived.

Panel (b) of Table 1 then computes the average yield by bond insurer and Moody’s bond rating (categories with fewer than 100 observations are omitted). It seems that Ambac, FSA, and MBIA have a different business model from FGIC. While the first three monolines have long maintained their AAA credit rating and are able to “rent out” this top rating to the bonds that they insure, FGIC is content with letting its credit rating fluctuate below the top and insuring lower quality bonds. Comparing the average yields of insured and uninsured bonds, it seems that uninsured bonds have lower yields. This is not the “yield inversion” phenomenon documented by Bergstresser et al. (2010), however. All it means here is that perhaps bond insurance enables an insured bond to barely clear the threshold for the rating that it received, hence its yield is expected to come in at the high end among all bonds with that rating. To examine yield inversion, we should really compare yields on insured and uninsured bonds that share the same issuer, which we will do below in our panel regressions, or at least compare yields on insured and uninsured bonds with the same issuer credit quality, which is what Bergstresser et al. (2010) did when they used S&P’s underlying issuer credit rating.⁷

⁷It is still possible that the intrinsic credit quality of issuers of insured bonds is systematically lower than that of issuers of uninsured bonds, even when they share the same intrinsic issuer credit rating. For example, if investors treat all insured Aaa-rated bonds the same and insurers charge the same fee for lending their Aaa rating to Aa-rated issuers, then the lower quality ones among Aa-rated issuers will be more likely to demand

The bond rating that we use to generate Panel B comes from Mergent and represents Moody’s bond rating at issuance, which may not be an accurate representation of the credit quality of a bond at the time of observation. We have two responses to this concern. First, the estimation of our intensity-based model uses the time-series of bond price observations for a single issuer and does not require bond rating at all. Second, for a selected subset of bond issuers, we have a superior proxy for credit quality—the CDS premium of the issuer. This will be useful in the panel regressions below.

2.3 Regression Analysis of Municipal Bond Yields

To get a closer look at the determinants of municipal bond yields, we conduct the following panel regression:

$$y_{ijkt_s} = \alpha_j + \beta_s + d_i(a_0 + a_1 CDS_{k_s}) + \gamma' \text{bond characteristics}_{it} + \varepsilon_{ijkt_s}, \quad (1)$$

where i , j , k , t , and s denote bond, issuer, insurer, time of trade, and date of trade, respectively, $d_i = 1$ for insured bonds and zero otherwise, and CDS_k denotes the five-year CDS premium of the monoline k that provides insurance for an insured bond i . Among the bond characteristics are coupon rate, maturity, issue size, trade size, and age. We take the logarithm of issue size and trade size because of the large positive skewness documented in Table 1. Of course, we entertain the intuitive conjecture that $a_0 < 0$ and $a_1 > 0$, meaning that bond insurance generally reduces the municipal bond yield, but that the effect is weakened by the insurer’s declining credit quality.

Note that the issuer-level fixed effects account for the average issuer credit quality during our (relatively short) sample period, the trading day-level fixed effects incorporate market-wide changes that affect the yields on all municipal bonds, and the bond-level and trade-level characteristics allow for heterogeneity in yields across bonds that share the same issuer. Still, one might be concerned that there is no explanatory variable that captures the time-varying issuer credit quality—to the extent that issuer credit quality is correlated with monoline credit quality, we can see a biased estimate of a_1 . To address this issue, we identify 30

bond insurance, because the potential cost savings are greater for them. However, Bergstresser et al. (2015) seem to rule out this possibility using evidence from subsequent rating migrations.

municipal bond issuers that have CDS coverage during 2007-08, and we include their CDS premiums in the enhanced specification below:

$$y_{ijkt_s} = \alpha_j + \beta_s + d_i (a_0 + a_1 CDS_{ks}) + a_2 CDS_{is} + \gamma' \text{bond characteristics}_{it} + \varepsilon_{ijkt_s}. \quad (2)$$

The trade-off is that the sample size becomes much smaller.

Yet another criticism for yield regressions in general is that a linear specification cannot account for the inherently nonlinear relation between the bond yield and its explanatory variables. We address this concern by creating Aaa-, Aa-, and A-rated subsamples of short-term (maturity less than or equal to five years) and long-term (maturity greater than five years) bonds using Moody’s initial bond ratings from Mergent, where the level of credit risk within each subsample is relatively homogeneous. Although some bonds could be misclassified because their ratings may have changed, the extent of such changes is expected to be small based on Moody’s historical rating transitions for municipal bonds (see Exhibit 4 in Tudela et al., 2015).

To our knowledge, there are two existing studies of municipal bond yields and insurance-related effects using the MSRB data and a regression-based methodology. Chung et al. (2015) use a panel regression like ours with the monoline CDS premium, but they include neither uninsured bonds nor issuer-level fixed effects. Therefore, their results are silent on the relation between insured and uninsured bond yields, and could be affected by omitted variables that capture the differences in credit quality across the issuers. Bergstresser et al. (2010) conduct monthly cross-sectional regressions of the bond yield on a “troubled insurer” dummy using groups of both insured and uninsured bonds with the same S&P’s issuer credit rating. They are mainly focused on the differences between insured and uninsured bond yields over time, especially during the financial crisis. The main issue with their approach, of course, is that the credit quality of insured and uninsured bond issuers can be systematically different even within the same rating category.

The results of our regressions are presented in Table 2. Focusing first on the insurance effects, we find that the coefficient on the insurance dummy is mostly negative and significant and ranges from 4 to 11 bps. The coefficient on the monoline five-year CDS premium is mostly

positive and significant, though the magnitude is small given the large swing of monoline credit quality during 2007-08. Nonetheless, the size of the coefficient implies a 2 to 12 bps of increase in the municipal bond yield for every 1,000 bps of increase in the monoline five-year CDS premium. Interestingly, the larger values among the estimates of a_0 are associated with A-rated bonds, indicating that there is greater room for yield reduction from bond insurance among the lower quality issuers. Furthermore, the larger values among the estimates of a_1 are associated with Aaa-rated bonds, suggesting that insured bonds that attain the top rating are the most sensitive to their insurers' credit quality.

In the regression with issuer CDS premiums, we group all bonds together because the sample size is much smaller. This regression shows a close to one-to-one relation between the issuer CDS premium and the municipal bond yield—the coefficient is 0.01 percent, or one bp, for every one bp increase in the issuer CDS premium. Meanwhile, the insurance dummy has a coefficient of -0.078 , which is significant at the one-percent level. Interestingly, the effect of the monoline CDS premium weakens and is no longer statistically significant. We do know through additional analysis, however, that this is more likely to be attributed to the reduction in sample size than the correlation between monoline and issuer CDS premiums.

Moving on to variables usually considered as indicators of bond liquidity, we find that there is no consistent relation between the bond yield and the issue size of the bond. Trade size, on the other hand, has a consistently negative and significant coefficient among short-term bonds. This fits well with the empirical findings that municipal bond trading costs and dealer markups are a downward-sloping function of trade size (Harris and Piwowar, 2006; Green et al., 2007). However, trade size is not significant when we use a long-term bond sample. In contrast, bond age is positive and significant for short-term bonds, and is positive and significant in the Aaa-rated long-term bond sample, which contains the majority of the long-term bonds. It appears that bond age is the most robust proxy for bond liquidity among the three that we have considered.

Overall, our panel regressions offer many interesting insights into the necessary ingredients for pricing municipal bonds; they demonstrate that issuer default risk, monoline default risk, and liquidity effects all play a role. Yet, a linear specification can only be stretched so far. For example, it would be silly to extrapolate the monoline CDS premium into the

10,000 bp range using the estimate for a_1 just because that level was reached by FGIC during 2008. Similarly, the one-to-one relation between the municipal bond yield and the issuer CDS premium is likely to break down when the monoline CDS premium is extremely low. While these can all be addressed in ad hoc ways in a regression-based framework, we will now turn to the construction of a formal municipal bond pricing model.

3 Model

Our model for the pricing of municipal bonds contains several novel features beyond the standard reduced-form model with latent issuer default factors (e.g., Duffee, 1999). First, we incorporate the credit enhancement provided by the monolines for insured bonds. Second, we estimate an aggregate liquidity factor and allow it influence the pricing of individual municipal bonds. Third, our model accounts for the correlation between the aggregate liquidity factor and the default intensities of the monolines and the bond issuers.⁸ Fourth, we allow for differential recovery rates and liquidity discounts for insured and uninsured bonds.

3.1 Aggregate Liquidity Factor

For now, it suffices to assume the existence of an aggregate municipal bond liquidity factor l_t , which follows a square-root diffusion with parameters $(\alpha_l, \beta_l, \sigma_l)$ under the risk-neutral measure:

$$dl_t = (\alpha_l - \beta_l l_t) dt + \sigma_l \sqrt{l_t} dZ_{l,t}. \quad (3)$$

The detailed procedure for the construction of this factor is explained in the next section on model estimation.

⁸This feature is important as municipal bond liquidity likely experienced negative shocks during the financial crisis period of 2007-08, when the market became increasingly concerned about corporate and municipal default risk. In this sense, our modeling framework is an improvement over that of Longstaff et al. (2005) and Driessen (2005), who treat their bond liquidity factor as independent of the credit risk factors.

3.2 Monoline CDS Pricing

We assume that the default intensity of the monoline $\lambda_{m,t}$ is given by the following:

$$\lambda_{m,t} = c_0 + c_1 l_t + \lambda'_{m,t}, \quad (4)$$

with $\lambda'_{m,t}$ following a square-root diffusion under the risk-neutral measure with parameters $(\alpha_m, \beta_m, \sigma_m)$:

$$d\lambda'_{m,t} = \left(\alpha_m - \beta_m \lambda'_{m,t} \right) dt + \sigma_m \sqrt{\lambda'_{m,t}} dZ_{m,t}, \quad (5)$$

and $Z_{m,t}$ being independent of $Z_{l,t}$. The parameter c_1 incorporates the potential correlation between the monoline default intensity and the aggregate liquidity factor.

For simplicity, we also assume that the default-free interest rate r_t is independent of other random sources in the model, and that the price of a default-free zero-coupon bond with a face value of \$1 and maturity t is $D(t)$. With payment dates denoted as $0 < t_1 < t_2 < \dots < t_n$, the default time τ_m , and the CDS premium c_m , the present value of future quarterly CDS premium payments is:

$$\frac{c_m}{4} E \left(\sum_{i=1}^n e^{-\int_0^{t_i} r_s ds} \mathbf{1}_{\{\tau_m > t_i\}} \right) = \frac{c_m}{4} \sum_{i=1}^n D(t_i) E \left(e^{-\int_0^{t_i} \lambda_{m,s} ds} \right) = \frac{c_m}{4} \sum_{i=1}^n D(t_i) \Phi_m(t_i), \quad (6)$$

where

$$\begin{aligned} \Phi_m(t) &= A_m(t) A_l(t; c_1) \exp \left(-c_0 t + B_m(t) \lambda'_{m,0} + B_l(t; c_1) c_1 l_0 \right), \\ A_m(t) &= A(t; \alpha_m, \beta_m, \sigma_m, 1), \\ B_m(t) &= B(t; \beta_m, \sigma_m, 1), \\ A_l(t; z) &= A(t; \alpha_l, \beta_l, \sigma_l, z), \\ B_l(t; z) &= B(t; \beta_l, \sigma_l, z), \end{aligned}$$

and the functions A and B are defined in Appendix A.

