JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 47, No. 5, Oct. 2012, pp. 1125–1153

doi:10.1017/S0022109012000336

Loss Allocation in Securitization Transactions

Günter Franke, Markus Herrmann, and Thomas Weber*

Abstract

This paper analyzes the loss allocation to first, second, and third loss positions in European collateralized debt obligation transactions. The quality of the underlying asset pool plays a predominant role for the loss allocation. A lower asset pool quality induces the originator to take a *higher* first loss position, but, in a synthetic transaction, a *smaller* third loss position. The share of expected default losses, borne by the first loss position, is largely independent of asset pool quality but lower in securitizations of corporate loans than in those of corporate bonds. Originators with a good rating and low Tobin's Q prefer synthetic transactions.

I. Introduction

The global annual issuance volume of securitizations has grown from roughly 270 bn USD in 1997 to about 2100 bn USD in 2006 (Herrmann, Sun, Jha, Rudolph, Beckmann, and Bishko (2007)). The recent financial crisis depressed it. Securitizations, in particular those of mortgage-backed loans, are viewed as one driver of this crisis (The Financial Crisis Inquiry Commission Report (2011)). Yet, there are now many attempts to revive securitization because, in principle, it allows a better allocation of default risks across banks and nonbanks. A major impediment to the transfer of default risks is information asymmetries between the seller and the buyer of debt claims. In a securitization transaction, credit enhancements such as first loss positions (FLPs) serve to protect the buyers of rated tranches against adverse selection and moral hazard of the originator. Improved alignment of incentives of originators and buyers is a crucial ingredient for financial stability (International Monetary Fund (2011), ch. III). In a securitization

*Franke, guenter.franke@uni-konstanz.de, University of Konstanz, Department of Economics, POB 147, D-78457 Konstanz, Germany; Herrmann, markus.herrmann@lbbw.de, Landesbank Baden-Württemberg, Am Hauptbahnhof 2, D-70173 Stuttgart, Germany; and Weber, thomas.weber@axpo.ch, Axpo Holding AG, Parkstrasse 23, 5401 Baden, Switzerland. We are indebted to Michael Brennan, Wolfgang Bühler, Michel Habib, Dennis Hänsel, Julia Hein, Jens Jackwerth, Jan Krahnen, Peter de Marzo, Nandu Nayar, Richard Stapleton, Monika Trapp, Branko Urosevic, Christian Wilde, and the participants of the 2007 European Skinance Workshop, the 2007 European Financial Management Association Conference in Vienna, and the 2007 European Finance Association Meeting in Ljubljana for valuable comments. In particular, Hendrik Bessembinder (the editor), Francis Longstaff (associate editor and referee), an anonymous referee, and Kjell Nyborg suggested many very helpful improvements. The Appendix is available from the first author.

transaction this alignment is closely related to the allocation of default risks to FLPs, to rated securitization tranches (second loss positions (SLPs)), and, in case of synthetic transactions, to the nonsecuritized super-senior third loss positions (TLPs).

This paper investigates the determinants of FLPs in securitization transactions. Conventional wisdom suggests that the originator should retain the FLP because it is the most information-sensitive tranche and sell the information-insensitive senior tranches (Boot and Thakor (1993)). Yet, we observe many synthetic transactions in which the super-senior tranche is not securitized. We also address this puzzle to find out what drives the originator's choice between true sale and synthetic transactions. Understanding loss allocation is crucial for the risk management of originators of securitization transactions and for asset allocation of investors buying rated tranches, but also for regulators trying to make banks less vulnerable to default risks embedded in securitizations. Moreover, studying loss allocation provides new insights into the mechanisms used in modern structured finance to deal with market imperfections.

We study the loss allocation in a subset of European securitization transactions, called collateralized debt obligation (CDO) transactions. In collateralized loan obligation (CLO) transactions, also called balance sheet transactions, a bank securitizes part of its corporate loan portfolio. In collateralized bond obligation (CBO) transactions, also called arbitrage transactions, the originator of the transaction, a bank or an investment manager, buys corporate bonds, pools them in a portfolio, and securitizes it. Our data set does not include CDOs of mortgage-backed loans or bonds,[1] which were at the center of the recent financial crisis. Corporate loans and bonds underlying CLOs and CBOs turned out to be rather stable, as we argue later. Therefore, our findings should remain valid for the design of future securitization transactions.

Loss allocation in securitization transactions would be irrelevant in a perfect capital market. In imperfect markets, loss allocation is driven by balancing the interests of originators and investors, subject to various market imperfections such as regulation, information asymmetries, funding cost differentials, and transaction costs. We attempt to find out how these drivers and the quality of the securitized asset pool affect loss allocation. The analysis is refined by comparing different types of transactions, such as true sale versus synthetic transactions and CLO versus CBO transactions.

First, we restate some basics about securitizations. The 3 players governing securitization are banks and investment managers as originators of these transactions, investors buying securitization tranches, and the rating agencies as information intermediaries. Given the models of the rating agencies, originators and investors determine the loss allocation in securitization transactions. The originator selects a set of loans and/or bonds[2] as the underlying asset pool of the transaction. In all transactions, loss allocation is governed by strict subordination. Default losses are solely borne by the FLP, also called equity tranche, until this tranche is

[1]This does not rule out that a corporate loan or bond is also partly collateralized by a mortgage or other types of collateral.

[2]The bonds may include a few tranches of other securitizations or structured finance products.

completely absorbed by these losses. Then, the next losses are solely borne by the tranche with the lowest rating until it is fully absorbed, and so on. Strict subordination generates very strong quality differences between the tranches. In a *true sale transaction*, the originator sells all loans/bonds without recourse to the special purpose vehicle, which issues the equity tranche and various rated bond tranches, usually including a large Aaa tranche. The par values of all tranches add up to the par value of all loans/bonds securitized or slightly less. Therefore, the transaction is fully funded. The originator can freely use the proceeds from tranche issuance. In a *synthetic transaction* the originator retains ownership of the loans/bonds and transfers part of the default risk through a credit default swap to the special purpose vehicle. This swap covers default risks beyond a threshold that defines the FLP; the coverage is limited by the par value of the issued tranches, which is usually much smaller than the par value of the underlying asset pool (partially funded transaction). Investors buying these tranches take an SLP. Default losses beyond the FLP and SLP are borne by the nonsecuritized super-senior tranche, a TLP held by the originator. She may insure its risk by buying a senior default swap. In a synthetic transaction the originator obtains no funding. The issuance proceeds need to be invested in Aaa securities or other almost default-free assets in order to provide strong collateral for the investors taking the SLP.

In a CLO transaction the originator usually owns the loans to be securitized, while in a CBO transaction she buys bonds, pools them, and securitizes the asset pool. In a CLO transaction the originator acts as the servicer of the loans so that her monitoring and collection policies affect loan defaults. In a CBO transaction the originator lacks this role so that her impact on defaults is very limited. The impact of the originator on the evolution of default losses of the asset pool also depends on whether the transaction is static or dynamic. In a static deal, the asset pool is determined from the outset; new assets cannot be added later on. In a dynamic (managed) deal, the originator may change the asset pool over time, subject to various restrictions in the securitization contract to preserve the quality of the asset pool.

To the best of our knowledge, this paper is the first to empirically study the impact of the quality of the asset pool and originator characteristics on the loss allocation in securitization transactions. To measure the quality of the asset pool, we rely on the most important quality characteristics published by Moody's, the assets' weighted average default probability (WADP) and Moody's diversity score (DS) as well as the more sophisticated adjusted diversity score (ADS). A lower WADP and/or a higher DS indicate a better asset pool quality. Relying on Moody's assessment may be viewed as problematic. Rating agencies, being important players in the securitization business, have come under strong attack in the subprime crisis. Their initial ratings of bonds issued in securitizations of subprime loans strongly underestimated the default risk. A huge wave of downgrades of these bonds followed in 2007. However, ratings of corporate bonds and of CDOs like those used in our sample escaped severe criticism. These ratings were quite stable until the end of 2007.[3] This is also supported by the evidence in

[3]Worldwide, the percentage of Moody's downgrades over the previous 12 months did not change significantly from the first to the last quarter of 2007, being around 9.5%, compared to an average of

Newman, Fabozzi, Lucas, and Goodman (2008). Clearly, Moody's quality assessments are subject to estimation error. We assume that they are unbiased indicators of default risk.

The new findings of the paper can be summarized as follows:

First, we show that ratings based on the probability of default (PD) imply the observed strict subordination of securitization tranches.

Second, the empirical evidence strongly confirms an *inverse* relation between the quality of the securitized asset pool and the size of the FLP, but a *positive* relation between asset pool quality and the size of the TLP in synthetic transactions. Since information asymmetry is likely to be stronger for asset pools of lower quality, the inverse relation between the size of the FLP and asset pool quality suggests that the FLP grows with the perceived extent of information asymmetry.[4] The positive relation between the size of the TLP and asset pool quality suggests that loss allocation to the TLP is driven primarily by the originator's risk and funding strategy.

Third, we ask whether there exists a measure of investor protection against default losses that is invariant to asset pool quality. Such a measure of loss allocation might balance the interests of the originator and investors. Since the size of the FLP is a crude measure of protection, we investigate 2 other measures of investor protection, assuming a lognormal distribution for the default loss rate of the underlying asset pool. The 1st measure is the share of expected default losses absorbed by the FLP, called the *loss share*. The 2nd measure is the probability that all default losses are exclusively borne by the FLP (i.e., investors are not hit by default losses). We denote it as the *support probability* of the FLP.

Empirically, it turns out that the *loss share* is largely independent of the asset pool quality, in contrast to the support probability. The loss share is, on average, about 84% for CLO transactions and about 90% for CBO transactions. In CLO transactions there is much room for the originator's moral hazard; in CBO transactions there is not. Hence, one would expect a higher loss share in CLO transactions, in contrast to the empirical evidence. This suggests that moral hazard may lead to high reputation costs for the originator, implying a strong monitoring effort. Therefore, investors may accept a smaller loss share in CLO transactions.

Fourth, the attractiveness of a synthetic relative to a true sale transaction increases with asset pool quality. Better quality implies a lower default risk of the super-senior tranche, making it less attractive for the originator to sell this tranche. A TLP is in stark contrast to the literature, which argues that the originator should sell the least information-sensitive tranche. The preference for synthetic transactions appears to be stronger for originators with a better rating. This may also be driven by funding costs. Highly rated originators can obtain cheaper funding issuing standard bonds rather than highly rated bonds in true sale securitizations.

