18

Another view on the pricing of MBSs, CMOs and CDOs of ABSs

Jean-David Fermanian
CREST–ENSAE

Since the 1980s, securitisation has fuelled the creation of multiple families of financial products that we will call "structures". The reasons for this boom are well known and have been widely studied in the literature: see, for example, Fabozzi and Kothari (2008) or Takavoli (2008). Apparently, most of these structures follow the same logic: gather more-or-less tradable assets together in an ad hoc vehicle, create different tranches that would provide different risk/return profiles for different classes of investors and sell them.

Let us keep in mind the two icons of this business.

- Synthetic corporate collateralised debt obligations (CDOs), based on about 100 credit default swaps: the main underlying risks are the potential default events of the underlying firms, their spread variations and the recovery amounts after any default.
- Asset-backed securities (ABSs): cash instruments based on thousands of individual loans coming from retail banking. Here, losses can occur due to prepayment (some underlying assets can be repaid more quickly than expected, inducing a marked-to-market loss for investors) or to default risk (inability to reimburse some coupons or principal fully). When the underlying loans are related to mortgages, we use the term mortgage-backed securities (MBSs).

All these structures are exposed to an interest rate risk, but mortgages face particular exposure. Indeed, lower interest rates are an incentive to prepay current loans and to enter into new ones under better financial conditions. When interest rates rise, the weight of periodic reimbursements will become heavier for weak-floating-rate borrowers, pushing them to bankruptcy more easily.

In addition to CDOs and ABSs, a lot of (more-or-less) exotic structures have appeared. They differ in the way risks are allocated between investors and as time passes. In particular, collateralised mortgage obligations (CMOs) borrow the "tranching" idea from CDOs, but apply it to mortgages. Moreover, there exists a menagerie of exotic corporate structures (forward-starting CDOs, leveraged super senior tranches, CPDOs, etc) exemplifying the imagination of structuring desks.

The children of securitisation share a common high degree of complexity and its corollary: the difficulty in assessing continuously their fair prices in any market environment and to measure and to understand their underlying risks. Moreover, their final payouts depend strongly on the tranching process that allocates risk among investors, and that generates the so-called "waterfalls" of payments. But beyond their visible similarities, such products show significant differences on the model side and in terms of practical management. In securitisation, two factors are particularly significant:

- the distinction between cash and synthetic structures;
- the number of underlyings.

Indeed, one of the key points in a cash structure is the full description of its future cashflows. Obviously, these cashflows will be random, due to possible future default events, prepayments, interest rate moves, etc. Otherwise, such a product would be risk free. But even in a fully deterministic environment, the task of cashflow calculation is painful. It often implies a significant cost in terms of IT infrastructure/systems. More seriously, the creativity of structurers and the risk appetite of clients have converged to complicate the picture. Now, a lot of structures involve non-linearities, triggers, tranching, optionalities (call rights), regular tests linked with external processes like ratings, etc. Thus, the degree of complexity can be so high that the formal pricing of any cash structure becomes a permanent challenge. Everybody knows the classical principles of derivatives valuation (to exhibit a relevant stochastic processes that would describe the underlyings and then the calculation of an expectation under a risk-neutral probability, possibly in closed form). Here, such an agenda appears inaccessible for most practitioners. This feeling has been strengthened with the latest avatars of the securitisation universe, the famous CDOs of ABSs, which provide another level of securitisation. Formally, as "CDO squared", they produce securitisations of already securitised products, inducing highly complex modelling issues. The risk management of such structures is particularly challenging. These CDOs of ABSs were under the spotlight during the financial crisis of 2007–9. The losses they suffered have offered arguments against the business of securitisation and its excesses.

Apparently, synthetic structures offer more guarantees. Indeed, their underlyings are derivatives, typically swaps, for which cashflows are simpler and relatively standardised. Then, there are some hopes the waterfalls of synthetic structures could be modelled explicitly, at least in the simplest cases. Indeed, this is often the case, particularly with corporate CDOs based on credit default swaps. For these products, the remaining challenge is the famous "correlation puzzle", or the way in which individual default risks may be correlated. This problem has generated a huge amount of literature. At the time of writing, the market standard is the so-called "Gaussian copula model" (Li 2000), an apparently static model that could be justified by some underlying credit-spread dynamics (Fermanian and Vigneron 2010). Its shortcomings are numerous and well known (Finger 2004; Lipton and Rennie 2007).[1] Its wide and rash use by practitioners is believed to be one of the sources of the 2007–9 credit crisis. Therefore, its key formula has been called "the formula that killed Wall Street" (Salmon 2009).