On the other hand, the present value of the CDS seller's protection payment is:

$$E \left(e^{-\int_0^{\tau_m} r_s ds} \mathbf{1}_{\{\tau_m \leq t_n\}} W_{m,\tau_m} \right), \quad (7)$$

where

$$W_{m,\tau_m} = w_m - c_m \left(t - \frac{\lfloor 4\tau_m \rfloor}{4} \right).$$

In this expression, w_m is a constant reflecting the LGD of the bonds underlying the monoline CDS contracts.⁹ The term being subtracted from w_m reflects the accrued CDS premium from the previous payment date to the time of default, with $\lfloor 4\tau_m \rfloor$ denoting the largest integer smaller than $4\tau_m$. We can rewrite (7) as:

$$\int_0^{t_n} D(t) \Psi_m(t) W_{m,t} dt, \quad (8)$$

where

$$\begin{aligned} \Psi_m(t) &= [c_0 A_m(t) A_l(t; c_1) + A_m(t) (G_l(t; c_1) + H_l(t; c_1) c_1 l_0) \\ &\quad + A_l(t; c_1) (G_m(t) + H_m(t) \lambda'_{m,0})] \exp \left(-c_0 t + B_m(t) \lambda'_{m,0} + B_l(t; c_1) c_1 l_0 \right), \\ G_m(t) &= G(t; \alpha_m, \beta_m, \sigma_m, 1), \\ H_m(t) &= H(t; \alpha_m, \beta_m, \sigma_m, 1), \\ G_l(t; z) &= G(t; \alpha_l, \beta_l, \sigma_l, z), \\ H_l(t; z) &= H(t; \alpha_l, \beta_l, \sigma_l, z), \end{aligned}$$

and the functions G and H are defined in Appendix A. Note that we have relied on the independence between l_t and $\lambda'_{m,t}$ when deriving the expressions for $\Phi_m(t)$ and $\Psi_m(t)$.

Equating (6) and (8), we can solve for the CDS premium as a function of the current value of the factors, $\lambda'_{m,0}$ and l_0 . In general, if we choose a finer set of grid points, $0 = s_0 < s_1 < \dots < s_m = t_n$, Δt apart from each other, to approximate the integral in (8), the CDS premium will be:

$$c_m = \frac{w_m \Delta t \sum_{j=1}^m D(s_j) \Psi_m(s_j)}{\frac{1}{4} \sum_{i=1}^n D(t_i) \Phi_m(t_i) + \Delta t \sum_{j=1}^m D(s_j) \Psi_m(s_j) \left(s_j - \frac{\lfloor 4s_j \rfloor}{4} \right)}. \quad (9)$$

When these grid points coincide with the quarterly CDS premium payment dates, $\Delta t = 1/4$,

⁹For pricing CDS contracts, it is conventional to assume a constant fractional recovery of the bond's par value. We assume 40 percent, which means that we set $w_m = 0.6$. Our later results are not sensitive to changes in this parameter.

the accrued premium term vanishes, and (9) simplifies to:

$$c_m = \frac{w_m \sum_{i=1}^n D(t_i) \Psi_m(t_i)}{\sum_{i=1}^n D(t_i) \Phi_m(t_i)}. \quad (10)$$

Alternatively, if we follow Berndt et al. (2008) to use the midpoints between the quarterly payments, then letting $s_i = (t_{i-1} + t_i) / 2$, $i = 1, 2, \dots, n$, we have:

$$c_m = \frac{w_m \sum_{i=1}^n D(s_i) \Psi_m(s_i)}{\sum_{i=1}^n D(t_i) \Phi_m(t_i) + \frac{1}{8} \sum_{i=1}^n D(s_i) \Psi_m(s_i)}. \quad (11)$$

In practice, the second approximation is much more accurate and we use it in the rest of our empirical analysis.

3.3 Municipal Bond Pricing

We assume that a municipal issuer with intrinsic default time τ_i and default intensity λ_i issues zero-coupon bonds with maturity T . Moreover, these bonds are insured by a monoline insurer, so that the promised principal will be repaid if either the issuer or the monoline remains healthy before T (assuming zero recovery given default). Therefore the payoff of the bond is:

$$X_T = 1_{\{\tau_i > T \text{ or } \tau_m > T\}}. \quad (12)$$

With municipal bonds, we need to use the after-tax discount factor $M(T)$, which represents the present value of an after-tax dollar at time T , to discount the payoffs. Let the marginal tax rate be η . According to Green (1993, equation 33), $M(T)$ can be obtained through the pre-tax discount factor $D(T)$, which in turns comes from bootstrapping Treasury bond yields:

$$M(T) = \frac{D(T)}{1 - \eta(1 - D(T))}. \quad (13)$$

The price of the insured municipal zero-coupon bonds with maturity T is then given as:

$$\begin{aligned} \frac{v(T)}{M(T)} &= E(1_{\{\tau_i > T \text{ or } \tau_m > T\}}) \\ &= E\left(e^{-\int_0^T \lambda_{i,s} ds} + e^{-\int_0^T \lambda_{m,s} ds} - e^{-\int_0^T (\lambda_{i,s} + \lambda_{m,s}) ds}\right). \end{aligned} \quad (14)$$

The above result invokes the independence between the risk-free rate and the default intensities. In addition, we follow the standard assumption in reduced-form credit risk models—that the default times of the issuer and the monoline are independent given their intensities.

To allow for a liquidity discount driven by the aggregate liquidity factor, we extend the above pricing relation slightly so that:

$$\frac{v(T)}{M(T)} = E \left[e^{-\int_0^T \gamma_{i,s} ds} \left(e^{-\int_0^T \lambda_{i,s} ds} + e^{-\int_0^T \lambda_{m,s} ds} - e^{-\int_0^T (\lambda_{i,s} + \lambda_{m,s}) ds} \right) \right], \quad (15)$$

where the liquidity discount

$$\gamma_{i,t} = c_2 + c_3 l_t. \quad (16)$$

To further evaluate the expectation in the above equation, we allow for the possibility that the default intensity of the municipal issuer is correlated with the aggregate municipal bond liquidity factor,

$$\lambda_{i,t} = c_4 + c_5 l_t + h_{i,t}, \quad (17)$$

with h_i specified as a square-root process under the risk-neutral measure:

$$dh_{i,t} = (\alpha_i - \beta_i h_{i,t}) dt + \sigma_i \sqrt{h_{i,t}} dZ_{i,t}, \quad (18)$$

where Z_i is independent of Z_m driving λ'_m as well as Z_l driving the liquidity factor l . Together, (17) and (4) complete a one-factor specification for the default intensities of the monoline and bond issuer, with the liquidity factor as the source of correlated intensities. Diagnostics tests in the results section will show that this is an adequate description of the comovement of the intensities.

Under these assumptions, we have:

$$\begin{aligned} & E \left(e^{-\int_0^T (\lambda_{i,s} + c_3 l_s) ds} \right) \\ &= A_l(T; c_3 + c_5) A_i(T) \exp(-c_4 T + B_l(T; c_3 + c_5) (c_3 + c_5) l_0 + B_i(T) h_{i,0}), \end{aligned} \quad (19)$$

$$\begin{aligned} & E \left(e^{-\int_0^T (\lambda_{m,s} + c_3 l_s) ds} \right) \\ &= A_l(T; c_3 + c_1) A_m(T) \exp\left(-c_0 T + B_l(T; c_3 + c_1) (c_3 + c_1) l_0 + B_m(T) \lambda'_{m,0}\right), \end{aligned} \quad (20)$$

$$\begin{aligned} & E \left(e^{-\int_0^T (\lambda_{i,s} + \lambda_{m,s} + c_3 l_s) ds} \right) \\ &= A_l(T; c_3 + c_1 + c_5) A_m(T) A_i(T) \times \\ & \exp\left(- (c_0 + c_4) T + B_l(T; c_3 + c_1 + c_5) (c_3 + c_1 + c_5) l_0 + B_m(T) \lambda'_{m,0} + B_i(T) h_{i,0}\right). \end{aligned} \quad (21)$$

As with previous definitions, we have

$$A_i(t) = A(t; \alpha_i, \beta_i, \sigma_i, 1),$$

$$B_i(t) = B(t; \beta_i, \sigma_i, 1).$$

We now extend the notation above to denote the price of a municipal bond with maturity T , coupon rate c , and recovery rate δ as $v(T, c, \delta)$. Here, it is assumed that for each unit of promised payment at t (including both principal and coupon payments), a fraction δ units are paid at t when default occurs before t . This is the well-known “recovery of Treasury” assumption introduced in Jarrow and Turnbull (1995) (refer to the appendix of Duffee (1999) for a similar treatment of corporate bonds). It implies that:

$$v(T, c, \delta) = \delta M(T, c) + (1 - \delta) v(T, c, 0), \quad (22)$$

where

$$M(T, c) = \sum_{i=1}^n M(T_i) c + M(T) \quad (23)$$

is the price of a hypothetical default-free municipal bond with coupon dates T_i , $i = 1, 2, \dots, n$, maturity T and coupon rate c , and

$$v(T, c, 0) = \sum_{i=1}^n v(T_i, 0, 0) c + v(T, 0, 0) \quad (24)$$

is the price of a defaultable municipal bond with coupon dates T_i , $i = 1, 2, \dots, n$, maturity T , coupon rate c , and zero recovery. With $v(T, 0, 0)$ given in (15), (22)-(24) complete the specification of insured municipal coupon bond pricing.

We should note that in (22), the first term on the right hand side reflects cash flows that will be received by bondholders regardless of default. Because of its risk-free nature, we do not apply a liquidity discount to this term. The second term, however, is affected by the default risk of the bond, and is therefore subject to a liquidity discount through (24) and (15). This treatment differs from Longstaff et al. (2005), who apply a liquidity discount to the entire bond price, including promised coupon and principal payments, as well as the fractional recovery of the bond's face value. While this difference is likely to be immaterial for our results, it does imply a less direct connection between the assumed liquidity discount $\gamma_{i,t}$ and the liquidity component of municipal bond yields in our case, since the latter is also affected by the recovery rate through (22). A discussion of this subtlety can be found in Section 4.4 below.

So far, we have focused on the pricing of insured bonds. The pricing of an uninsured municipal zero-coupon bond with zero recovery is given by:

$$\frac{v(T)}{M(T)} = E \left(e^{-\int_0^T (\lambda_{i,s} + \gamma_{i,s}) ds} \right), \quad (25)$$

which can simply be considered as a special case of (15) with $\lambda_{m,t}$ set to a very large value. The corresponding uninsured coupon bond with non-zero recovery can be dealt with similarly using (22)-(24). In the empirical estimation of our model, we will allow the liquidity discount and recovery rate to differ across insured and uninsured bonds. This means that the liquidity discount in (15) and (25), as well as the recovery rate in (22), will subsequently be labelled by “in” for insured bonds and “un” for uninsured bonds.

3.4 Identification Issues

3.4.1 Liquidity Discounts

Comparing the pricing of insured bonds in (15) and uninsured bonds in (25), we see that insured bond pricing naturally incorporates variations of monoline credit quality as reflected

in the monoline default intensity λ_m , while this does not (and it should not) enter the pricing of uninsured bonds. Also, the same issuer default intensity λ_i enters both equations, although in different ways as it interacts with the effect of λ_m in the case of insured bonds. On the other hand, we are completely free to allow different liquidity discounts, $\gamma_i^{un} = c_2^{un} + c_3^{un}l_t$ and $\gamma_i^{in} = c_2^{in} + c_3^{in}l_t$, for the two types of bonds. In fact, Bergstresser et al. (2010) suggest that lower liquidity on insured bonds could be the reason why their yields were higher than those on otherwise similar uninsured bonds during as well as after the financial crisis.