12.6% over the period 1985 through the 3rd quarter of 2008 (Moody's (2008)). In the United States and Canada, there was even a slight decline in downgrades, while in Europe there was some increase. Not surprisingly, corporate downgrades increased in 2008. Similar results are obtained for Standard & Poor's (S&P) as shown by Bloomberg and Fitch (Fitch (2008)).

[4]Originators rarely announce to what extent they retain the FLP. The current financial market reforms require the originator to retain at least 5% of all default losses to mitigate adverse selection and moral hazard.

Fifth, the last important surprising finding is that characteristics of the originator such as her total capital ratio or Tobin's Q, which may proxy for her securitization motives, add little to the explanatory power of the regressions.

The paper is structured as follows: Section II reviews the literature. In Section III we model the originator's choice problem and derive hypotheses. The empirical findings are presented and discussed in Section IV. Section V concludes.

II. Literature Review

Several papers analyze the optimality of FLPs. In the absence of information asymmetries, Arrow (1971) (see also Gollier and Schlesinger (1996)) analyzes the optimal insurance contract for a setting in which the protection buyer is risk averse but the protection sellers are risk neutral. If they bound their expected loss from above, then an FLP of the protection buyer is optimal. Townsend (1979) considers risk sharing between a risk-averse entrepreneur and investors in the presence of information asymmetries about the entrepreneur's ability to pay. The optimal contract is a standard debt contract implying an FLP of the entrepreneur (see also Gale and Hellwig (1985)). Malamud, Rui, and Whinston (2009) study optimal tranching in a securitization transaction, given a perfect market where agents have homogeneous expectations but heterogeneous preferences. In their model, optimal risk sharing is achieved by issuing multiple tranches and investors buying different portfolios of tranches.

The literature on security design distinguishes between information-sensitive and -insensitive securities. Boot and Thakor (1993) argue that a risky cash flow should be split into a senior and a subordinated security. The senior information-insensitive security can be sold to uninformed investors, while the information-sensitive subordinated security should be sold to informed investors. This should raise the sales revenue. Riddiough (1997) extends this reasoning by showing that loan bundling allows for asset pool diversification, which softens information asymmetries.[5] DeMarzo and Duffie (1999) analyze the security design assuming a trade-off between the retention cost of holding cash flows and the liquidity cost of selling information-sensitive securities. They prove that a standard debt contract is optimal. DeMarzo (2005) shows that pooling of assets has an information-destruction effect, since it prohibits the seller from selling asset cash flows separately. But pooling improves diversification. Tranching allows selling more liquid, information-insensitive claims. Summarizing, these papers demonstrate the optimality of an FLP and argue that the senior information-insensitive tranches should be sold to investors. This is in strong contrast to synthetic transactions where the large least information-sensitive tranche, the TLP, is not sold.

Duffie and Gârleanu (2001) discuss adverse selection and moral hazard in securitizations. They also discuss Moody's DS and illustrate the sensitivity of the portfolio loss rate distribution to various parameters.[6] Duffie, Eckner, Horel,

[5]Gorton and Pennacchi (1995) consider a bank that optimizes the fraction of a single loan to be sold and the guarantee against loan default through a repurchase agreement.

[6]Plantin (2003) shows that sophisticated institutions with high distribution costs buy and sell the junior tranches, leaving senior tranches to retail institutions with low distribution costs.

and Saita (2009) argue that unobservable, nonstationary risk factors raise default correlations and imply higher tail risks for loan portfolios. This could explain the rather high credit spreads on Aaa tranches in securitizations. Albrecher, Ladoucette, and Schoutens (2007) propose a generic 1-factor Lévy model to derive the portfolio loss rate distribution. Burtschell, Gregory, and Laurent (2009) compare models with one latent factor and different copulas to derive default intensities for CDOs. In an empirical study, Krekel (2008) proposes a Gaussian base correlation model with correlated recovery rates to improve the empirical model fit. Longstaff and Rajan (2008) estimate loss distributions of Credit Default Index tranches with multifactor models and find that a 3-factor model leads to a 3-modal loss rate distribution, where the 2nd (3rd) mode has a much smaller density than the 1st (2nd). They interpret the 2nd and 3rd factors as default clustering factors.

Among the empirical studies of securitizations, Downing and Wallace (2005) find in commercial mortgage-backed securities transactions that FLPs are higher than what might be expected looking at the actual performance of mortgages. Downing, Jaffee, and Wallace (2009) observe that securitized assets have lower quality than nonsecuritized. Consistent with this, Piskorski, Seru, and Vig (2010) find a lower foreclosure rate associated with bank-held loans compared to similar securitized loans. Purnanandam (2011) observes significantly higher mortgage-related write-offs for banks that are more engaged in true sale securitizations, consistent with adverse selection and moral hazard. Loutskina and Strahan (2009) argue that securitization raises banks' willingness to approve mortgages that are hard to sell. Cebenoyan and Strahan (2004) document that banks securitizing loans hold less capital than other banks and have more risky assets relative to total assets. Franke and Krahnen (2006) find that securitization tends to raise the bank's stock market beta, indicating more systematic risk.

III. Hypotheses

A. Strict Subordination

In this section we present the hypotheses to be tested. Since strict subordination is perhaps the most striking property of loss allocation in securitization transactions, we first motivate it by the strong role played by the rating agencies. S&P and Fitch rate according to the PD, while Moody's rates according to the expected loss rate. The empirical evidence shows that Moody's ratings are similar to PD ratings. The following lemma shows that under weak conditions, PD rating leads to strict subordination.

Lemma 1. Assume that i) the owner of a tranche enjoys limited liability, ii) the issuer of securitization tranches minimizes the credit spreads to be paid, iii) the credit spread paid on a tranche is inversely related to its rating, and iv) inverse loss sharing is ruled out. Then PD rating implies strict subordination of tranches.

The first 3 assumptions of the lemma need no explanation. Inverse loss sharing is defined by a loss sharing such that at least one agent bears *less* default losses when the default loss of the underlying asset pool *increases*. This agent would benefit from higher losses of the asset pool and, thus, have an incentive to raise losses. So inverse loss sharing is ruled out. Then PD rating implies strict subordination. The proof is straightforward. Consider a transaction with 2 tranches only. If both tranches share the 1st unit of default losses as well as additional units, then both tranches have the same PD. To minimize the PD of the 2nd tranche, holding the PD of the 1st tranche constant, the latter has to exclusively bear all losses until it is exhausted. This maximizes the rating of the 2nd tranche and, hence, minimizes the credit spread of this tranche. The same argument applies to multiple rated tranches, proving the lemma.

B. Information Asymmetry and Asset Pool Quality

The PD-based rating implies a minimum size of the FLP under the usual condition that the lowest rating of a rated tranche is B or better. While ratings provide important information to investors, they cannot remove information asymmetries completely. Originators usually have better information on loans and bonds underlying a transaction than investors. This creates room for adverse selection and moral hazard of the originator, so that investors presumably demand higher credit spreads for rated securitization tranches. The originator may try to reduce this penalty by credit enhancements, in particular by an FLP. The stronger the information asymmetry, the higher the FLP should be.

To model information asymmetries, we distinguish between the published and the true quality of the underlying asset pool. Asset pool quality is measured by the WADP of the loans/bonds in the asset pool and by asset pool diversification. Diversification is summarized by Moody's DS or the improved ADS. This score can be interpreted as the diversification-equivalent number of equal-sized loans whose defaults are uncorrelated. A 3rd characteristic of the asset pool quality is the weighted average expected loss given default. To simplify modeling we assume that the loss given default λ is nonrandom. Then the expected loss rate of the asset pool is λ WADP. Given λ, we characterize asset pool quality by WADP and DS or ADS.

We assume that rating agencies publish unbiased information about the underlying asset pool quality.[7] The true quality differs from the published quality by a noise term ε,

$$\text{published asset pool quality} \quad = \quad \text{true asset pool quality} + \varepsilon.$$

The standard deviation of the noise term, $\sigma(\varepsilon)$, is a measure of quality uncertainty. It should be inversely related to the true asset pool quality. The intuition for this is that errors in estimating WADP are likely to be *proportional* to the true WADP. If the true WADP is very small (high), then errors in estimating WADP are likely to be small (high). Also, $\sigma(\varepsilon)$ should be inversely related to the true DS. As pointed out by DeMarzo (2005) and others, a high DS reduces information asymmetries because the idiosyncratic risks of the assets tend to be

[7] As argued in the Introduction, the criticism of rating agencies concerns their rating of securitizations of mortgage-backed loans, not that of corporate loans or bonds. In any case, our sample ends in 2005. At that time, confidence in ratings was still very strong.

diversified away. *Hence, asset pool quality and quality uncertainty should be inversely related.*

Quality uncertainty creates room for adverse selection and moral hazard and therefore should be priced. One might argue that investors can buy tranches of many different transactions and thereby diversify their ε-risk. But buying a tranche incurs management and other transaction costs constraining diversification.[8] Moreover, there is a nondiversifiable risk that the rating models used by the rating agencies are flawed. Given only 3 relevant agencies that appear to use similar models, this creates a systematic risk. Therefore the ε-risk is likely to be priced.

C. Allocation of Losses to the FLP

As discussed before, the FLP should be higher, the stronger the information asymmetry. This and the inverse relationship between information asymmetry and asset pool quality motivate

Hypothesis 1. The lower the quality of the asset pool, the higher is the FLP.

This hypothesis can also be motivated by the rating methodology. Given the rating of the lowest-rated tranche, a lower asset pool quality requires a higher FLP. The size of the FLP[9] is a crude measure of loss allocation to the FLP because it does not take into account the loss distribution of the asset pool. In equilibrium, the loss exposure of the FLP *relative* to that of the rated tranches should be balanced. An intuitive measure for relative loss exposure is the *loss share*. It is defined as the expected loss borne by the FLP, divided by the expected loss of the asset pool. Investors might view the loss share as an important signal of balancing interests. Alternatively, consider the *support probability* of the FLP, the probability that the asset pool losses are smaller than the FLP, and *(1-support probability)*, the probability that rated tranches are hit by default losses. According to S&P and Fitch, this probability determines the rating of the tranche, which is subordinate to all other rated tranches. It relates the loss allocation to quantile considerations, as does the value at risk, which is commonly used to assess tail risk.

There is no equilibrium model that relates the loss share or the support probability to the asset pool quality. Investors may pay more attention to their expected losses than to the support probability. Hence, we conjecture that there might be an equilibrium loss share that is largely independent of asset pool quality. This motivates

Hypothesis 2. The loss share of the FLP is independent of the asset pool quality.