The uninsured bond's liquidity parameters (c_2^{un}, c_3^{un}) , however, cannot be identified from uninsured bond prices alone. This can be seen through (25), where we have the sum of the issuer intensity and the uninsured bond's liquidity discount:

$$\lambda_{i,t} + \gamma_{i,t}^{un} = c_2^{un} + c_4 + (c_3^{un} + c_5)l_t + h_{i,t}. \quad (26)$$

Using uninsured bond prices, therefore, we can only identify $c_2^{un} + c_4$ and $c_3^{un} + c_5$, but not each coefficient separately. With insured bonds, however, the crucial middle term in (15) contains:

$$\lambda_{m,t} + \gamma_{i,t}^{in} = c_2^{in} + c_0 + (c_3^{in} + c_1)l_t + \lambda'_{m,t}. \quad (27)$$

Since c_0 and c_1 have been identified in an earlier stage of monoline intensity estimation using CDS data, the insured bond's liquidity parameters (c_2^{in}, c_3^{in}) can be identified using insured bond prices alone. This means that the issuer intensity parameters (c_4, c_5) are now identifiable using insured bonds, and consequently (c_2^{un}, c_3^{un}) can also be pinned down. This discussion therefore shows the importance of using insured and uninsured bonds jointly in our model estimation.

3.4.2 Recovery Rates

Another ingredient of our model that can help bring insured and uninsured bond yields closer is the recovery rate. On the surface, it seems paradoxical for insured and uninsured bonds issued by the same municipality to command different recovery rates. After all, there are no systematic differences in seniority between these bonds, and both are backed by the same municipal assets and tax revenues. However, the (much less likely) states of the world

in which an insured bond suffers a loss of capital are probably quite different from those (more general cases) in which a municipality is defaulting on its uninsured obligations. For instance, when Orange County, California, defaulted in 1994, it ranked as the largest municipal bankruptcy at the time. While it was the Fed's interest rate actions that precipitated the loss on the Orange County Investment Pool, the bankruptcy was ultimately attributed to the investment decisions of a rogue county treasurer. What this means is that municipal default risk is generally not very systematic.¹⁰

While a significant number of the 80 municipal defaults during 1970-2014 (see Tudela et al., 2015) involved insured bonds, all but two were repaid in full by the monolines. The two exceptions are the sewer bonds of Jefferson County, Alabama insured by XLCA and FGIC, and the Las Vegas Monorail bonds insured by Ambac. Both defaults occurred during or shortly after the great financial crisis, and Moody's ultimate recovery rates on these bonds are 54 percent and two percent, respectively. The fact that the monolines had never failed to fully repay any defaulted insured bonds from their inception in 1971 to 2008, a period of nearly 40 years, suggests two implications. First, there was virtually no historical data to estimate recovery rates on insured bonds reliably before the financial crisis. The two observations we now have are certainly consistent with the bimodal distribution of muni recovery rates estimated by Tudela et al. (2012, Exhibit 17), with the two modes very close to either zero or one. Second, investors probably considered the joint default of a municipal issuer and the monoline that insures its bonds to be so unlikely that it is essentially an "end-of-the-world" scenario where nothing can be recovered. In any case, we will let the empirical data speak to the risk-neutral recovery rates implied from insured and uninsured bonds.

¹⁰This general sense also comes through by inspecting the number of municipal defaults during 1970-2014 (Exhibit 11 in Tudela et al., 2015). Annual defaults were generally few and there was no dramatic increase in the early 1990s, the early 2000s, or during 2008-09 for the recent great recession. This stands in sharp contrast with the temporal pattern of corporate defaults (Exhibit 1 in Ou, 2011).

4 Empirical Results

4.1 Liquidity Factor

The first step of our estimation is to construct an aggregate municipal bond liquidity factor. We do this by exploiting the relation between the municipal bond yield and the age of the bond as shown in Table 2. Our procedure follows Driessen (2005), who identifies a similar relation between bond yield and bond age among corporate bonds. Building upon the regression results, we plot the bond yield against bond age separately for short-term (maturity ≤ 5 years) and long-term (maturity > 5 years) bonds. Figure 2 shows that, for both short-term and long-term bonds, bond yield initially increases with age, but then the relation becomes downward-sloping or flat when age is greater than three years (the median bond age is 2.91 in Table 1). The same figure also reveals that bond yield increases with age among the older bonds in our sample.

Taking advantage of these more detailed findings, we form four municipal bond portfolios for each day in our sample. Specifically, the first portfolio P1 contains all bonds with maturity ≤ 5 years and $7 \text{ years} \leq \text{age} < 10 \text{ years}$, the second portfolio P2 contains all bonds with maturity > 5 years and $7 \text{ years} \leq \text{age} < 10 \text{ years}$, the third portfolio P3 contains all bonds with maturity ≤ 5 years and $\text{age} \geq 10 \text{ years}$, and the fourth portfolio P4 contains all bonds with maturity > 5 years and $\text{age} \geq 10 \text{ years}$. The first two portfolios are equally-weighted, while the latter two are weighted to match the average duration and rating distribution of the first two portfolios. This way, the average yield differences between P3 and P1 and between P4 and P2 reflect only differences in bond liquidity and not interest rate risk or credit risk. We plot the short-term liquidity spread (P3-P1) and long-term liquidity spread (P4-P2) in Figure 3. We can see that both spreads exhibit an upward trend over the sample period, and that the long-term spread is generally higher than the short-term spread.

Next, we treat these daily “liquidity spreads” as if they were zero-coupon bond yields with maturities set to equal the variable durations of P1 and P2, respectively. We then assume that these yields are generated by a one-factor CIR model and proceed to estimate the liquidity factor and its parameters using a standard Kalman Filter, since the mapping from the factor to the zero-coupon bond yield is an affine function. We assume an extended

affine market price of risk (Cheridito et al., 2007) so that the liquidity factor process under the P -measure is:

$$dl_t = (\alpha_l^P - \beta_l^P l_t) dt + \sigma_l \sqrt{l_t} dZ_{l,t}^P, \quad (28)$$

We estimate the liquidity factor parameters as: $\alpha_l^P = 0.004$, $\beta_l^P = 2.423$, $\sigma_l = 0.106$, $\alpha_l = 0$, and $\beta_l^P = -0.487$. The physical parameters α_l^P and β_l^P imply a mean-reverting process under P with a long-run mean of around 17 bps. In contrast, the risk-neutral parameters α_l and β_l suggest an explosive process that explains the higher liquidity spread at the long maturities. The time-series of the filtered liquidity factor is represented in Figure 3. Like the liquidity spreads, it also exhibits an upward trend through the sample period, starting around 5 bps in July 2007 and rising above 20 bps in April and June 2008.

4.2 Monoline Default Intensities

The second step of our estimation is to extract the default intensities of the four major monolines using their CDS contracts. Equation (11) maps the state variables $\lambda'_{m,t}$ and l_t into the monoline CDS premium at time t . For the value of l_t , we simply substitute for the filtered liquidity factor from the preceding estimation, assuming that it is observed without error. Therefore, this is effectively a one-factor model of $\lambda'_{m,t}$ that can again be estimated using the Kalman Filter. However, the standard KF cannot be used here because the mapping (11) is nonlinear. In this case, one approach is to linearize the so-called “measurement equation” around the long-run mean or the one-period-ahead forecast of the state variable and then apply the standard KF. This “Extended Kalman Filter” (EKF) has been used by Duffee (1999) to estimate an intensity-based model for corporate bond pricing. We, however, use another related approach called the “Unscented Kalman Filter” (UKF), which better addresses nonlinearities in the measurement equation (see Appendix B for details of the UKF).

We specify the dynamics of $\lambda'_{m,t}$ under the physical measure using the same extended affine market price of risk as what we did for the liquidity factor:

$$d\lambda'_{m,t} = (\alpha_m^P - \beta_m^P \lambda'_{m,t}) dt + \sigma_m \sqrt{\lambda'_{m,t}} dZ_{m,t}^P, \quad (29)$$

where $Z_{m,t}^P$ is independent of $Z_{i,t}^P$. This allows us to write down a discretized “transition equation” for the square-root process, also required for the KF estimation. Finally, we conduct the estimation using monoline CDS premiums of the following maturities: six months, one, two, three, four, five, seven, and ten years. We assume that the five-year CDS premium, usually the most highly traded maturity, is observed without error, and use it to infer the value of the state variable $\lambda'_{m,t}$. The CDS premiums of other maturities are assumed to contain measurement errors.

Table 3 presents the estimated parameters for the four monolines and Figure 4 compares the actual and model-implied three-year and ten-year CDS premiums. Graphically, we can see that the fitting performance is quite good for Ambac, FSA, and MBIA even though their CDS premiums exhibited very large variations over the sample period. This good performance is also reflected in “VR,” or one minus the ratio of the variance of CDS pricing errors to the variance of actual CDS premiums, which is above 0.99 in these three cases. We also compute the “relative RMSE,” or the root mean squared CDS pricing error divided by the sample average of actual CDS premiums, which is below 12 percent in these three cases. For comparison purposes, the bid-ask spread in the CDS market is around five to ten percent of the CDS premium in recent years, and this percentage can become even higher for top quality obligors with very low CDS premiums. Therefore, the pricing errors in these three cases are certainly within the margin of tolerance.

We do note the relatively poor fitting performance for FGIC both graphically and as reflected in the lower VR and higher relative RMSE. FGIC has the highest level of CDS premiums among the four monolines, and its short-term CDS premiums (six-month and one-year) exceeded 10,000 bps in June 2008. Therefore, we did not expect great performance from such a simple specification to begin with. In fact, while a one-factor model like ours can match the level of the CDS curve well (e.g., the five-year CDS premium), it has limited ability in matching the slope of the CDS curve. Specifically, the risk-neutral drift of $\lambda'_{m,t}$ is $\alpha_m - \beta_m \lambda'_{m,t}$, which for positive α_m and β_m will generate a positive slope when the CDS level is low and a negative slope when the CDS level is high. But the relation between the CDS level and slope within our one-factor model is deterministic—explaining deviations from this relation would likely require an additional stochastic factor.

Among the model parameters, we see that $\alpha_m^P > 0$ and $\beta_m^P < 0$ imply an explosive $\lambda'_{m,t}$ process under P for Ambac, FSA, and MBIA, which is certainly consistent with their rising CDS levels during our sample period (see Figure 1). For FGIC, the drift of $\lambda'_{m,t}$ under P is positive except for its highest values in excess of 10,000 bps, which is also consistent with the sample path of its CDS premium. For the risk-neutral parameters, α_m and β_m are both positive for Ambac, FSA, and MBIA, implying mean-reverting behavior with implications for the slope of the CDS curve as we have just discussed. In the case of FGIC, it has $\alpha_m > 0$ and $\beta_m < 0$, rendering the drift always positive and suggesting that the model may have trouble simultaneously matching a positive slope when the CDS level is low and a negative slope when the CDS level is high.

For the volatility parameter σ_m , we find similar estimates among Ambac, FSA, and MBIA, and a much higher one for FGIC, again consistent with casual observations of their CDS premiums. Lastly, we recall that the parameter c_1 measures the sensitivity of the monoline default intensity to the liquidity factor. Our estimates of c_1 (positive in all cases) are close to zero for FSA, similar and close to one for Ambac and MBIA, and five times larger still for FGIC.

4.3 Issuer Parameters

The third and final step of our estimation is to identify the parameters associated with the municipal bond issuers. For each issuer i , the parameters to be estimated include $(\alpha_i, \beta_i, \sigma_i, \alpha_i^P, \beta_i^P)$ for $h_{i,t}$, the issuer-specific component of the default intensity $\lambda_{i,t}$; (c_4, c_5) , the part of the default intensity linear in the liquidity factor; $(c_2^{in}, c_3^{in}, c_2^{un}, c_3^{un})$, the liquidity discounts of uninsured and insured bonds; $(\delta^{in}, \delta^{un})$, the recovery rates of insured and uninsured bonds; and η , the marginal tax rate.¹¹ Consistent with previous assumptions, the physical dynamics of $h_{i,t}$ follows:

$$dh_{i,t} = (\alpha_i^P - \beta_i^P h_{i,t}) dt + \sigma_i \sqrt{h_{i,t}} dZ_{i,t}^P, \quad (30)$$

¹¹Besides these 14 parameters, we need the initial value of the state variable (assuming an initial variance of zero) to start the UKF recursion, and the two variances of pricing errors for insured and uninsured bonds, respectively. Therefore, we have a total of 17 parameters for each issuer.

where $Z_{i,t}^P$ is independent of both $Z_{l,t}^P$ and $Z_{m,t}^P$.

The pricing formulas (15) and (25) map l_t , $\lambda'_{m,t}$, and $h_{i,t}$ into prices of insured and uninsured muni zeros. As before, we substitute out l_t and $\lambda'_{m,t}$ using their filtered values, leaving this as a one-factor model of $h_{i,t}$ that can be estimated using the UKF. While the UKF procedure does not require regularly-spaced bond price observations, there needs to be enough of them to span the sample period as well as provide enough variation to pin down the large number of model parameters. Our procedure for selecting the issuers is as follows: First, for bonds with multiple dealer sales within the same trading day, we consolidate them into one observation using simple averaging. Second, we screen for issuers with at least 100 observations that include both insured and uninsured bonds, which are relatively evenly distributed over our sample period. To speed up computation, for cases where there are more than 100 observations, we truncate the data so that only bonds with the highest number of observations are kept. This procedure results in a sample for model estimation that contains 61 issuers.

Table 4 summarizes the estimation results. Because of the large number of issuers, we only present the first quartile, median, third quartile, and mean of the parameter estimates. First, we focus on the general performance of our model as measured by the VR and the relative RMSE. Across the 61 issuers, the median VR is 0.964 for insured bonds and 0.985 for uninsured bonds, showing that the model-generated bond price leaves less than four percent of the variation of the actual bond price unexplained for the median municipal bond issuer. Perhaps more direct for interpretation is the median value of the relative RMSE, which equals 0.617 percent and 0.439 percent for insured and uninsured bonds, respectively. Table 1 shows that the average maturity of the bonds included in our regression analysis is 5.6 years. Assuming a modified duration of five years, the corresponding RMSE in terms of bond yields would be around 12 bps for insured bonds and 9 bps for uninsured bonds. In comparison, the median bond yield RMSE across 161 corporate bond issuers is around 10 bps in Duffee (1999). Interestingly, Duffee's model contains two default-free factors filtered from a previous stage from Treasury bonds that together with an issuer-specific factor drive corporate bond yields. Ours, on the other hand, contains one liquidity factor and one monoline default factor filtered from previous estimations that together with an issuer-specific factor drive municipal

bond yields. We now know that our model performs just as well as Duffee’s model—at least for explaining the yields on the most highly traded municipal bonds.

Turning our attention to the parameters of the $h_{i,t}$ process, we see that the physical parameters α_i^P and β_i^P are mostly positive, with the median parameters implying a long-run mean for $h_{i,t}$ of 55 bps (0.076/13.842). On the other hand, the risk-neutral parameter α_i is uniformly close to zero and β_i is mostly negative, suggesting that $h_{i,t}$ is explosive under the Q measure. Thus the representative profile for our municipal issuer resembles a highly rated corporate issuer with low and upward-sloping credit spreads (see Section 5 in Duffee, 1999, for relevant discussions). This is not surprising when our 61 issuers include 17 states, 17 cities, and 17 school districts that all have high credit quality.

Our issuer default intensity also includes a constant term c_4 , which is close to zero for most issuers, plus c_5 times the liquidity factor. We see that c_5 seems to be negatively skewed based on its Q1, median, and Q3, yet its mean is slightly positive. In any case, its magnitude is much smaller than c_1 , which captures a strong positive dependence on the liquidity factor for the monoline default intensity. This suggests that the correlation between issuer credit risk and aggregate bond liquidity should not be a compelling concern for the modeler.

The remaining issuer parameters pertain to liquidity discounts, recovery rates, and the marginal tax rate, and their estimates offer interesting insights into the pricing of municipal bonds. For the liquidity discount parameters, we find $\gamma_{i,t}^{un} = c_2^{un} + c_3^{un}l_t$ to be much larger than $\gamma_{i,t}^{in} = c_2^{in} + c_3^{in}l_t$ across virtually the entire distribution of the parameters (e.g., Q1, median, Q3, and mean). Also, c_2 and c_3 are mostly positive for both insured and uninsured bonds, consistent with a discount of the bond value due to liquidity risk. Do these results imply, then, that insured bonds are generally much more liquid than uninsured bonds in the sense that they are discounted less? Surprisingly, the answer based on our decomposition of the muni yield is “No.” We will return to this question in the next subsection.

We find robust evidence that the recovery rate of insured bonds is heavily concentrated at zero with an average value across 61 issuers of 0.142, suggesting that investors expect little when the bond issuer and the bond insurer both default. In fact, both the Q1 and the median of the recovery rate for insured bonds are equal to zero, suggesting that our UKF procedure settled for a corner solution for the recovery rate of insured bonds in many

cases. In contrast, the recovery rate for uninsured bonds seems to follow a rather tight distribution around the mean value of 0.526 and median value of 0.581. Interestingly, a recent Moody’s municipal bond default and recovery study (Tudela et al., 2012) computes an average recovery rate of 65 percent based on 71 defaults from 1970 to 2011. It also shows that the recovery rate follows a bimodal distribution with one mode close to zero and the other mode close to one. There is no doubt that many cases where the recovery rate is exactly one include insured bonds that are made whole by the monolines. Therefore, the recovery rate on uninsured bonds is likely to be lower. Our estimate of the recovery rate for uninsured bonds is therefore consistent with Moody’s data.

Our estimate of the marginal tax rate η also follows a tight distribution with a mean value of 0.504 and a median value of 0.501. Longstaff (2011) estimates the marginal tax rate from municipal swaps that exchange the seven-day MSI rate for a fraction of the LIBOR rate. His estimate is time-varying and averages to about 40-50 percent during the 2007-08 period, which is consistent with our estimates obtained from a difference source—the municipal bond market. Borrowing an argument from Longstaff (2011), many of our 61 issuers are states, counties, cities, and local school districts. The investors in bonds issued by these entities likely differ in terms of their geographic locations. Therefore, our tax rate estimate probably reflects the aggregate federal, state, and local income tax burden of the marginal investor in these bonds. It is not difficult to conjure up an example where the total tax rate faced by an investor is close to our median estimate of 50 percent.

4.4 Decomposing Insured and Uninsured Bond Yields

Using the estimated model parameters, we now perform a decomposition of the municipal bond yield into default-free, default, liquidity, and insurance components (if the bond is insured). Our methodology is as follows. First, we note that the bond yield depends on the bond maturity and coupon rate, which are generally not the same across issuers. Therefore, we hold fixed a “hypothetical bond” with a maturity of 4.2 years and coupon rate of 5.25 percent, and use the model parameters to compute the municipal yield components. This way, the decomposition is solely attributed to the (issuer-specific) model parameters and not

differences in bond characteristics.

For each issuer and each day in the sample period, given the estimated tax rate and the Treasury discount factors, we use (13) to calculate the yield of a “default-free” municipal bond yield, y_0 . We then switch on the issuer default intensity (with the recovery rate for insured bonds) while keeping the liquidity discount at zero and allowing for no bond insurance. This generates an updated bond yield of y_1 . Then, while still keeping the liquidity discount at zero, we introduce the effect of the bond insurer. We assume that the hypothetical bond is insured by either FSA or MBIA, because we want to examine the effect of monoline credit quality on the yield decomposition, knowing that MBIA’s credit quality became much worse than FSA’s during our sample period. We update the bond yield to y_2 after this step. Finally, we turn on the liquidity discount for insured bonds and compute the bond yield in the full model as y_3 . We define the “default” component as $y_1 - y_0$, the “default+insurance” component as $y_2 - y_0$, and the “liquidity” component as $y_3 - y_2$.¹² The total of the “default+insurance” and “liquidity” components is $y_3 - y_0$, or the total yield spread of the hypothetical insured municipal bond. For the hypothetical uninsured bond, we have the same default-free yield, but use the recovery rate for uninsured bonds with the same issuer default intensity to compute a different default component, and we turn on the liquidity discount for uninsured bonds after that to obtain the liquidity component.

Table 5 summarizes the sample averages of the municipal yield components. First looking across all 61 issuers, we see that the liquidity component takes up a dominant share of the total yield spread for municipal bonds. For uninsured bonds, the median components show that the fraction of yield spread attributed to liquidity is 79 percent ($1.240 / (1.240 + 0.329)$), with 21 percent attributed to default. For insured bonds, since bond insurance has the effect of sharply reducing the default component (compare the “def” column with the “def+ins” column), the share of the yield spread attributed to liquidity is even higher—for the median case, this is 91 or 96 percent depending on the insurer.¹³ For a similar decompositions of corporate bond yield spreads, Longstaff et al. (2005, Table 4) find that AAA/AA-rated

¹²We find that the liquidity component is not sensitive to whether FSA or MBIA is used in the second step. In other words, while the choice of the insurer affects y_3 , it does not have a significant effect on $y_3 - y_2$.

¹³For example, for the case of MBIA, the liquidity share of the yield spread for the median case is $1.392 / (1.392 + 0.143) \approx 0.91$.

issuers have an average default component of 0.533 percent out of an average total yield spread of 1.036 percent relative to Treasury yields. While our municipal issuers have a lower default component of 0.329 percent, the liquidity component is much larger, boosting the total yield spread against a hypothetical default-free muni yield to 1.569 percent.

One interesting result is that the default component of insured bonds (not yet having incorporated insurance effects) is higher than that of uninsured bonds. This is apparently because the recovery rate on insured bonds are much lower than that on uninsured bonds, while they share the same issuer default intensity. In fact, different recovery rates are also the reason why the liquidity component for insured bonds is somewhat higher than that of uninsured bonds even though the c_2 and c_3 estimates of uninsured bonds are much larger than those of insured bonds (see Table 4). Specifically, (22) shows that the price of a municipal zero with recovery rate δ is equal to $1 - \delta$ times the price of a municipal zero with zero recovery plus δ times the price of a default-free zero. While the first part is subject to a liquidity discount following (15), the default-free part is not. Therefore, while it is true that c_2 and c_3 are much bigger for uninsured bonds, the associated liquidity discount is only applied to part of the total bond price when the median recovery rate estimate for uninsured bonds is about 0.5. In contrast, the median recovery rate estimate for insured bonds is zero, and the liquidity discount for insured bonds is applied to the entire bond price. In the end, our calculation reveals that the liquidity component is about 15 bps larger for insured bonds of the median issuer than for uninsured bonds (1.392 vs. 1.240).

Our 61 issuers include 17 states, 17 cities, and 17 unified school districts.¹⁴ Table 5 also summarizes the sample averages of the yield components for each of these three groups. There is little difference that sets these three types apart from each other, perhaps because they all have high credit qualities. We do see, however, that state-issued (city-issued) uninsured bonds have the lowest (highest) level of default component among the three groups. Also, city-issued uninsured bonds have the lowest liquidity component, suggesting that they are the most liquid among the three groups.

¹⁴Also included are five counties and five state-level public finance entities.