Before testing this hypothesis, we present some theoretical properties. A decline in DS or ADS, holding WADP constant, is modeled as a mean-preserving

[8] Gabaix, Krishnamurthy, and Vigneron (2007) find that prepayment risk is priced in mortgage-backed securities transactions, even though it is a pure redistribution risk. They attribute this to limits of arbitrage.

[9] The FLP can take different forms. In a true sale transaction, the FLP is the most junior tranche. It may be supplemented by a reserve account in which interest surplus (interest revenue from the asset pool minus interest expense on tranches) accrues over time. Default losses are then absorbed first by the reserve account.

spread in the loss rate distribution of the asset pool. An increase in WADP, holding diversification constant, is modeled as a 1st-order stochastic dominance shift in the loss rate distribution. Without loss of generality, we assume that the asset pool has a par value of 1€. Then the default loss of the pool equals the portfolio loss rate. Lemma 2 presents comparative statics of the loss allocation measures with respect to asset pool quality, *given* the size of the FLP.

Lemma 2. Consider a securitization transaction, given the size of the FLP.

a) A decline in asset pool diversification

   1) implies a lower expected loss for the FLP and a higher expected loss for the rated tranches (including the TLP in case of a synthetic transaction); hence, it reduces the share of expected losses of the asset pool borne by the FLP;

   2) reduces the support probability of the FLP if the FLP is higher than the loss rate at which the 2 cumulative probability distributions intersect. This condition holds for a lognormal loss rate distribution and a nonrandom loss given default λ for each loan/bond if and only if

$$\text{FLP} \geq \lambda\,\text{WADP}\sqrt{1+\frac{\frac{1}{\text{WADP}}-1}{\text{DS}}}. \tag{1}$$

b) An increase in the WADP of the asset pool

   1) raises the expected loss of both, the rated tranches (including the TLP in the case of a synthetic transaction) and the FLP,

   2) reduces the support probability of the FLP,

   3) reduces the share in expected losses of the asset pool borne by the FLP, given a lognormal loss rate distribution, if

$$1-N(h+\sigma)-n(h+\sigma)\,\frac{1}{2\sigma}\,\frac{1}{1+\text{WADP}(\text{DS}-1)} \geq 0, \tag{2}$$

with $h = \frac{\ln\frac{\text{FLP}}{\lambda\,\text{WADP}}}{\sigma} - \frac{\sigma}{2}$ and $\sigma^2 = \ln\left(1+\frac{\frac{1}{\text{WADP}}-1}{\text{DS}}\right)$.

$N(\cdot)$ and $n(\cdot)$ denote the standard normal distribution function and the standard normal probability density function, respectively.

Lemma 2a and 2b are well known except for the results regarding a lognormal loss rate distribution, which are proved in Appendix 1.2 and 1.3.[10] By Lemma 2a, a decline in asset pool diversification redistributes default losses from the FLP to the rated tranches. Hence, the loss share declines and the originator would benefit more from adverse selection and moral hazard. By Lemma 2b, an increase in WADP raises the expected losses of the FLP and of the rated tranches, and it reduces the support probability. This, again, might alert investors.

[10]The Appendix is available from the 1st author.

While these results are true for all probability distributions, the impact of DS on the support probability and of WADP on the loss share depend on the probability distribution. While Longstaff and Rajan (2008) derive a 3-modal loss rate distribution, we use a unimodal distribution because the models used by the rating agencies to simulate the default losses of the asset pool generate unimodal loss rate distributions. A distribution that approximates the probability distribution derived from simulation models reasonably well is the lognormal distribution. This was also used by Moody's Investor Service (2000).[11] For each securitization transaction we translate the expected value and the standard deviation of the loss rate of the asset pool into the 2 moments of a lognormal distribution, as shown in Appendix 1.1. For each claim in the asset pool we assume the same default probability WADP and the same loss given default λ. Using the Black-Scholes (1973) model, the loss share s is given by $s = N(h) +$ (FLP/λ WADP) $(1 - N(h+\sigma))$, and the support probability is $\gamma(\text{FLP}) = N(h+\sigma)$.

Lemma 2a states the condition for a decline of the DS to lower the support probability, and Lemma 2b states the condition for an increase in WADP to lower the loss share. These conditions are mostly satisfied in our sample. Therefore, we start from the premise that a decline in asset pool quality reduces both the loss share and the support probability, given the size of the FLP. Avoiding both effects requires an increase of the FLP as stated in Hypothesis 1. If the originator adjusts the FLP to asset pool quality so as to keep the loss share constant (Hypothesis 2), then the support probability will change as stated in Lemma 3, which is proved in Appendix 2.

Lemma 3. Assume a lognormal loss rate distribution. Suppose that the loss share of the FLP is independent of the asset pool quality. Then the support probability of the FLP is inversely related to the WADP and to the DS, if and only if $h < n(h+\sigma)/(1 - N(h+\sigma))$. Also, $\partial \ln\text{FLP}/\partial \ln\text{WADP} < 1$.

Lemma 3 states a surprising testable result. Given a constant loss share and the condition on h, the support probability of the FLP declines if one measure of portfolio quality, the WADP, worsens, but it also declines if the other measure of portfolio quality, the DS, *improves*. This shows that given a *higher* WADP, the depressed support probability increases relatively more slowly than the depressed loss share with the FLP, but given a *lower* DS, the depressed support probability increases relatively faster than the depressed loss share with the FLP.

D. True Sale versus Synthetic Transactions

Our conjecture that information asymmetry plays a major role for the allocation of losses to the FLP can be checked by comparing true sale and synthetic transactions. In synthetic transactions, by strict subordination, the TLP does not

[11]The lognormal distribution implies a positive probability of a loss rate above 1. But this probability is very small even for low-quality asset pools. Consider a transaction in which the WADP of loans is very high with 20% and the DS is very low with 10. Then the probability of the implied lognormal distribution for loss rates above 1 is 0.1%. In typical transactions, this probability would be much smaller.

serve as a credit enhancement for the rated tranches (SLP). Therefore, information asymmetry should play a minor role for the allocation of losses to the TLP. The choice between true sale and synthetic transactions is, however, driven also by other considerations. For example, borrowers may not agree to a transfer of the loan to another creditor. A true sale transaction would then be infeasible. A synthetic transaction also avoids the operational risk associated with a property transfer in a true sale transaction. Moreover, a synthetic transaction permits the originator to short-sell default losses.

The originator may choose a synthetic transaction if she prefers to retain the TLP risk and avoid the transaction cost of securitizing it. This preference should be stronger for a stronger asset pool quality because then the TLP risk is smaller. Therefore we state

Hypothesis 3. The preference for synthetic over true sale transactions increases with asset pool quality.

A better asset pool quality reduces the attachment point of an Aaa tranche. This is the contractually defined asset pool loss rate such that this tranche bears losses only when the loss rate exceeds this rate. Hence a better asset pool quality should raise the TLP (instead of lowering it as suggested for the FLP by Hypothesis 1). This motivates

Hypothesis 4. In a synthetic transaction the nonsecuritized super-senior tranche (TLP) increases with the quality of the asset pool.

Retaining the information-insensitive super-senior tranche is in stark contrast to the literature. One explanation of this puzzle may be the funding cost. In a true sale transaction the originator may freely use the proceeds from issuing tranches, while synthetic transactions provide no funding. Hence, a bank's choice between true sale and synthetic transactions may also depend on the funding costs in a true sale transaction versus those of standard bank bonds. We hypothesize that banks with a very good rating have little incentive to use CDO transactions for funding purposes, since they can obtain funds at low credit spreads anyway. This motivates

Hypothesis 5. Synthetic [true sale] transactions are preferably used by banks with a strong [weak] rating.

Interestingly, Hypotheses 3–5 can also be derived from optimal risk sharing in a model with heterogeneous investor preferences (Malamud et al. (2009)).

E. CLO versus CBO Transactions

Next, compare CLO and CBO transactions. In a CLO transaction, the originator should monitor the debtors so that a moral hazard problem exists. In a CBO transaction, the originator is, as any other bond investor, in a remote position vis-à-vis the bond obligors so that she cannot effectively monitor them. Therefore, credit spreads of CLOs may include a higher penalty for moral hazard. Also, the potential for adverse selection may be stronger in CLO transactions because more public information exists on bonds than on loans. These conjectures about adverse selection and moral hazard are confirmed by some of the empirical studies on mortgage-backed securities quoted previously. Higher credit spreads in CLO transactions should motivate the originator to mitigate this penalty through a higher FLP. This leads to

Hypothesis 6. Given the same quality of the asset pool, the loss share of the FLP is higher in CLO than in CBO transactions.

This hypothesis ignores reputation costs of the originator. Investors may react to default losses they have to bear by imposing higher credit spreads on the same originator in future securitizations. This generates a reputation cost. A high reputation cost may induce a strong originator effort that partially substitutes for the FLP. Hence, the FLP may be smaller in a CLO than in a CBO transaction. Higgins and Mason (2004) document various cases in which banks voluntarily absorbed default losses in credit card securitizations beyond the FLP so as to reduce their reputation cost. This cost might invalidate Hypothesis 6. Second, this hypothesis ignores the observation that default rates tend to be higher for bonds than for loans of the same debtor (Emery and Cantor (2005)). Therefore, expected losses could be higher for securitized bonds than for loans. This difference would not show up in debtor ratings as opposed to debt claim ratings. Hence, investors might insist on higher loss shares in CBO transactions.

IV. Empirical Findings

These hypotheses will be tested on a set of 169 European CDO transactions of corporate loans and bonds, excluding other types of collateral such as mortgages. Our data set includes all European CDO transactions from the end of 1997 to the end of 2005 for which we know Moody's DS and can derive WADP.[12] Information about transactions is taken from offering circulars, from presale reports issued by Moody's, and from the Deutsche Bank's European Securitization Almanac. The sample represents about ½ of all European CDO transactions issued in the observation period.

A. Derivation of Asset Pool Quality

Since asset pool quality is essential for our analysis, we provide information on how WADP and DS are derived. Rating agencies assign a rating to each asset in the pool. For publicly rated assets, such as corporate bonds, current bond ratings are used if derived in-house. The rating of another agency is lowered by a defined number of rating steps (notches), mainly to account for the uncertainty about the underlying rating model. For asset pools without publicly rated assets, such as corporate loans, 3 rating approaches are typically employed.[13] Moody's publishes a weighted average rating for the asset pool, based on the individual

[12]We include a few transactions without a rating from Moody's where the average quality of the underlying assets is known and also their diversification.