4.5 Time Variation of Yield Components

Examining the sample averages of the yield components gives us a good sense of the relative importance of default vs. liquidity and the effect of bond insurance a la FSA vs. MBIA. In the latter comparison, Table 5 shows that having FSA (instead of MBIA) as the bond insurer can lower the bond yield by around 10 bps (0.071 vs. 0.172) for the average issuer. In this subsection, we focus on how these yield components change over time.

Figure 5 plots the median yield spread components for the hypothetical bond over our sample period. The left-hand-side (right-hand-side) panels assume MBIA (FSA) as the insurer. We begin our discussion with the case of MBIA. In Panel (a), we see that the default component of the uninsured bond is much larger than the def+ins component of the insured bond at the beginning of our sample period. This is because the CDS premium of MBIA is low during 2007. However, the gap between the two components begins to narrow at the end of 2007 because of MBIA's deteriorating financial conditions.¹⁵ Panel (a) also shows that the liquidity component of the insured bond is visibly higher than that of the uninsured bond during the first half of our sample period. However, the uninsured bond is more sensitive to the aggregate liquidity factor, which causes its liquidity component to catch up to that of the insured bond toward the end of our sample period.¹⁶ With these components coming so close to each other, it is not surprising that Panel (e) shows that yield inversion occurs more frequently after April 2008. In contrast, we do not see as many incidences of yield inversion in Panel (f) for the case of FSA because its credit quality did not decline as much as MBIA's. Notably, the def+ins component of Panel (b) is much lower than that of Panel (a).

Our purpose here is to articulate the conditions under which yield inversion can occur. According to the numerical example just presented, those conditions have to do with a rapidly

¹⁵Another reason why the def+ins component of insured bonds can rival or even exceed the default component of uninsured bonds is the lower *constant* recovery rate we estimate for insured bonds. When the insurer is "teetering on default," as was the case with Ambac, MBIA, and FGIC in the middle of 2008, the recovery rate on insured bonds must be very similar to that of uninsured bonds by definition. Therefore, if we properly model this behavior through time-varying recovery rates, our model would not have produced such an "inversion" of the default components. We leave this model extension to future research.

¹⁶Note that in Table 4, the median value of c_3^{un} is 1.029, much larger than the median value of c_3^{in} , which is 0.107.

declining insurer credit quality as well as liquidity discounts that are greater for insured bonds. While our median parameter values are consistent with the liquidity component of insured bonds being about 15 bps higher than that of uninsured bonds, we must point out that these estimates are based on 61 most liquid municipal issuers. Our approach cannot be used to uncover yield inversion in the general population of hundreds of thousands of municipal bonds that are rarely traded. Having said that, Bergstresser et al. (2010, Table 15)’s regression estimate of the additional dealer markup for bonds insured by “troubled monolines” is around 10-15 bps during the first half of 2008, quite similar to the median estimate of the liquidity gap between insured and uninsured bonds in our smaller sample.

4.6 Bond-Implied vs CDS-Implied Default Intensities

Although liquidity is the largest component by far in the municipal yield spread, an examination of its default component can also yield many insights. This is especially true when we compare the bond-implied default intensity side-by-side with its CDS-implied counterpart, whenever we have CDS trading for municipal issuers. Unfortunately, CDS trading activities in municipalities only picked up at the end of 2007, most likely as investors began reacting to the monoline debacle. This means that we can only conduct this comparison for a selected few issuers. Specifically, we assume the same functional form for the CDS-based default intensity as in (17),

$$\lambda_{i,t}^{cds} = c_4^{cds} + c_5^{cds} l_t + h_{i,t}^{cds},$$

and of course we estimate its parameters from the CDS premiums of the municipality, not from its bonds. Procedurally, this estimation is identical to the estimation of the default intensities of the monolines.

Figure 6 plots the CDS-implied default intensity against the bond-implied default intensity that we estimated earlier. The figure presents three cases: the State of California, the City of New York, and the State of Illinois. NYC and California are the top two issuers in our sample by the number of bonds traded. All three are selected because they have good CDS coverage from July 2007 to June 2008.¹⁷ All three panels of Figure 6 reveal qualitatively

¹⁷Several other states also had CDS trading during this period. However, the trading was usually very thin, with only one or two maturities available, and with zero or spotty coverage before 2008.

similar patterns. Before November 2007, the CDS intensities of all three states were low and stable, presumably because investors did not worry about the default risk of large municipalities. In contrast, the muni bond market was showing higher volatility during this period, and one could find portfolio managers giving advice about “opportunities” in municipal bonds because their yields were rising relative to Treasuries (see Lim, 2007). Between November 2007 and April 2008, however, the CDS intensities began to rise and shadow movements in the bond intensities closely. Comparing the two intensities during this period, it seems that the bond intensity is much more volatile than the CDS intensity, perhaps due to the measurement noise in the muni bond prices. Nevertheless, the close relation between the two gives us confidence that one could still meaningfully extract default-related information from muni bonds, even when this is a market known for its illiquidity.

Figure 6 also shows an interesting divergence between the two intensities after April 2008. While the bond intensity precipitously declined, the CDS intensity flattened and even increased. An inspection of the ratio of the Bloomberg 10-year AAA muni yield to the 10-year Treasury yield confirms a decline into June 2008. Therefore, while the CDS market remained pessimistic about the credit risk of large municipalities, the municipal bond market “disagreed.” Interestingly, a fixed income arbitrageur who traded on this divergence would probably have profited, because the same muni-to-Treasury yield ratio rose dramatically during July to December 2008. In this case, the CDS market turned out to be correct.

4.7 Testing the Independence Assumptions

Several important assumptions in our model for the dynamics of issuer and monoline default intensities, as summarized in (17) and (4), are the independence between $\lambda'_{m,t}$ and l_t , $h_{i,t}$ and l_t , and $\lambda'_{m,t}$ and $h_{i,t}$. We therefore compute these correlations using the filtered estimates of the state variables. Since some of the intensities exhibit non-stationary behavior under the physical measure, we take the first difference of the state variables before computing the correlations.

Table 6 provides supportive evidence of the independence assumptions. The last row shows the correlations between l_t and each individual $\lambda'_{m,t}$ as well as the first and second

principal components (PCs) of the four $\lambda'_{m,t}$'s. The absolute values of these correlations are all below 0.18. The last column shows the summary statistics of the correlations between l_t and $h_{i,t}$, with a Q1, median, Q3, and mean of -0.059 , 0.01 , 0.069 , and 0.006 , respectively. Finally, the correlations between $h_{i,t}$ and each $\lambda'_{m,t}$ as well as the two PCs have quartiles and means that never exceed 0.13 in absolute value. These low correlations suggest that our one-factor structure describes the comovement of the liquidity factor, the four monoline intensities, and the 61 issuer intensities well.

5 Concluding Remarks

We construct an intensity-based model of municipal yields, similar to that of Duffee (1999) for pricing corporate bonds, but with features unique to the municipal bond market. One important feature is that a significant percentage of municipal bonds carry bond insurance, and their pricing is likely influenced by the credit quality of the insurer besides that of the municipal issuer. Another important feature is that municipal bonds trade infrequently, and their pricing is probably subject to substantial liquidity discounts.

We estimate the model in steps, first constructing an aggregate liquidity factor and estimating the default intensities of bond insurers from their CDS curves, and then disentangling issuer default risk, bond recovery rates, and liquidity discounts using insured and uninsured bond transactions of 61 issuers from July 2007 to June 2008. We find that the model generates pricing errors similar to that of Duffee (1999) for 161 issuers of corporate bonds, with a median RMSE of 9-12 bps. We use the model to decompose the municipal bond yield into default-free, default (with or without bond insurance), and liquidity components. Results indicate that the liquidity component accounts for more than 80 percent of the municipal yield spread, and that the median liquidity component is about 15 bps higher for insured bonds than uninsured bonds. The latter finding can generate “yield inversion” towards the end of our sample period, when bond insurers’ credit quality worsened. Other interesting findings include an estimated marginal tax rate that is consistent with Longstaff (2011)’s estimates based on municipal swaps, recovery rates on uninsured bonds consistent with Moody’s historical estimates but much lower recovery rates on insured bonds, and bond-implied default

intensities that resemble CDS-implied counterparts for issuers that have CDS trading.

These findings suggest that an intensity-based model like ours can be used to extract meaningful information about marginal tax rates, default probabilities, bond recoveries, and liquidity premiums from municipal bond prices. While many aspects of our model can be improved, we hope that it can serve as a useful tool and benchmark for academic researchers and industry practitioners interested in assessing the risks of municipal bonds.

Appendices

A Definition of Functions A , B , G , and H

Following Longstaff et al. (2005), for a default time τ with intensity λ described by a square-root process,

$$d\lambda_t = (\alpha - \beta\lambda_t) dt + \sigma\sqrt{\lambda_t}dZ_t,$$

the expectations below, which are useful in the computation of CDS premiums and municipal bond prices, can be evaluated as:

$$\begin{aligned} E\left(e^{-\int_0^t c\lambda_s ds}\right) &= A(t; \alpha, \beta, \sigma, c) e^{B(t; \beta, \sigma, c)c\lambda_0}, \\ E\left(c\lambda_t e^{-\int_0^t c\lambda_s ds}\right) &= (G(t; \alpha, \beta, \sigma, c) + H(t; \alpha, \beta, \sigma, c) c\lambda_0) e^{B(t; \beta, \sigma, c)c\lambda_0}. \end{aligned}$$

When $\beta^2 + 2c\sigma^2 > 0$,

$$\begin{aligned} A(t; \alpha, \beta, \sigma, c) &= e^{\frac{\alpha(\beta+\phi)}{\sigma^2}t} \left(\frac{1-\kappa}{1-\kappa e^{\phi t}}\right)^{\frac{2\alpha}{\sigma^2}}, \\ B(t; \beta, \sigma, c) &= \frac{\beta-\phi}{c\sigma^2} + \frac{2\phi}{c\sigma^2(1-\kappa e^{\phi t})}, \\ G(t; \alpha, \beta, \sigma, c) &= \frac{\alpha c}{\phi} (e^{\phi t} - 1) e^{\frac{\alpha(\beta+\phi)}{\sigma^2}t} \left(\frac{1-\kappa}{1-\kappa e^{\phi t}}\right)^{\frac{2\alpha}{\sigma^2}+1}, \\ H(t; \alpha, \beta, \sigma, c) &= e^{\frac{\alpha(\beta+\phi)+\phi\sigma^2}{\sigma^2}t} \left(\frac{1-\kappa}{1-\kappa e^{\phi t}}\right)^{\frac{2\alpha}{\sigma^2}+2}, \end{aligned}$$

and

$$\phi = \sqrt{2c\sigma^2 + \beta^2}, \kappa = \frac{\beta + \phi}{\beta - \phi}.$$

Note that $H(t) = -A(t)B'(t)$ and $G(t) = -A'(t)$.