[13]The 3 rating approaches are: i) The 1st approach relies on the quality of the originator's internal rating system. It maps the originator's rating scale to the agency's rating scale. Mappings are monitored regularly and adjusted if necessary. ii) The 2nd approach uses historical default loss data, often

assets' ratings. We use Moody's tables to translate the weighted average rating of an asset pool into the WADP.

The expected default loss rate of the asset pool equals λ WADP. Since we mostly do not have transaction-specific information on loss given default, we assume λ to be 50%, with few exceptions. This is in line with Acharya, Bharath, and Srinivasan (2007), who document recovery rates for various loans and bonds in the United States. They find an average recovery rate slightly above 50%. For 2 transactions with secured loans, we use $\lambda = 25\%$. For mezzanine transactions with subordinated and unsecured underlyings, we use $\lambda = 100\%$ as the rating agencies do.

The loss allocation in a transaction also depends on the diversification of the asset pool. Moody's DS measures the diversification of the assets within and across industries, taking into account also variations in asset size. DS is defined as the number of claims of equal size and uncorrelated defaults, which gives the same standard deviation of the asset pool loss rate distribution as that actually observed. DS is defined by Moody's as

$$\mathrm{DS} = \sum_{k=1}^{m} G\left\{\sum_{i=1}^{n_k} \min\left\{1, F_i/\bar{F}\right\}\right\}.$$

Here, m denotes the number of industries, n_k the number of claims against obligors in industry k, F_i the par value of claim i, $\bar{F}$ the average par value of all claims, and $G(y)$ is an increasing concave function starting at $G(1) = 1$ with a maximum of 5 attained at $y = 20$. Hence, the maximum DS within an industry is 5. The DS ranges between 1 and 135. DS $= 1$ indicates "no diversification," and DS $= 135$ indicates "excellent diversification."

The DS has been criticized on various grounds (Fender and Kiff (2004)). Therefore, in 2000, Moody's started to use an *adjusted* DS (ADS). The ADS explicitly takes into account asset correlations of obligors within an industry, ρ_{in}, and between industries, ρ_{ex}. Given equal-sized loans, Fender and Kiff show that

$$\mathrm{ADS} = \frac{n^2}{n + \rho_{\text{ex}}\, n\,(n-1) + (\rho_{\text{in}} - \rho_{\text{ex}}) \sum_{k=1}^{m} n_k\,(n_k - 1)}.$$

For example, consider a transaction with 15 industries and 10 loans of equal size in each industry, assuming $\rho_{\text{in}} = 20\%$. Then the ADS with $\rho_{\text{ex}} = 0\%$ is about 39. But with $\rho_{\text{ex}} = 2\%$ and 4% it would be about 23 and 16, respectively, indicating the sensitivity of the ADS to ρ_{ex}. Given 10 industries and 6 loans of equal size in each industry, the DS is similar to the ADS if $\rho_{\text{in}} = 20\%$ and $\rho_{\text{ex}} = 0\%$ (Fender and Kiff (2004)).

Deriving ADS with 2 correlation coefficients only is clearly a simplification. This shrinkage approach may have some merits in view of the difficulties of

called the statistical/actuarial approach, in line with the agency's published recovery assumptions. iii) The 3rd approach uses proprietary credit risk information systems (i.e., a database/scoring model of corporate accounting data) in order to calculate an expected default rate for each obligor, using recovery assumptions. Examples include Moody's KMV RiscCalc™ and S&P's Credit Risk Tracker. The 3 approaches can be used in combination.

estimating correlations. An approach with 2 correlation parameters would follow from a KMV-portfolio model like that of Gordy (2003) in which the asset returns of the indebted companies are driven by 1 systematic macro factor, orthogonal industry factors, and idiosyncratic factors. For each company, the correlation between the asset return and the macro factor would equal ρ_{ex} and the (additional) correlation between the asset return and the industry factor would be $(\rho_{in} - \rho_{ex})$.

We know Moody's DS for all 169 transactions. For 92 transactions we have enough information about industry diversification to carefully derive the ADS. In line with Moody's and S&P, we derive the ADS assuming an intraindustry asset correlation ρ_{in} of 20% and an interindustry asset correlation ρ_{ex} of 2%. Alternatively, we use interindustry correlations of 0% and 4% for a robustness check. We use Moody's DS to analyze the full sample and, in addition, the ADS to analyze the reduced sample of 92 transactions. We refer to this subset whenever we use the lognormal loss rate distribution.

B. Descriptive Statistics and Methodology

Table 1 gives the distribution of the 169 transactions across CLO/CBO and true sale/synthetic transactions and across years. Here, 57% of the transactions are CBO transactions; 54% are synthetic. Some 136 transactions are arranged by banks, and 33 by investment firms. The latter buy existing bonds and securitize them. Of these 33 transactions, 15 also include some loans. We classify these transactions as CBO transactions. Of the CBO transactions, ⅓ are originated by investment firms.

TABLE 1

Number of Transactions

Panel A of Table 1 reports the number of transactions in the sample differentiating CLO and CBO transactions as well as true sale and synthetic transactions. Panel B gives the annual distribution of transactions.

Panel A. Number of Transactions Differentiated for CLO/CBO and True Sale/Synthetic

	True Sale	Synthetic	Total
CLO	30	43	73
CBO	48	48	96
Total	78	91	169

Panel B. Distribution of Transactions over Time

	Year								
	1997	1998	1999	2000	2001	2002	2003	2004	2005
No. of transactions	1	1	12	26	40	42	16	19	12

Table 2 presents the means and standard deviations of WADP, DS, FLP, and TLP in synthetic transactions. Using a loss given default of 50% with a few exceptions, the expected default loss of an asset pool is about ½ the WADP. The data are presented separately for the 4 subsamples of true sale/CLO, synthetic/CLO, true sale/CBO, and synthetic/CBO transactions. Table 2 indicates several interesting properties. The mean of the WADP is much higher for true sale than synthetic transactions. It is also clearly higher for synthetic CLO than synthetic CBO transactions. On average, CLO transactions are much better diversified

https://doi.org/10.1017/S0022109012000336 Published online by Cambridge University Press

than CBO transactions. In a CLO transaction a bank can easily securitize many loans from its loan book to obtain a high DS. Buying bonds is often costly, since the bond market is rather illiquid. Therefore the DS tends to be smaller in CBO transactions.

TABLE 2

Descriptive Statistics

Table 2 reports the means and standard deviations of transaction characteristics differentiating CLO and CBO transactions as well as true sale (TS) and synthetic (SYN) transactions. WADP and Moody's DS are the weighted average default probability and Moody's diversity score of the asset pool, respectively. FLP is the initial size of the FLP, and TLP the nonsecuritized senior tranche as a percentage of the asset pool volume in synthetic transactions. The numbers in parentheses for TLP in CLO transactions are obtained if 3 fully funded synthetic Geldilux transactions are excluded.

	TS/CLO	SYN/CLO	TS/CBO	SYN/CBO
WADP: mean	7.5%	3.8%	13.2%	1.9%
WADP: std.	7.5%	3.1%	9.8%	3.2%
DS: mean	87	89	34	56
DS: std.	46	30	11	26
FLP: mean	6.1%	2.9%	12.1%	3.6%
FLP: std.	4.8%	1.5%	6.2%	2.6%
TLP: mean		80% (86%)		87%
TLP: std.		23% (7%)		7%

The average size of the FLP is higher for true sale than for synthetic transactions, and within these subsets it is higher for CBO than for CLO transactions. The average size of the FLP exceeds the average expected loss, which is about ½ of the WADP. Thus, the averages satisfy the condition in Lemma 2a. Also, the average size of the FLP is smaller than the average WADP, except for synthetic CBO transactions. The condition in Lemma 2b and the condition $h < n(h+\sigma)/(1-N(h+\sigma))$ in Lemma 3 always hold, based on an ADS with $\rho_{\text{ex}} = 2\%$. The average TLP in synthetic transactions is about 87% of the asset pool volume with a standard deviation of only 7% if we exclude 3 atypical Geldilux transactions. These are the only fully funded synthetic CLO transactions (i.e., TLP = 0).

In the following we test the hypotheses presented in Section III. We hypothesize that loss allocation depends on asset pool quality, on other characteristics of the securitization transaction, and on exogenous factors such as the attitudes of investors and rating agencies, market imperfections, and originator characteristics. The function relating loss allocation to all these factors is assumed to be the same for all CDO transactions. We try to find out the properties of this function. One difficulty of this approach is that the originator simultaneously chooses asset pool quality and loss allocation. Asset pool quality cannot be viewed as an exogenous determinant of loss allocation, and vice versa. Yet, if a function exists that relates loss allocation to asset pool quality in equilibrium, then the best the originator can do is to adhere to this function. Originator characteristics are likely to affect loss allocation. Therefore, we include various characteristics of originating banks as controls. Moreover, we run the regressions for loss allocation separately for the 4 subsets of true sale/CLO, true sale/CBO, synthetic/CLO, and synthetic/CBO transactions to find out whether these subsets differ systematically.

Banks and investment firms are originators. While many characteristics of banks are known, those of investment firms are largely unknown. We include the following bank control variables, obtained from the Bank Scope Database: i) Basel ratios: the tier 1 capital ratio and the total capital ratio; ii) capital structure: equity/total assets; iii) asset structure: loans/total assets; iv) profitability: return on average equity capital in the transaction year, the average return for 1994–2004, and the standard deviation of these returns as a proxy for profitability risk; v) Tobin's Q to proxy for the bank's growth potential; and vi) the bank's rating to proxy for its funding cost. Rating is always captured by an integer variable that equals -1 for an Aaa rating and declines by 1 for every notch, with -16 for a rating of B3.

For each characteristic we attach a residual dummy RD $= 1$ to those originators for which the characteristic is not known, otherwise RD $= 0$. We use regressions of the type

$$y \quad = \quad a + b\, x_1 + c(1 - \mathrm{RD})\Delta x_2 + d\,\mathrm{RD} + \varepsilon, \tag{3}$$

where x_1 is the vector of explaining variables other than originator characteristics, Δx_2 the vector of (bank characteristics $-$ its sample average), and ε the usual error term. This approach implies that for banks with a known characteristic, the variation in this zero-mean characteristic is taken into consideration, while for the other originators a fixed effect is assumed. If a variable does not add to the explanatory power of a regression, then we often eliminate it from the regression. We always run the regressions with constants, but to save space we do not always report their values.

C. The Quality of the Asset Pool

The quality of the underlying asset pool is a core variable. Asset pool quality and loss allocation may be interdependent. Therefore, we ask 3 questions: i) Does the originator follow a homogeneous quality policy, that is, is a low (high) WADP associated with a high (low) DS? ii) Does loss allocation affect WADP and DS? iii) Do originator characteristics affect the choice of asset pool quality? Figure 1 illustrates the relation between WADP and DS (Graph A) and the ADS2 (Graph B), respectively. Neither figure indicates a strong relation between WADP and asset pool diversification. Hence, the figures provide, at best, partial support for homogeneous quality choice.

Next, we run ordinary least squares (OLS) regressions to check the relation between WADP and DS. In Table 3, we first regress WADP for all 169 transactions on the inverse lnDS and originator characteristics. DS turns out to be insignificant. Investment firms and banks with a higher total capital ratio tend to choose asset pools with higher WADP. Banks with a strong equity buffer might take higher default risks, which they then securitize. Other originator characteristics have no significant impact. In accordance with Table 2, true sale CBO transactions have significantly higher WADPs.

Since the choice of WADP might differ for subsets of transactions, we also run the 1st regression separately for TS/CLO, TS/CBO, SYN/CLO, and SYN/CBO transactions. The regression coefficient of the inverse lnDS is significant only for the 30 true sale CLO transactions (2nd regression in Table 3); a higher DS tends to lower WADP.

FIGURE 1

Diversity Scores and Weighted Average Default Probabilities

Graph A of Figure 1 shows a scatter plot of Moody's diversity score (DS) and the weighted average default probability (WADP) for 169 transactions. Graph B shows the adjusted diversity score (ADS) with $\rho_{ex} = 0.02$ and the WADP for 92 transactions.

Graph A. DS and WADP *Graph B. ADS2 and WADP*

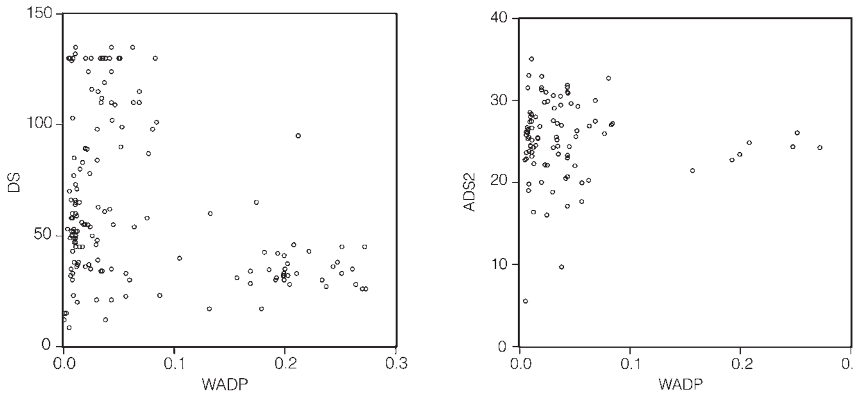

TABLE 3

Results for WADP

Table 3 reports the coefficients (Newey-West (1987) heteroskedasticity-adjusted p-values in parentheses) of OLS regressions of WADP on various variables, without showing the regression constants. Inverse lnADS2 is the inverse log-adjusted diversity score assuming a default correlation between industries of 2%. The investment firm dummy variable is 1 if an investment firm is the originator, and 0 otherwise. The CBO dummy variable is 1 for a CBO transaction, and 0 otherwise. The synthetic dummy variable is 1 for a synthetic transaction, and 0 otherwise. The ΔTotal capital ratio is the total capital ratio of the originating bank in the transaction year minus the average total capital ratio in the sample (see equation (3)). The adjusted R^2 is shown in the last row.

Explained Variable	WADP (%)						
Sample of Transactions	All 169	30 TS, CLO	48 TS, CBO	43 SYN, CLO	48 SYN, CBO	92	92
Inverse lnDS	−4.7 (0.6860)	92.2 (0.0209)	3.0 (0.9096)	−17.2 (0.4890)	−0.54 (0.9582)		
Inverse lnADS2						8.63 (0.3791)	−3.78 (0.6785)
Investment firm dummy	9.5 (0.0000)	0.40 (0.9284)	13.17 (0.0000)		−0.26 (0.7481)	5.41 (0.0028)	8.57 (0.0247)
Loss share × CBO							−7.87 (0.2695)
Loss share × (1 – CBO)							−9.41 (0.2623)
Support Prob × CBO						−27.1 (0.0000)	
Support Prob × (1 – CBO)						−28.9 (0.0000)	
CBO dummy (1 − Synthetic dummy)	6.9 (0.0000)						
ΔTotal capital ratio	1.2 (0.0021)	0.95 (0.1633)	3.056 (0.0000)	0.541 (0.1536)	0.306 (0.1649)	0.259 (0.2517)	0.898 (0.0951)
Adj. R^2	0.508	0.221	0.602	0.0397	−0.0553	0.650	0.180

Also, loss allocation might affect the choice of asset pool quality. As Table 2 indicates, a higher WADP tends to lead to a higher FLP. When we include FLP in the regression, we find a strongly significant, positive regression coefficient (not shown). But due to endogeneity issues, it would be dangerous to conclude that a higher FLP generates a higher WADP. The loss share and the support probability are less prone to endogeneity problems. Therefore, we use them as regressors. Since both are sensitive to diversification, we use the subsample of 92 transactions for which we know the more reliable ADS2, using an interindustry asset correlation ρ_{ex} of 2%. Since loss sharing may differ for CLO and CBO transactions, we multiply the loss measures by the CBO or (1 − CBO) dummy variables. The last 2 regressions in Table 3 indicate that the ADS2 coefficient is insignificant, while the investment firm dummy variable has a positive and significant impact on WADP. The support probability coefficient is strongly significant and negative for both CBO and CLO transactions, but the loss share is insignificant. By Lemma 3, one of these loss measures should have a clear impact if the other one does not. We defer a more detailed discussion of these findings to the next subsection.

Table 4 reports regressions to explain DS and ADS2. For the full sample of 169 transactions, the inverse lnDS does not depend significantly on WADP, and the investment firm dummy variable is insignificant. As suggested by Table 2, DS is lower for CBO transactions, but higher for synthetic transactions. The total capital ratio coefficient is only weakly significant. Other originator

TABLE 4

Results for DS and ADS2

Table 4 reports the coefficients (Newey-West (1987) heteroskedasticity-adjusted p-values in parentheses) for OLS regressions of 1/lnDS and 1/lnADS2 on various variables, without showing the regression constants. The investment firm dummy variable is 1 if an investment firm is the originator, and 0 otherwise. The CBO dummy variable is 1 for a CBO transaction, and 0 otherwise. The synthetic dummy variable is 1 for a synthetic transaction, and 0 otherwise. The ΔTotal capital ratio is the total capital ratio of the originating bank in the transaction year minus the average total capital ratio in the sample (see equation (3)). The adjusted R^2 is shown in the last row.

Explained Variable	Inverse lnDS					Inverse lnADS2	
Sample of Transactions	All 169	30 TS, CLO	48 TS, CBO	43 SYN, CLO	48 SYN, CBO	92	92
WADP (%)	4.85 (0.3376)	18.5 (0.0579)	1.27 (0.9069)	−8.46 (0.5791)	−0.69 (0.9585)	−0.017 (0.8475)	−0.095 (0.2047)
Investment firm dummy	0.004 (0.6763)	0.047 (0.0141)	−0.0118 (0.3448)		−0.0155 (0.0358)	−0.002 (0.7910)	−0.0013 (0.8733)
Loss share × CBO							0.059 (0.0550)
Loss share × (1 − CBO)							0.042 (0.1243)
Support Prob × CBO						0.048 (0.0163)	
Support Prob × (1 − CBO)						0.030 (0.1225)	
CBO dummy	0.040 (0.0000)						
Synthetic dummy	−0.018 (0.0645)					−0.030 (0.0258)	−0.029 (0.0225)
ΔTotal capital ratio	−0.008 (0.0219)	−0.001 (0.7187)	−0.011 (0.0242)	−0.0016 (0.2935)	−0.019 (0.1355)	0.000 (0.9745)	−0.000 (0.9821)
Adj. R^2	0.347	0.472	0.049	−0.013	0.209	0.182	0.186

characteristics appear to be irrelevant. In separate regressions for each of the 4 subsamples, we only find a significant WADP coefficient for the subsample of 30 true sale CLO transactions (2nd regression in Table 4), in line with the significant coefficient of DS on WADP in Table 3. We also check the subsample of 92 transactions to find out whether loss allocation has any relation to ADS2 (last 2 regressions in Table 4). The support probability coefficient is significant for the CBO, but not for the CLO transactions. The loss share coefficient is weakly significant only in CBO transactions. Hence, similar to Table 3, there is very little support for dependence between DS and WADP.

While DS is primarily driven by CLO versus CBO transactions, an endogeneity problem regarding the impact of DS on WADP might exist. Therefore, we run a 2-stage least squares (2SLS) regression to explain WADP. As the DS is much higher in CLO than in CBO transactions, we use the CBO dummy variable as an instrumental variable for DS. The findings of this exercise do not support an endogeneity effect.

In summary, the choices of WADP and DS appear to be independent. Investment firms and banks with higher total capital ratios appear to prefer transactions with higher WADPs. Otherwise, originator characteristics do not seem to be relevant. While WADP is partly explained by the support probability, the loss share explains neither WADP nor ADS.

D. Loss Allocation to the FLP

1. Size of FLP

Now the core hypotheses on loss sharing will be tested. Hypothesis 1 states that the size of the FLP is inversely related to the quality of the asset pool. This hypothesis is strongly confirmed by the regressions in Table 5. The first 2 regressions are based on the full sample of 169 transactions using the DS, while the last regression is based on the subsample of 92 transactions using ADS. The results indicate that even for an asset pool with excellent quality, given by a WADP close to 0 and a DS of 135, the FLP is positive. Using the coefficients of the 1st regression, the estimated FLP equals $-8.5+46.6/\ln 135 = 1.00$ (%). This suggests that investors are concerned about information asymmetry and other potential problems of securitization transactions even if the published asset pool quality is very good. It may also explain why the FLP increases with WADP at a slope below the loss given default of ½. Interestingly, the regression coefficient of the synthetic dummy variable is significantly negative in the 2nd regression, but insignificant in the last regression, which uses the more sophisticated diversification measure.