When $\beta^2 + 2c\sigma^2 < 0$,

$$\begin{aligned}
A(t; \alpha, \beta, \sigma, c) &= e^{\frac{\alpha\beta}{\sigma^2}t} \left(\frac{\bar{\phi}}{|\varphi(t)|} \right)^{\frac{2\alpha}{\sigma^2}}, \quad B(t; \beta, \sigma, c) = -\frac{2 \tan\left(\frac{\bar{\phi}t}{2}\right)}{\bar{\phi} + \beta \tan\left(\frac{\bar{\phi}t}{2}\right)}, \\
G(t; \alpha, \beta, \sigma, c) &= -\frac{\alpha\bar{\phi}^{\frac{2\alpha}{\sigma^2}} e^{\frac{\alpha\beta}{\sigma^2}t}}{\sigma^2 |\varphi(t)|^{(1+\frac{2\alpha}{\sigma^2})}} \\
&\quad \left\{ \beta |\varphi(t)| + \bar{\phi} \text{signum}(\varphi(t)) \left[\bar{\phi} \sin\left(\frac{\bar{\phi}t}{2}\right) - \beta \cos\left(\frac{\bar{\phi}t}{2}\right) \right] \right\}, \\
H(t; \alpha, \beta, \sigma, c) &= e^{\frac{\alpha\beta}{\sigma^2}t} \left(\frac{\bar{\phi}}{|\varphi(t)|} \right)^{\frac{2\alpha}{\sigma^2}} \frac{\tan^2\left(\frac{\bar{\phi}t}{2}\right) + 1}{\left(1 + \frac{\beta \tan\left(\frac{\bar{\phi}t}{2}\right)}{\bar{\phi}}\right)^2},
\end{aligned}$$

where $\varphi(t) = \beta \sin\left(\frac{\bar{\phi}t}{2}\right) + \bar{\phi} \cos\left(\frac{\bar{\phi}t}{2}\right)$, $\bar{\phi} = \sqrt{-2c\sigma^2 - \beta^2}$, and

$$\text{signum}(z) = \begin{cases} \frac{z}{|z|} & \text{if } z \neq 0 \\ 0 & \text{if } z = 0 \end{cases}.$$

B Unscented Kalman Filter

The Unscented Kalman Filter (UKF) is a well-developed technique, widely applied in state estimation, neural networks, and nonlinear dynamic systems (see, e.g., Haykin et al., 2001 and Simon, 2006). Since the measurement equations in the state space formulation are nonlinear (for both CDS premiums and municipal bond prices) in this paper, the UKF is the natural choice for our estimation procedure. The state space (one-dimensional transition and m -dimensional measurements) is given by the following system:

Transition equation

$$X_t = \mathcal{T}X_{t-\Delta t} + \Theta + \sqrt{\mathbf{V}_t}e_t, \quad e_t \sim N(0, 1),$$

where X_t represents either $\lambda'_{m,t}$ or h_t depending on estimating the CDS model or the municipal bond model, \mathcal{T} and Θ are given by discretizing λ'_m or h_i , and \mathbf{V}_t is given by matching the variance of the innovation term in the discretized state dynamic.

Measurement equation

$$y_t = \mathcal{G}(X_t) + \zeta_t, \quad \zeta_t \sim N(\mathbf{0}_{m \times 1}, \mathcal{S}_{m \times m}),$$

where $\mathcal{G}(\cdot)$ is a nonlinear function of the state variable, giving the CDS premiums or the municipal bond prices, ζ_t is the pricing error, and \mathcal{S} is a diagonal covariance matrix with positive and distinct diagonal elements.

The essence of the UKF (Chow et al., 2007) used in this paper can be summarized briefly as follows. For each measurement time t , a set of deterministically selected points, termed *sigma points*, are used to approximate the distribution of the current state estimated at time t using a normal distribution with a mean vector $X_{t|t-\Delta t}$ and a covariance matrix that is a function of the state covariance matrix, $P_{X,t-\Delta t|t-\Delta t}$, and the conditional covariance \mathbf{V}_t .¹⁸ Sigma points are specifically selected to capture the dispersion around $X_{t|t-\Delta t}$, and are then projected using the measurement function $\mathcal{G}(\cdot)$, weighted, and then used to update the estimates in conjunction with the newly observed measurements at time t to obtain $X_{t|t}$ and $P_{X,t|t}$.

Next, we outline the detailed UKF algorithm:

1. Initialization¹⁹

$$X_{0|0} = \text{Constant},$$

$$P_{0|0} = 0.$$

¹⁸In the typical UKF setting, both transition and measurement equations are nonlinear. Hence, to compute the ex-ante predictions of the state variables' mean and variance, sigma points are needed to approximate the distribution of the previous state estimates. However, in our paper, the transition equations are linear. Therefore, we can directly compute the ex-ante predictions as in the standard Kalman Filter, without needing sigma points at this stage.

¹⁹Normally, the UKF is initialized at the unconditional mean and variance of the state variable. However, in the current case, the state variable is typically a non-stationary (explosive) process with no finite unconditional mean and variance. So in our case, we initialize the UKF at a constant parameter to be estimated with zero variance. This essentially assumes that by searching for the starting point of the UKF using the MLE criterion, we know exactly where the UKF starts.

2. Ex-ante predictions of the state

$$\begin{aligned} X_{t|t-\Delta t} &= \mathcal{T} X_{t-\Delta t|t-\Delta t} + \Theta, \\ P_{X,t|t-\Delta t} &= \mathcal{T} P_{X,t-\Delta t|t-\Delta t} \mathcal{T}' + \mathbf{V}_t. \end{aligned}$$

3. Selecting sigma points

Given the *ex-ante* prediction of the state variable $X_{t|t-\Delta t}$, three sigma points are selected as follows:

$$\chi_{t|t-\Delta t} = \begin{bmatrix} \chi_{0,t-\Delta t} & \chi_{+,t-\Delta t} & \chi_{-,t-\Delta t} \end{bmatrix},$$

where

$$\begin{aligned} \chi_{0,t-\Delta t} &= X_{t|t-\Delta t}, \\ \chi_{+,t-\Delta t} &= X_{t|t-\Delta t} + \sqrt{(1+\vartheta)} \left(\mathcal{T} \sqrt{P_{X,t-\Delta t|t-\Delta t}} + \sqrt{\mathbf{V}_t} \right), \\ \chi_{-,t-\Delta t} &= X_{t|t-\Delta t} - \sqrt{(1+\vartheta)} \left(\mathcal{T} \sqrt{P_{X,t-\Delta t|t-\Delta t}} + \sqrt{\mathbf{V}_t} \right). \end{aligned}$$

The term ϑ is a scaling constant given by

$$\vartheta = \eta^2 (1 + \varrho) - 1,$$

where η and ϱ are user-specified constants in this paper, with $\eta = 0.001$, and $\varrho = 2$. Since the values of these constants are not critical in our case, we omit a detailed description for the sake of saving space. Readers are referred to Chow et al. (2007) or Chapter 7 in Haykin et al. (2001) for further details.

4. Transformation of sigma points by way of the measurement function (predictions of measurements)

$\chi_{t|t-\Delta t}$ is propagated through the nonlinear measurement function $\mathcal{G}(\cdot)$:

$$\mathbf{Y}_{t|t-\Delta t} = \mathcal{G}(\chi_{t|t-\Delta t}),$$

where the dimension of $\mathbf{Y}_{t|t-\Delta t}$ is $m \times 3$. Then, define the set of weights for the

covariance estimates as:

$$W^{(c)} = \text{diag} \left[\frac{\vartheta}{1+\vartheta} + 1 - \eta^2 + 2, \frac{1}{2(1+\vartheta)}, \frac{1}{2(1+\vartheta)} \right]_{3 \times 3},$$

and obtain weights for the mean estimates as follows:

$$W^{(m)} = \begin{bmatrix} \frac{\vartheta}{1+\vartheta} \\ \frac{1}{2(1+\vartheta)} \\ \frac{1}{2(1+\vartheta)} \end{bmatrix}_{3 \times 1}.$$

The predicted measurements and associated variance and covariance matrices are computed as follows:

$$\begin{aligned} y_{t|t-\Delta t} &= \mathbf{Y}_{t|t-\Delta t} W^{(m)}, \\ P_{y_t|t-\Delta t} &= [\mathbf{Y}_{t|t-\Delta t} - \mathbf{1}_{1 \times 3} \otimes y_{t|t-\Delta t}] W^{(c)} [\mathbf{Y}_{t|t-\Delta t} - \mathbf{1}_{1 \times 3} \otimes y_{t|t-\Delta t}]', \\ P_{X_t, y_t} &= [\mathbf{X}_{t|t-\Delta t} - \mathbf{1}_{1 \times 3} \otimes X_{t|t-\Delta t}] W^{(c)} [\mathbf{Y}_{t|t-\Delta t} - \mathbf{1}_{1 \times 3} \otimes y_{t|t-\Delta t}]'. \end{aligned}$$

5. Kalman gain and ex-post filtering state update

With the output from Step 4, actual observations are finally brought in, and the discrepancy between the model's predictions and the actual observations is weighted by a Kalman gain Ξ_t to yield ex-post state and covariance estimates as follows:

$$\begin{aligned} \Xi_t &= P_{X_t, y_t} P_{y_t|t-\Delta t}^{-1}, \\ X_{t|t} &= X_{t|t-\Delta t} + \Xi_t (y_t - y_{t|t-\Delta t}), \\ P_{X_t, t|t} &= P_{X_t, t|t-\Delta t} - \Xi_t P_{y_t|t-\Delta t} \Xi_t', \\ y_{t|t} &= \mathcal{G}(X_{t|t}). \end{aligned}$$

References

- Ang, A., V. Bhansali, and Y. Xing (2014). The muni bond spread: Credit, liquidity, and tax. *Columbia Business School Research Paper* (14-37).
- Bergstresser, D., R. Cohen, and S. Shenai (2010). Financial guarantors and the 2007-2009 credit crisis. Working papers, Brandeis University.
- Bergstresser, D., R. Cohen, and S. Shenai (2015). Skin in the game: The performance of insured and uninsured municipal debt. Working papers, Brandeis University.
- Berndt, A., R. Douglas, D. Duffie, M. Ferguson, and D. Schranz (2008). Measuring default risk premia from default swap rates and EDFs. Working papers, Stanford University.
- Chalmers, J. M. (1998). Default risk cannot explain the muni puzzle: Evidence from municipal bonds that are secured by U.S. Treasury obligations. *Review of Financial Studies* 11(2), 281–308.
- Cheridito, P., D. Filipović, and R. L. Kimmel (2007). Market price of risk specifications for affine models: Theory and evidence. *Journal of Financial Economics* 83(1), 123–170.
- Chow, S., E. Ferrer, and J. Nesselroade (2007). An unscented Kalman filter approach to the estimation of nonlinear dynamical systems models. *Multivariate Behavioral Research* 42(2), 283–321.
- Chung, S.-L., C.-W. Kao, C. Wu, and C.-Y. Yeh (2015). Counterparty credit risk in the municipal bond market. *Journal of Fixed Income* 25(1), 7–33.
- Downing, C. and F. Zhang (2004). Trading activity and price volatility in the municipal bond market. *Journal of Finance* 59(2), 899–931.
- Driessen, J. (2005). Is default event risk priced in corporate bonds? *Review of Financial Studies* 18(1), 165–195.
- Duffee, G. R. (1999). Estimating the price of default risk. *Review of Financial Studies* 12(1), 197–226.

- Green, R. C. (1993). A simple model of the taxable and tax-exempt yield curves. *Review of Financial Studies* 6(2), 233–264.
- Green, R. C., B. Hollifield, and N. Schürhoff (2007). Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies* 20(2), 275–314.
- Harris, L. E. and M. S. Piwowar (2006). Secondary trading costs in the municipal bond market. *The Journal of Finance* 61(3), 1361–1397.
- Haykin, S. et al. (2001). *Kalman Filtering and Neural Networks*. Wiley Chichester.
- Jarrow, R. A. and S. M. Turnbull (1995). Pricing derivatives on financial securities subject to credit risk. *Journal of Finance* 50(1), 53–85.
- Lando, D. (1998). On cox processes and credit risky securities. *Review of Derivatives research* 2(2-3), 99–120.
- Lim, P. J. (2007). Don't let a bouncing bond market make you dizzy. *New York Times*, July 1.
- Longstaff, F. A. (2011). Municipal debt and marginal tax rates: Is there a tax premium in asset prices? *Journal of Finance* 66(3), 721–751.
- Longstaff, F. A., S. Mithal, and E. Neis (2005). Corporate yield spreads: Default risk or liquidity? New evidence from the credit default swap market. *Journal of Finance* 60(5), 2213–2253.
- Moldogaziev, T. T.-u. (2013). The collapse of the municipal bond insurance market: How did we get here and is there life for the monoline industry beyond the great recession? *Journal of Public Budgeting, Accounting and Financial Management* 25(1), 199.
- Ou, S. (2011). Corporate default and recovery rates, 1920–2010. *Special Comment, Moody's Investors Service*.
- Renick, O. and M. Bonello (2014). Bond insurance then and now: The revival of an industry. *Bond Buyer*, April 30.

- Simon, D. (2006). *Optimal State Estimation: Kalman, H_∞ , and Nonlinear Approaches*. Wiley-Interscience.
- Tudela, M., C. Jacoby, and A. Mediolì (2015). U.S. municipal bond defaults and recoveries, 1970-2014. *Special Comment, Moody's Investor Service*.
- Tudela, M., A. Mediolì, and A. V. Praagh (2012). U.S. municipal bond defaults and recoveries, 1970-2011. *Special Comment, Moody's Investor Service*.
- Wang, J., C. Wu, and F. X. Zhang (2008). Liquidity, default, taxes, and yields on municipal bonds. *Journal of Banking and Finance* 32(6), 1133–1149.
- Wells Fargo (2008). Deterioration of monoline insurance companies and the repercussions for municipal bonds. Technical report.

Table 1: Summary statistics of insured and uninsured bonds

This table presents the summary statistics of various characteristics of insured and uninsured bonds used in our regression analysis. Panel (a) shows the first quarter (Q1), median, third quarter (Q3), mean, and standard deviation (Stdev) of the yield, coupon, maturity, offering amount (in \$1,000), trade size (in \$1,000), and age (in years); Panel (b) shows the average yield and number of observations of insured bonds (insured by Ambac, FGIC, FSA, and MBIA) and uninsured bonds across different ratings: Aaa, Aa1, Aa2, Aa3, A1, A2, A3, Baa1, Baa2, Baa3.

(a) Bond characteristics

	Q1	Median	Q3	Mean	Stdev
Yield (%)	3.00	3.42	3.75	3.39	0.67
Coupon (%)	3.50	4.00	5.00	3.87	1.63
Maturity (years)	2.42	4.58	7.35	5.60	4.63
Offering amount (\$1,000)	921	2,620	10,300	13,700	93,200
Trade size (\$1,000)	20	50	100	224	1,111
Age (years)	0.12	2.91	5.13	3.63	4.18

(b) Average yields and observations (in parentheses) across ratings

	Ambac	FGIC	FSA	MBIA	Uninsured
Aaa	3.5 (29,067)		3.39 (91,132)	3.46 (64,857)	3.29 (71,252)
Aa1		3.81 (2,659)			3.22 (25,582)
Aa2		3.47 (5,070)			3.21 (24,176)
Aa3		3.52 (10,677)			3.34 (58,350)
A1		3.61 (7,275)			3.42 (23,240)
A2		3.53 (4,357)			3.13 (1,347)
A3		3.43 (3,668)			3.36 (768)
Baa1		3.91 (722)			3.44 (295)
Baa2		3.59 (708)			3.85 (186)
Baa3		3.73 (13,755)			4.03 (1,761)

Table 2: Panel regressions of municipal bond yields

This table presents the results of the panel regressions of municipal bond yields. The results for bonds with maturities shorter than five years (larger than five years) are shown in columns two to four (five to seven). The results of the regression with issuer CDS premiums are shown in the last column. Robust standard errors adjusted for clustering at the issuer level are in parentheses. *, **, and *** denote significance at the ten-percent, five-percent, and one-percent level, respectively.

	Maturity \leq 5 yrs			Maturity $>$ 5 yrs			All with CDS
	A	Aa	Aaa	A	Aa	Aaa	
Coupon	-0.012* (0.0068)	-0.032*** (0.0089)	-0.03*** (0.0039)	-0.035 (0.021)	-0.027** (0.011)	-0.047*** (0.0051)	-0.024** (0.011)
Maturity	0.12*** (0.011)	0.084*** (0.0087)	0.081*** (0.0023)	0.11*** (0.0086)	0.1*** (0.0048)	0.096*** (0.0027)	0.13*** (0.011)
Ln(amount)	-0.012*** (0.0041)	-0.0078 (0.0059)	0.0069* (0.0036)	-0.00063 (0.0065)	0.017*** (0.0049)	0.0056 (0.0064)	0.0021 (0.013)
Ln(tradesize)	-0.041*** (0.0051)	-0.044*** (0.0067)	-0.037*** (0.0023)	0.0046 (0.0036)	0.00094 (0.0017)	-0.0015 (0.002)	-0.033*** (0.0072)
Age	0.007*** (0.00095)	0.011*** (0.001)	0.013*** (0.0016)	-0.00046 (0.0011)	0.0062 (0.0045)	0.0096*** (0.0024)	0.0049 (0.0047)
Insured	-0.11*** (0.025)	-0.036** (0.017)	-0.043* (0.023)	-0.1*** (0.028)	-0.037 (0.027)	-0.037 (0.041)	-0.078*** (0.027)
Insurer 5yr CDS	0.000032*** (0.0000077)	0.000017 (0.000012)	0.000072*** (0.000013)	0.000074*** (0.0000097)	0.000068*** (0.0000075)	0.00012*** (0.000022)	0.000025 (0.000016)
Issuer 5yr CDS (bp)							0.01*** (0.0015)
Time FE	yes	yes	yes	yes	yes	yes	yes
Issuer FE	yes	yes	yes	yes	yes	yes	yes
Adj. R^2	0.61	0.62	0.61	0.87	0.84	0.86	0.73
Obs	21,930	73,147	129,718	16,512	45,623	110,406	12,692
Issuers	975	1,388	5,700	426	716	3,791	22

Table 3: Parameter estimates of the monoline default intensities

This table presents the parameter estimates of the default intensities, $\lambda_{m,t}$, of the four monolines: Ambac, FGIC, FSA, and MBIA. The standard errors are presented in parentheses. “ave VR” is the variance ratio averaged across CDS maturities, where the variance ratio is defined as one minus the ratio of the variance of pricing errors to the variance of actual CDS premiums. “rel RMSE” is the ratio of the root mean squared CDS pricing error to the sample average of actual CDS premiums. “rel RMSE” is in percentages.

	α	β	σ	α^P	β^P	c_0	c_1	ave VR	rel RMSE %
Ambac	0.002 (0.00)	0.010 (0.00)	0.348 (0.00)	0.006 (0.02)	-3.339 (0.84)	-0.001 (0.00)	1.036 (0.11)	1.00	6.05
FGIC	0.004 (0.00)	-0.386 (0.05)	0.915 (0.02)	0.634 (0.30)	0.604 (1.47)	-0.011 (0.01)	5.312 (1.54)	0.96	33.24
FSA	0.002 (0.00)	0.041 (0.01)	0.304 (0.01)	0.039 (0.03)	-1.171 (1.96)	0.000 (0.00)	0.029 (0.01)	0.99	8.85
MBIA	0.003 (0.00)	0.045 (0.00)	0.348 (0.00)	0.011 (0.02)	-4.391 (0.68)	0.002 (0.00)	1.006 (0.29)	0.99	11.71

Table 4: Summary of the parameter estimates of the municipal bond pricing model

This table summarizes the parameter estimates and fitting performance of the municipal bond pricing model for 61 issuers. “Q1” (“Q3”) is the 25th (75th) percentile of the sample. “(in)”, “(un)”, and “(all)” represent results of insured bonds, uninsured bonds, and all bonds, respectively. “VR” denotes the variance ratio, defined as one minus the ratio of the variance of pricing errors to the variance of actual bond prices. “rel RMSE” is the ratio of root mean squared error of all (cross section and over time) pricing errors to the sample average of all actual bond prices. “rel RMSE” is in percentages.

	Q1	Median	Q3	Mean
α_i	-0.001	-0.000	0.000	-0.001
β_i	-0.699	-0.371	-0.201	-0.538
σ_i	0.126	0.227	0.297	0.232
α_i^P	0.033	0.076	0.177	0.161
β_i^P	8.231	13.842	23.022	18.926
c_4	-0.002	-0.000	0.001	-0.009
c_5	-0.531	-0.107	0.048	0.146
c_2^{in}	0.011	0.014	0.020	0.019
c_3^{in}	-0.048	0.107	0.744	0.492
c_2^{un}	0.016	0.025	0.041	0.034
c_3^{un}	0.290	1.029	2.750	1.673
δ^{in}	0.000	0.000	0.067	0.142
δ^{un}	0.413	0.581	0.671	0.526
η	0.444	0.501	0.551	0.504
VR(all)	0.957	0.981	0.990	0.970
VR(in)	0.947	0.964	0.981	0.961
VR(un)	0.946	0.985	0.996	0.917
rel RMSE(all) %	0.478	0.610	0.777	0.660
rel RMSE(in) %	0.496	0.617	0.785	0.652
rel RMSE(un) %	0.284	0.439	0.735	0.610

Table 5: Summary of the average yield components of 61 issuers

This table summarizes the sample averages (from July 2007 to June 2008) of different yield components of 61 issuers' insured and uninsured bonds. "Q1" ("Q3") is the 25th (75th) percentile of the sample. This yield decomposition is done by applying each issuer's calibrated municipal bond pricing model to a hypothetical bond, which has a semi-annual coupon rate of 5.25% and a maturity of 4.2 years. For the insured bond, we consider the default component, liquidity component, and default+insurance component (both FSA and MBIA). For the uninsured bond, we consider the default component and liquidity component. All numbers in the table are in percentages.

Groups		Insured				Uninsured	
		Def	Liq	Def+Ins (FSA)	Def+Ins (MBIA)	Def	Liq
All (61 issuers)	Q1	0.373	1.141	0.036	0.071	0.194	0.958
	Median	0.653	1.392	0.057	0.143	0.329	1.240
	Q3	1.141	1.608	0.104	0.260	0.527	1.578
	Mean	0.778	1.375	0.071	0.172	0.413	1.289
States (17 issuers)	Q1	0.320	1.162	0.027	0.042	0.096	0.997
	Median	0.572	1.253	0.052	0.140	0.314	1.532
	Q3	0.848	1.538	0.082	0.205	0.550	1.747
	Mean	0.606	1.342	0.056	0.134	0.341	1.459
Cities (17 issuers)	Q1	0.539	0.988	0.044	0.086	0.337	0.769
	Median	0.858	1.299	0.079	0.193	0.452	1.123
	Q3	1.339	1.484	0.127	0.309	0.626	1.202
	Mean	0.965	1.255	0.088	0.210	0.560	1.003
Sch. Dist. (17 issuers)	Q1	0.345	1.064	0.027	0.085	0.173	0.998
	Median	0.553	1.530	0.046	0.123	0.321	1.420
	Q3	1.205	1.680	0.121	0.323	0.474	1.557
	Mean	0.813	1.485	0.077	0.195	0.419	1.344

Table 6: Correlation among the liquidity factor, the monoline default intensities, and the issuer default intensities

This table summarizes the correlations among the first order differences of l , λ'_m , and h_i . The 1st and 2nd PCs are the first and second principal components of the four λ'_m 's.

	1st PC	2nd PC	Ambac	FGIC	FSA	MBIA	l_t
$h_{i,t}$							
Q1	-0.098	-0.008	-0.044	-0.096	-0.082	-0.033	-0.059
Median	-0.024	0.052	0.013	-0.033	-0.004	0.028	0.010
Q3	0.049	0.126	0.083	0.030	0.071	0.115	0.069
Mean	-0.021	0.038	0.009	-0.027	-0.003	0.029	0.006
Stdev	0.106	0.118	0.133	0.102	0.124	0.122	0.118
l_t	0.167	-0.170	-0.100	0.181	-0.045	0.026	N.A.