We also regress the FLP size on asset pool quality for the 4 subsets of transactions. The results (regressions 3–6) are very similar for true sale CLOs and for true sale CBOs. For synthetic CLOs and synthetic CBOs, WADP still remains strongly significant, but inverse lnDS turns insignificant. It may be that investors are less concerned about diversification in synthetic transactions because they interpret the TLP as a strong quality signal.

We also check whether the FLP is higher in a managed (dynamic) transaction than in a static transaction. In the 2nd and the last regression, the dummy variable being 1 for a managed transaction, and 0 otherwise, is insignificant. This may be

TABLE 5

Results for the Size of the FLP

Table 5 reports the coefficients (Newey-West (1987) heteroskedasticity-adjusted *p*-values in parentheses) for OLS regressions explaining the size of the FLP. WADP is the weighted average default probability of the asset pool. Inverse lnADS2 is the inverse log-adjusted diversity score assuming a default correlation between industries of 2%. The synthetic dummy variable is 1 for a synthetic transaction, and 0 otherwise. The CBO dummy variable is 1 for a CBO transaction, and 0 otherwise. The managed dummy variable is 1 for a managed transaction, and 0 otherwise. The adjusted R^2 is shown in the last row.

Explained Variable	Size of FLP (%)						
Sample of Transactions	169	169	30 TS, CLO	48 TS, CBO	43 SYN, CLO	48 SYN, CBO	92
Constant	−8.50 (0.0188)	−5.25 (0.1865)	−11.83 (0.0006)	−15.88 (0.0611)	2.60 (0.1953)	4.76 (0.0341)	−15.0 (0.0421)
WADP (%)	0.347 (0.0000)	0.296 (0.0000)	0.294 (0.0003)	0.237 (0.0027)	0.309 (0.0002)	0.447 (0.0159)	0.356 (0.0000)
Inverse lnDS	46.6 (0.0028)	40.5 (0.0117)	65.7 (0.0001)	79.5 (0.0056)	−4.19 (0.6068)	−8.15 (0.3278)	
Inverse lnADS2							60.4 (0.0095)
Synthetic dummy		−2.07 (0.0029)					−1.03 (0.1345)
CBO dummy		0.14 (0.7963)					0.82 (0.2022)
Managed dummy		−0.49 (0.4417)					−0.68 (0.3263)
Adj. R^2	0.546	0.574	0.766	0.343	0.393	0.328	0.625

due to the strict rules on replenishment/substitution. Also, originator characteristics have no impact on the size of the FLP. Originators with more valuable real options as indicated by Tobin's Q do not seem to prefer a smaller FLP.

We also run the FLP regressions including the issuance date and the iBoxx spread as regressors. The iBoxx spread is the difference between the iBoxx for BBB bonds and the government iBoxx for a maturity of 3–5 years. Both regressors have no significant impact (not shown). The insignificant issuance date coefficient does not support the claim that the rating agencies relaxed their rating standards over time, implying smaller FLPs. Finally, we run the regression for the subsample of 92 transactions with ADSs based on interindustry correlations ρ_{ex} of 0% or 4%. We do not report the results because they are very similar to those for $\rho_{\text{ex}} = 2\%$.

2. Loss Share and Support Probability

Next, we analyze the relation between asset pool quality and sophisticated loss-sharing measures. Using the interindustry correlation $\rho_{\text{ex}} = 2\%$, the loss share of the FLP has a mean of 86.1% and a standard deviation of only 8.4%. This indicates that the FLP takes a high share of the expected default losses. For the support probability, the mean is 87.6% and the standard deviation 14.7%. This mean is also quite high. In accordance with Lemma 2a, the average loss share declines from 91.6% to 82.3% if ρ_{ex} increases from 0% to 4%, and the average support probability declines slightly from 88.25% to 87.57%. Surprisingly, the support probability is almost constant. This indicates that the cumulative lognormal distributions, generated by different interindustry correlations, intersect at loss rates that are only slightly below the FLP.

When regressing the loss share on WADP and inverse lnADS2 only, it turns out that WADP is completely insignificant, while the coefficient of the inverse ADS2 is positive and significant (Table 6, 1st regression). However, the explanatory power of the regression is only about 5.6%. Adding the CBO dummy variable as a regressor strongly weakens the significance of the ADS2 coefficient. The CBO coefficient is strongly significant, positive indicating a higher loss share in CBO transactions. The average loss share is 83.7% in CLO and 90% in CBO transactions, invalidating Hypothesis 6. This is consistent with higher expected losses of securitized bonds versus loans as well as with a strong reputation cost for default losses in CLO transactions. This cost might induce a strong (unobservable) monitoring effort of the originator in these transactions that partially substitutes for the loss share. But it also points to operational risk in CLO transactions. If an originator does not intend future securitization transactions, she might not care about reputation costs and enjoy the private benefits of moral hazard. Again, the coefficients of the issuance date and the iBoxx spread are insignificant (not shown).

We also run the 1st regression in Table 6 separately for the 4 subsets of transactions (regressions 3–6). For true sale CLOs, synthetic CLOs, and synthetic CBOs, WADP and inverse lnADS2 have no significance, and the adjusted R^2s are negative or slightly above 0. Only for the small sample of 15 true sale CBOs do we find a strongly significant negative effect of WADP, driven by 4 transactions with a WADP above 20%, while all other transactions have a WADP below 6% (4th regression in Table 6). Apart from these few transactions, the empirical evidence clearly indicates that the loss share is largely independent of portfolio quality. This conclusion is also supported by the regressions in Tables 3 and 4, where the loss share adds nothing or little to the explanation of asset pool quality. Thus, Hypothesis 2 is clearly supported.

Next, we run similar regressions for the support probability (right-hand side of Table 6). For the full sample of 92 transactions, the WADP effect is strongly significant and negative. ADS2 is weakly significant, but insignificant if the CBO dummy variable is included. These 2 regressions have an impressive explanatory power of about 60%. These findings are not surprising in view of the previous finding that the loss share is independent of the asset pool quality. Lemma 3 then implies that the support probability should react inversely to WADP and to ADS. Again, the issuance date and the iBoxx spread have no significant effects. Analyzing the 4 subsets of transactions separately (last 4 regressions), WADP always has a strongly significant negative impact, while inverse lnADS2 has a positive impact that, however, is significant only for synthetic CBOs. Hence, we conclude that the support probability varies with asset pool quality, in particular WADP, while the loss share does not. A loss share that is largely independent of the asset pool quality appears to be the market norm. These findings need not necessarily be driven by information asymmetry, because we can only test for effects of asset pool quality.

More importantly, adding originator characteristics as regressors does not improve the explanatory power (not shown). Hence, loss allocation appears to be driven by market forces, not by originator characteristics. This may be surprising if one believes in the cooperation of originators and rating agencies to maximize

TABLE 6

Results for the Share of Expected Losses and the Support Probability of the FLP

Table 6 reports the coefficients (Newey-West (1987) heteroskedasticity-adjusted p-values in parentheses) for OLS regressions explaining the share of expected losses of the FLP and the support probability of the FLP. WADP is the weighted average default probability of the asset pool. Inverse lnADS2 is the inverse log-adjusted diversity score assuming a default correlation between industries of 2%. The CBO dummy variable is 1 for a CBO transaction, and 0 otherwise. Date is the issuance date. The adjusted R^2 is shown in the last row.

	Share of Expected Losses (%)						Support Probability (%)					
Sample of Transactions	92	92	20 TS, CLO	15 TS, CBO	34 SYN, CLO	23 SYN, CBO	92	92	20 TS, CLO	15 TS, CBO	34 SYN, CLO	23 SYN, CBO
Constant	67.3 (0.0000)	69.4 (0.0000)	71.3 (0.0532)	94.1 (0.0000)	120.9 (0.0018)	30.8 (0.4341)	74.8 (0.0000)	80.2 (0.0000)	72.4 (0.0936)	103.7 (0.0000)	78.7 (0.0687)	62.8 (0.0027)
WADP (%)	−0.124 (0.2950)	−0.121 (0.3556)	0.325 (0.1326)	−0.556 (0.0007)	−0.35 (0.6195)	0.187 (0.2308)	−1.99 (0.0000)	−2.00 (0.0000)	−1.46 (0.0000)	−2.31 (0.0000)	−4.07 (0.0001)	−0.63 (0.0000)
Inverse lnADS2	61.7 (0.0145)	40.4 (0.0687)	36.5 (0.7537)	6.63 (0.7277)	−122.0 (0.3243)	187.9 (0.1432)	64.6 (0.0613)	43.2 (0.1476)	65.9 (0.6275)	1.08 (0.9456)	69.3 (0.6308)	111.6 (0.0664)
CBO dummy		6.38 (0.0002)						5.99 (0.0002)				
Date		0.14 (0.2680)					0.08 (0.6777)					
Adj. R^2	0.056	0.166	−0.056	0.620	−0.023	0.014	0.593	0.632	0.395	0.958	0.486	0.035

their joint benefit. As a caveat, our set of originator characteristics may miss some relevant characteristics.

The assumption of a lognormal loss rate distribution is sometimes criticized. If one simulates the loss rate distribution of a loan portfolio period by period, it turns out that the distribution is in some cases better approximated by a gamma distribution. Therefore, we check robustness by using a 2-parameter gamma distribution. For each transaction, the expected loss rate and the loss rate variance, based on ADS2, are translated into the parameters of a gamma distribution. While for the 92 transactions the average loss share (average support probability) assuming a lognormal distribution is 86.1% (87.6%), it is 84.3% (85.7%) assuming a gamma distribution. Hence, it is not surprising that the regression results for loss shares and support probabilities based on the gamma distribution (not shown) are similar to those based on the lognormal distribution.

E. Loss Allocation to the TLP

Comparing true sale and synthetic transactions permits us to better understand the determinants of loss allocation. Hypothesis 3 claims that synthetic transactions are preferred for high-quality asset pools. This hypothesis is clearly supported by the 1st probit regression in Table 7. It is also supported if we analyze CLO and CBO transactions separately (2nd and 3rd regressions). For CLO transactions the DS effect is insignificant, perhaps because these transactions tend to be well diversified.

TABLE 7

Results for the Synthetic Dummy Variable

Table 7 reports the coefficients (with *p*-values in parentheses) of binary probit regressions explaining the synthetic dummy variable, without showing the regression constants. The synthetic dummy variable is 1 for a synthetic transaction, and 0 otherwise. iBoxx spread is the spread between the BBB iBoxx and the government iBoxx for 3–5 years maturity. The ΔOriginator's rating is the originator's rating minus the average originator rating in the sample (see equation (3)). The ΔTobin's Q and ΔTotal capital ratio are defined analogously. The Originator rating dummy variable is 1 for originators without a rating, and 0 otherwise. The last row shows the McFadden R^2.

Explained Variable	Synthetic Dummy					
Sample of Transactions	169	73 CLO	96 CBO	151	151	151
WADP (%)	−0.11 (0.0000)	−0.077 (0.0275)	−0.135 (0.0000)		−0.10 (0.0000)	−0.11 (0.0000)
Inverse lnDS	−6.71 (0.0050)	−3.47 (0.5212)	−11.43 (0.0014)		−6.86 (0.0085)	−6.19 (0.0523)
iBoxx spread (%)						0.879 (0.0098)
ΔOriginator's rating				0.225 (0.0033)	0.22 (0.0107)	0.33 (0.0013)
Originator rating dummy				−1.54 (0.0000)	−1.29 (0.0002)	−0.77 (0.0631)
ΔTobin's Q						−1.12 (0.0001)
ΔTotal capital ratio						0.198 (0.0947)
McFadden R^2	0.265	0.085	0.503	0.190	0.362	0.440

Hypothesis 5 claims that originators with a good rating are less interested in funding through securitization. The 4th regression supports this hypothesis,

https://doi.org/10.1017/S0022109012000336 Published online by Cambridge University Press

showing a significant, positive impact of the originator rating, while the originators without a rating, mostly investment firms, appear to prefer true sale transactions. Combining originator rating and portfolio quality in the 5th regression clearly improves the explanatory power of the regression.

In the last regression, the explanatory power is further improved by including the iBoxx spread, the originator's Tobin's Q, and her total capital ratio. The iBoxx spread has a significant, positive impact on the preference for synthetic transactions. Hence, it appears that originators are less interested in funding through securitization when credit spreads are high. The issuance date is insignificant (not shown). The significant, negative coefficient of Tobin's Q suggests that it may not pay for originators with attractive outside options to retain the risk of a TLP. Originators with a high total capital ratio may prefer to retain the TLP risk because their cost of retention may be smaller than that of selling.

We interpret the strong preference for synthetic transactions of originators with a good rating as evidence of a funding cost effect. This effect may dominate information-sensitivity effects and explain why the FLP, the most information-sensitive tranche, is at least partly sold, while the TLP, the least information-sensitive tranche, is not. Graph A of Figure 2 shows credit spreads of European bank bonds with a maturity of at least 4 years, rated Aa3 and better, issued between 2000 and 2005. Graph B shows the credit spreads of Aaa-rated CDO tranches with a maturity of at least 4 years in our sample. The mean credit spread is 9.1 basis points (bp) for the bank bonds and 40.6 bp for the Aaa tranches. The minimum (maximum) spread of the bank bonds is −27 (+22) bp, while it is +1 (+100) bp for the Aaa tranches.[14] Hence, credit spreads of highly rated bank bonds are often lower than those of Aaa tranches.

FIGURE 2

Credit Spreads of Bank Bonds and Aaa Tranches

Graph A of Figure 2 shows a histogram of the credit spreads over Euro Interbank Offered Rate/London Interbank Offered Rate (EURIBOR/LIBOR) of 118 European bank bonds with a maturity of at least 4 years, rated Aa3 or better, issued between 2000 and 2005. Data are obtained from DealScan. Graph B shows the credit spreads over EURIBOR/LIBOR of the 135 Aaa-rated CDO tranches with a maturity of at least 4 years of our securitization sample.

Graph A. Credit Spreads of Bank Bonds *Graph B. Credit Spreads of Aaa Tranches*

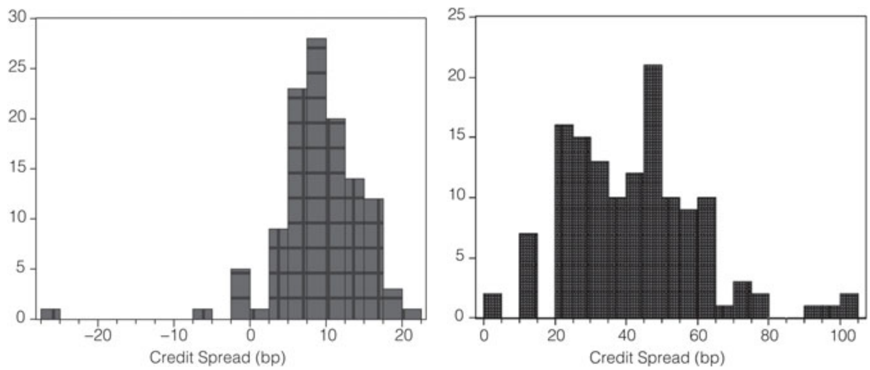

[14]Similar to Aaa tranches, standard bonds issued by large, well-rated banks also represent well-diversified risks. They predominantly default in very bad macro states, similar to Aaa tranches.

In Table 8, we regress the credit spreads of bank bonds and CDO tranches with a rating of at least Aa3 on the iBoxx spread, negative rating, squared rating, and bond maturity. Table 8 indicates that the spreads of bank bonds and CDO tranches are lower for a better rating and a shorter maturity. The tranche spreads increase with a higher iBoxx spread, while the bank bond spreads do not. To illustrate the funding decision, compare the estimated credit spread of an Aa2-bank bond and of an Aaa tranche for a true sale transaction, given 5 years maturity and an iBoxx spread of 2%. The estimated credit spread of the Aa2-bank bond is 7.9 bp, and that of the Aaa tranche is 42.7 bp. Thus, for highly rated banks, funding through a large Aaa tranche in a true sale transaction likely implies a higher funding cost than issuing standard Aa2 bonds. These banks may enjoy a strong reputation, which is not fully reflected in their bond rating, but in their credit spreads. Also, investors may believe that big banks are too big to fail. Conversely, an Aaa tranche in a securitization transaction may face some investor skepticism because securitization transactions are relatively new instruments with little performance history.

TABLE 8

Results for the Credit Spreads of Bank Bonds and CDO Tranches

Table 8 reports the coefficients (Newey-West (1987) heteroskedasticity-adjusted p-values in parentheses) for OLS regressions explaining the credit spreads of European bank bonds and CDO tranches with a rating of Aa3 and better. Rating is −1 for Aaa, −2 for Aa1, −3 for Aa2, and −4 for Aa3. iBoxx spread is the spread between the BBB iBoxx and the government iBoxx for 3–5 years maturity. The synthetic dummy variable is 1 for a synthetic transaction, and 0 otherwise. The adjusted R^2 is shown in the last row.

Explained Variable	Credit Spread of Bank Bonds	Credit Spread of CDO Tranches
No. of obs.	118	226
Constant	−7.99 (0.0273)	−21.3 (0.0419)
iBoxx spread (%)	−0.695 (0.5545)	11.85 (0.0055)
– Rating	5.01 (0.0390)	32.46 (0.0011)
Rating^2	−0.312 (0.5040)	−3.93 (0.0787)
Maturity of bond	1.002 (0.0000)	2.36 (0.0037)
Synthetic dummy		4.51 (0.2272)
Adj. R^2	0.286	0.420

Finally, we analyze the size of the TLP. We exclude the 3 atypical fully funded Geldilux transactions. In the 1st regression of Table 9, WADP has a strongly negative effect on the size of the TLP, while the impact of lnDS is U-shaped. For small DSs up to about 28, the estimated TLP declines with an increasing DS, and for higher DSs it increases. There are only a few transactions with a DS below 28. Therefore, the TLP mostly increases with the DS. Thus, Hypothesis 4 is clearly confirmed. Differentiating between CLO and CBO transactions (2nd and 3rd regressions) shows similar results. But the DS turns insignificant in the strongly diversified CLO transactions. The explanatory power of the 1st regression can be improved slightly by including the investment firm dummy variable (last regression). The negative coefficient indicates that investment firms tend to retain smaller TLPs. Other originator characteristics do not appear to have a significant impact on the size of the TLP. Also, the issuance date and the iBoxx spread appear to be irrelevant.

TABLE 9

Results for the Size of the TLP

Table 9 reports the coefficients (Newey-West (1987) heteroskedasticity-adjusted p-values in parentheses) of OLS regressions explaining the size of the TLP in synthetic transactions. The investment firm dummy variable is 1 if the originator is an investment firm, and 0 otherwise. The sample contains 86 transactions. The adjusted R^2 is shown in the last row.

Explained Variable	Size of TLP (%)			
Sample of Transactions	86	40 CLO	46 CBO	86
Constant	−0.066 (0.7861)	0.958 (0.6657)	−0.162 (0.4871)	−0.017 (0.9462)
WADP (%)	−0.015 (0.0000)	−0.018 (0.0000)	−0.014 (0.0000)	−0.015 (0.0000)
lnDS	0.14 (0.0000)	0.025 (0.9224)	0.16 (0.0000)	0.14 (0.0001)
Inverse lnDS	1.56 (0.0007)	−0.60 (0.8987)	1.73 (0.0001)	1.47 (0.0017)
Investment firm dummy				−0.06 (0.0007)
Adj. R^2	0.576	0.592	0.544	0.588

Comparing our findings for the FLP and TLP, the differences are striking. While the FLP size *reacts inversely* to asset pool quality, the TLP *increases*. This indicates that the choices are driven by different motives. Presumably, the FLP mitigates problems of information asymmetry; a lower asset pool quality induces a higher FLP, providing more protection to investors. The TLP does not protect investors so that investor protection is irrelevant. Better asset pool quality lowers the Aaa-attachment point so that the TLP increases. Yet, the TLP should be sold, since it is least information sensitive. But funding through Aaa tranches is not attractive for highly rated banks. This is reinforced in times of high credit spreads. Also, selling a large Aaa tranche transfers little default risk, so that banks may consider it a cost-ineffective tool for risk management.

F. Robustness Checks and Discussion

A potential critique of OLS regressions to explain the FLP and the TLP is that these variables are constrained to the (0,1) range. For a robustness check we transform the FLP and the TLP so that the transformed variable varies between plus and minus infinity. The regression results basically stay the same.

We already checked for potential endogeneity problems regarding the choice of asset pool quality. In the other regressions we see little potential for endogeneity. These regressions address loss allocation, given exogenous originator characteristics and attitudes of investors and rating agencies. While originator characteristics vary across transactions and are therefore used as controls, attitudes of investors and agencies should be similar for all transactions. Including a CBO dummy variable as a regressor does not lead to endogeneity, because CLO and CBO transactions represent 2 types of transactions. Including a synthetic dummy variable is more prone to endogeneity problems. But the findings of separate regressions for true sale and synthetic transactions are similar.