Figure 1: Monoline CDS premiums

This figure shows the dynamics of five-year CDS premiums of the four monolines, Ambac, FGIC, FSA, and MBIA, from July 2007 to June 2008.




Figure 2: Relation between bond yield and bond age

This figure plots the average bond yield against bond age separately for short-term bonds (maturity shorter than 5 years), long-term bonds (maturity longer than 5 years), and all bonds.

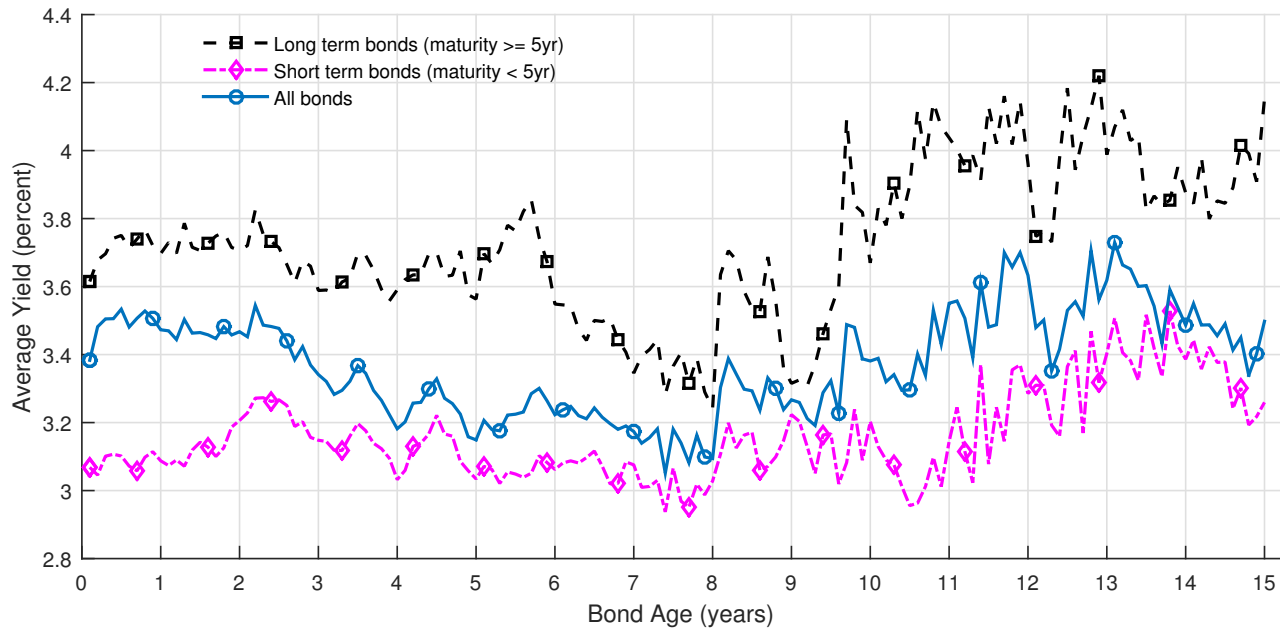


Figure 3: Liquidity spreads and the aggregate liquidity factor

This figure presents the dynamics of the short-term liquidity spread, long-term liquidity spread, and filtered liquidity factor l_t from July 2007 to June 2008.

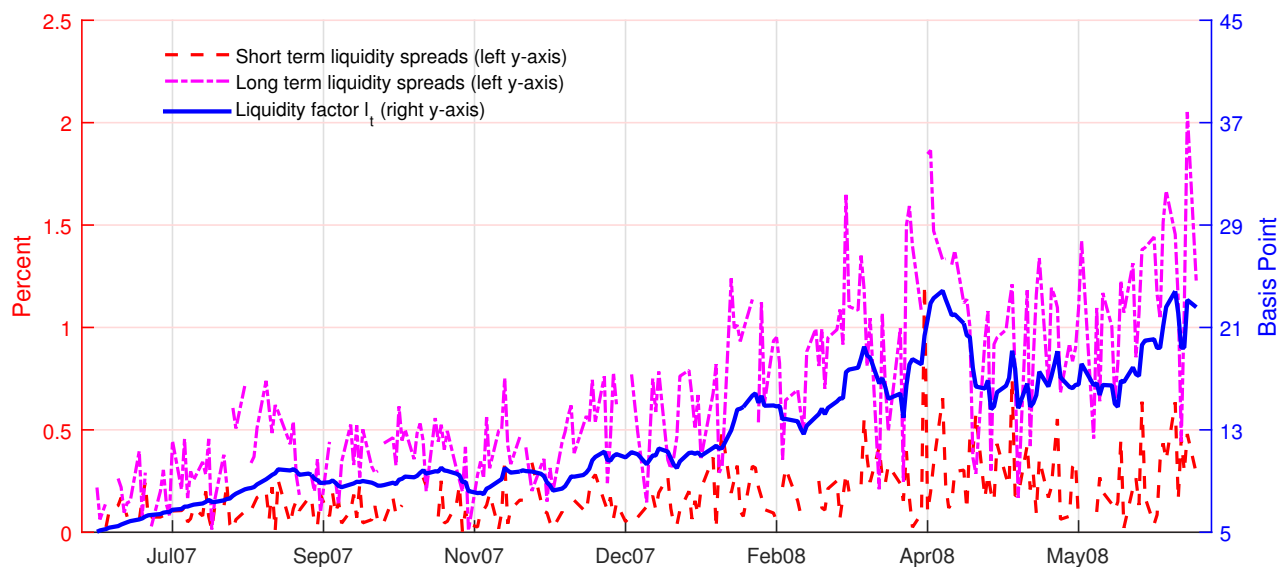


Figure 4: Monoline CDS pricing performance

This figure compares the model CDS premiums and the actual CDS premiums for the four monolines.

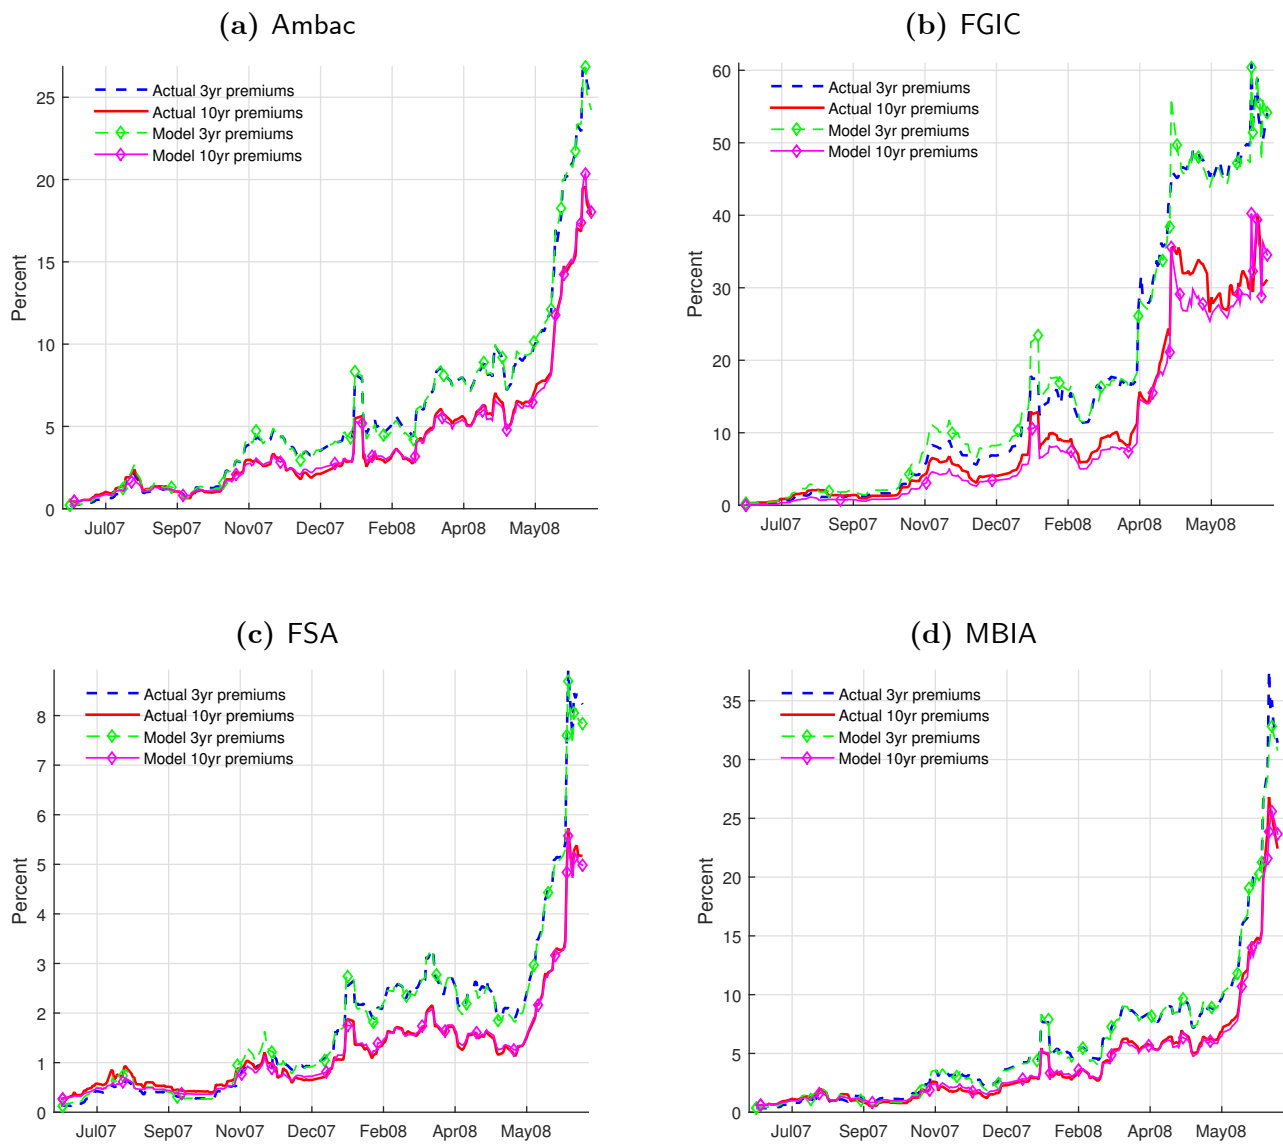




Figure 5: Yield decomposition of insured and uninsured bonds

This figure plots the median yield components across the 61 issuers from July 2007 to June 2008. This yield decomposition is done by applying each issuer's individual calibrated municipal bond pricing model to a hypothetical bond, which has a semi-annual coupon rate of 5.25% and a maturity of 4.2 years. For the insured bond, we consider the liquidity component and default+insurance component (both FSA and MBIA). For the uninsured bond, we consider the liquidity component and default component.


(a) Yield components of MBIA-insured and uninsured bonds




(b) Yield components of FSA-insured and uninsured bonds



(c) Yield comparison (MBIA-insured vs uninsured)



(d) Yield comparison (FSA-insured vs uninsured)



(e) Yield difference (uninsured minus MBIA-insured)



(f) Yield difference (uninsured minus FSA-insured)



Figure 6: Municipal bond-implied default intensity vs. CDS-implied default intensity

This figure compares the dynamics of the municipal bond-implied default intensity and the CDS-implied default intensity for three issuers, the State of California, the City of New York, and the State of Illinois, from July 2007 to June 2008.