Our assumption that Moody's quality assessments of the asset pool are unbiased, is difficult to test. Assessment biases might be driven by joint interests of originators and rating agencies, as suggested by the Congressional hearings on mortgage-backed securitizations of Oct. 22, 2008. Our findings, however, indicate that the rating methods have been rather stable over the sampling period. We are not aware of any major changes in rating methodologies for corporate loans/bonds and CDO transactions after our sampling period.

The discussion about the best way to measure the DS has led us not only to consider Moody's DS, but also ADS based on interindustry default correlations of 0%, 2%, or 4%. The regression results are similar. Sometimes the results are somewhat stronger for 4%. This may indicate that 4% was considered more realistic. If Moody's over- or underestimated all WADPs by the same multiple, this would not affect our findings, apart from the level of the regression coefficients.

The assumption of a constant loss given default ignores that this parameter changes with business conditions that might impact loss allocation. Uncertainty in WADP and correlations should affect the results because the formulas for the loss allocation measures indicate nonlinear effects, as illustrated in Tarashev (2010), for the value at risk. For example, ADS is a declining convex function of correlation so that correlation uncertainty should upward bias the expected ADS. Given an increasing convex function of the loss share in ADS, the expected loss share would also be upward biased. Similarly, uncertainty in WADP should upward bias the expected loss share, since the loss share tends to be a declining convex function in WADP. Parameter uncertainty might also upward bias the support probability.

On a general level, our sample of 169 European CDO transactions is modest even though it covers about ½ of the European transactions in the sampling period. This clearly limits the generality of our findings. The number of those transactions in the United States was much larger. It would be useful to compare our findings with those from the United States. It is dangerous to apply our findings to securitizations of other claims like mortgage-backed loans, credit card claims, etc. Asset pool quality in these securitizations is measured by other criteria; often, diversification of these pools is much stronger. This may lead to other loss allocations.

V. Conclusion

This paper investigates loss allocation to first, second, and third loss positions in securitizations of corporate loans and bonds. The originator cooperates with investors and rating agencies to determine the quality of the underlying asset pool and the loss allocation. A sample of European securitization transactions of corporate loans and bonds shows that the size of the first loss position (FLP) is strongly and inversely related to asset pool quality, while the third loss position (TLP) is positively related. The FLP likely serves to mitigate information asymmetry problems, in contrast to the TLP. The general guideline for structuring transactions appears to be that the FLP should cover a high share of the expected default losses, *largely independent* of the asset pool quality and of originator characteristics. The support probability of the FLP (i.e., the probability that the FLP absorbs all losses) is inversely related to the WADP of the asset pool, as predicted by theory.

The loss share is higher in CBO than in CLO transactions. In the latter, originators can mitigate default losses through their loan monitoring. The smaller loss share in CLO transactions suggests that rating agencies and investors anticipate a strong monitoring effort of the originator in CLO transactions, due to high reputation costs. Also, higher default rates of bonds relative to loans may motivate a higher loss share in CBO transactions.

Asset pool quality *and* originator rating positively affect the originator's preference for a synthetic transaction with a large TLP. Its size increases with the quality of the asset pool, in contrast to the FLP. Retaining the TLP is also in conflict with the claim that the originator should sell the least information-sensitive tranche. Selling this tranche does not achieve a substantial risk transfer; moreover, credit spreads of Aaa tranches tend to be higher than those of standard Aa-bank bonds so that an originator with a good rating may consider funding through Aaa tranches too expensive.

Originator characteristics have a surprisingly small impact on loss allocation except for the choice between true sale and synthetic transactions. This indicates that choices are largely driven by attitudes of investors and rating agencies. We do not find evidence that rating standards declined over time.

This paper sheds some light on the regulation of securitizations. Forcing the originator to retain 5% of all tranches would impose a loss share of the same percentage on her. Forcing her to retain 10% of the FLP would impose a loss share of 7%–9% on her. Hence, the latter is presumably more effective to constrain adverse selection and moral hazard. Clearly, more empirical research is needed to better understand the various motives driving loss allocation in securitization transactions.

References

Acharya, V. V.; S. T. Bharath; and A. Srinivasan. "Does Industry-Wide Distress Affect Defaulted Firms? Evidence from Creditor Recoveries." *Journal of Financial Economics*, 85 (2007), 787–821.

Albrecher, H.; S. Ladoucette; and W. Schoutens. "A Generic One-Factor Lévy Model for Pricing Synthetic CDOs." In *Advances in Mathematical Finance,* M. C. Fu, R. A. Jarrow, J.-Y. J. Yen, and R. Elliott, eds. Boston, MA: Birkhaeuser (2007), 259–277.

Arrow, K. *Essays in the Theory of Risk Bearing.* Chicago, IL: Markham (1971).

Black, F., and M. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy,* 81 (1973), 637–654.

Boot, A. W. A., and A. V. Thakor. "Security Design." *Journal of Finance,* 48 (1993), 1349–1378.

Burtschell, X.; J. Gregory; and J.-P. Laurent. "A Comparative Analysis of CDO Pricing Models under the Factor Copula Framework." *Journal of Derivatives,* 16 (2009), 9–37.

Cebenoyan, A. S., and P. E. Strahan. "Risk Management, Capital Structure and Lending at Banks." *Journal of Banking and Finance,* 28 (2004), 19–43.

DeMarzo, P. "The Pooling and Tranching of Securities: A Model of Informed Intermediation." *Review of Financial Studies,* 18 (2005), 1–35.

DeMarzo, P., and D. Duffie. "A Liquidity-Based Model of Security Design." *Econometrica,* 67 (1999), 65–99.

Downing, C.; D. Jaffee; and N. Wallace. "Is the Market for Mortgage-Backed Securities a Market for Lemons?" *Review of Financial Studies,* 22 (2009), 2457–2494.

Downing, C., and N. Wallace. "Commercial Mortgage Backed Securities: How Much Subordination Is Enough?" Working Paper, Rice University (2005).

Duffie, D.; A. Eckner; G. Horel; and L. Saita. "Frailty Correlated Default." *Journal of Finance,* 64 (2009), 2089–2123.

Duffie, D., and N. Gârleanu. "Risk and the Valuation of Collateralized Debt Obligations." *Financial Analysts Journal,* 57 (2001), 41–59.

Emery, K. M., and R. Cantor. "Relative Default Rates on Corporate Loans and Bonds." *Journal of Banking and Finance,* 29 (2005), 1575–1584.

Fender, I., and J. Kiff. "CDO Rating Methodology: Some Thoughts on Model Risk and Its Implications." BIS Working Paper No. 163 (2004).

Fitch. "Fitch Ratings Global Corporate Finance 2007 Transition and Default Study." www.fitchratings.com (2008).

Franke, G., and J. P. Krahnen. "Default Risk Sharing between Banks and Markets: The Contribution of Collateralized Debt Obligations." In *The Risks of Financial Institutions*, M. Carey and R. M. Stulz, eds. Chicago, IL: University of Chicago Press (2006), 603–631.

Gabaix, X.; A. Krishnamurthy; and O. Vigneron. "Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market." *Journal of Finance*, 62 (2007), 557–595.

Gale, D., and M. Hellwig. "Incentive-Compatible Debt Contracts: The One-Period Problem." *Review of Economic Studies,* 52 (1985), 647–663.

Gollier, C., and H. Schlesinger. "Arrow's Theorem on the Optimality of Deductibles: A Stochastic Dominance Approach." *Economic Theory*, 7 (1996), 359–363.

Gordy, M. B. "A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules." *Journal of Financial Intermediation,* 12 (2003), 199–232.

Gorton, G. B., and G. G. Pennacchi. "Banking and Loan Sales: Marketing Nonmarketable Assets." *Journal of Monetary Economics,* 35 (1995), 389–411.

Herrmann, M.; P. Sun; M. Jha; J. Rudolph; G. Beckmann; and C. Bishko. "Global ABS 2007." HSBC Bank plc London, electronic publication (Jan. 2007).

Higgins, E. J., and J. R. Mason. "What Is the Value of Recourse to Asset-Backed Securities? A Clinical Study of Credit Card Banks." *Journal of Banking and Finance,* 28 (2004), 875–899.

International Monetary Fund. "Global Financial Stability Report." Washington (April 2011).

Krekel, M. "Pricing Distressed CDOs with Base Correlation and Stochastic Recovery." Working Paper, Uni Credit Markets & Investment Banking (2008).

Longstaff, F. A., and A. Rajan. "An Empirical Analysis of the Pricing of Collateralized Debt Obligations." *Journal of Finance,* 63 (2008), 529–563.

Loutskina, E., and P. E. Strahan. "Securitization and the Declining Impact of Bank Finance on Loan Supply: Evidence from Mortgage Originations." *Journal of Finance*, 64 (2009), 861–889.

Malamud, S.; H. Rui; and A. Whinston. "Optimal Securitization with Heterogeneous Investors." Working Paper, Swiss Finance Institute (2009).

Moody's. "Moody's Rating Actions, Reviews and Outlooks: Quarterly Update—Fourth Quarter 2007." www.moodys.com (2008).

Moody's Investor Service. "The Lognormal Method Applied to ABS Analysis. International Structured Finance." Special Report (2000).

Newey, W. K., and K. D. West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica,* 55 (1987), 703–708.

Newman, D.; F. J. Fabozzi; D. J. Lucas; and L. S. Goodman. "Empirical Evidence on CDO Performance." *Journal of Fixed Income* (2008), 32–40.

Piskorski, T.; A. Seru; and V. Vig. "Securitization and Distressed Loan Renegotiation: Evidence from the Subprime Mortgage Crisis." *Journal of Financial Economics,* 97 (2010), 369–397.

Plantin, G. "Tranching." Working Paper, University of Toulouse (2003).

Purnanandam, A. "Originate-to-Distribute Model and the Subprime Mortgage Crisis." *Review of Financial Studies*, 24 (2011), 1881–1915.

Riddiough, T. J. "Optimal Design and Governance of Asset-Backed Securities." *Journal of Financial Intermediation,* 6 (1997), 121–152.

Tarashev, N. "Measuring Portfolio Credit Risk Correctly: Why Parameter Uncertainty Matters." *Journal of Banking and Finance,* 34 (2010), 2065–2076.

The Financial Crisis Inquiry Commission Report. Available at http://www.gpo.gov/fdsys/pkg/GPO-FCIC/pdf/GPO-FCIC.pdf (2011).

Townsend, R. M. "Optimal Contracts and Competitive Markets with Costly State Verification." *Journal of Economic Theory,* 2 (1979), 265–293.