Actually, the "correlation puzzle" is still present for cash structures too. In the US, many mortgage-related products are implicitly guaranteed by an agency (eg, Federal National Mortgage Association, Federal Home Loan Mortgage Corporation, Government National Mortgage Association). They convey virtually no default risk, and academics have often considered credit risk as a second-order effect. As a consequence, the amount of academic literature that deals with prepayment issues is significantly larger than the stream of papers that focus on the default risk of borrowers. Unfortunately, the subprime crisis has illustrated the weaknesses of the standard risk management practices in such cash structures. After the dramatic increase of default risk in every category of US mortgages, it is no longer reasonable to leave credit risk aside. Note that default risk is far from being negligible in continental European structures too. There, prepayment incentives are a lot weaker than in the US (partly due to associated penalties) and loans are not guaranteed most of the time.

Strangely, these deep questions have been somewhat hidden by the first immediate challenge: to recover expected cashflows associated with the at-hand cash structure. In practice, engineers were a little obsessed by the management of huge amounts of loan-level information which created a smokescreen covering the major problems on the purely quantitative side: the way default and prepayment events would be dependent and would modify waterfalls. Ultimately, the main concern is to price and hedge such derivatives. But, as a consequence, and as noticed by Brigo *et al* (2010),

> [cash CDOs] are complex products with sophisticated and path-dependent payouts that are often valued with extremely simplistic models.

Therefore, two methodologies have been developed, largely independently, without any clear link between them: on the one hand, synthetic and relatively standardised synthetic structures (based on CDSs or ABCDSs, their counterparts for asset-backed securities), daily market quotes, "correlation trading" desks and financial engineers with the same background as quants on other derivatives (FX, interest rates, stocks) and, on the other hand, a less well-defined world of cash structures, mainly US and mortgage related, rather than buy-side oriented, where quantitative models are more econometric and crucially dependent on traders' inputs.

In our opinion, there are no theoretical reasons to justify such a discrepancy, merely practice and history. The fundamental modelling issues remain whatever the nature or number of the underlyings (particularly cash or synthetic). The cornerstone is dependence between all the underlying risks.

A few attempts have tried to adapt the usual models of synthetic corporate CDOs towards ABSs: Lou (2007), Garcia and Goossens (2008a,b), Hull and White (2009), among others. These authors have amended the usual one-factor Merton model to deal with ABS-type structures. Such approaches will be outlined in the next section. Since the standard "static" Merton model does not focus on the underlying dynamics, these models cannot capture the uncertainty around the timing of cashflows without a heroic simplification of reality.

Another tempting approach is related to intensity-based models. Here, by building dependent individual default-intensity processes, we get a satisfying description of correlated default events and their timings to fit the CDO market. The extension of such reduced-form models to ABS structures seems to be particularly relevant. Indeed, duration models have become the standard framework for the latter products, since the publication of some seminal papers at the end of the 1980s (Richard and Roll 1989; Schwartz and Torous 1989). Thus, individual stochastic default and prepayment rate processes (possibly correlated) could be the main tools for pricing ABSs. In addition, the associated formalism could be very close to the classical approach of stochastic intensity models, especially with use of the arsenal of affine jump-diffusion processes available (Duffie *et al* 2000). We outline such an approach later in the chapter (see page 317 onwards).

As an alternative to the previous bottom-up models, we present the top-down approach of Fermanian (2010), where total exposures, losses and prepaid amounts are correlated stochastic processes. Therefore, in a parsimonious framework and under reasonable assumptions, it is possible to get closed-form (or at least semi-analytical) pricing formulas for CMOs and even CDOs of ABSs (see page 320). In the discussion section (see page 329), we discuss the relative advantages and shortcomings of all these approaches. Finally, we describe the skeleton of a core model to price every ABS-type structure. We quote some technical hurdles and difficulties that need to be overcome before we start to live in such a "perfect world" (see page 332).

THE MERTON MODEL AND ABS-TYPE PRODUCTS

As usual in a Merton-style approach (Merton 1974), a "default risk" asset value process $(X_{it})_{t>0}$ can be associated with any underlying i in a given pool, ie, with any debt (corporate, retail, other). The default event of i is recorded when its asset value falls below a debt threshold. Similarly, a "prepayment" asset value process $(Y_{it})_{t>0}$ can be built, and prepayment events would be generated exactly like default events. When the first of these two events occurs, the corresponding name leaves the pool and its current balance is removed.

It is theoretically possible to make all these processes dependent, to state relevant debt and prepayment boundaries and generate trajectories to get prices (by simulation, most of the time). Clearly, the lack of closed-form formulas and of instruments for calibration, the number of parameters and the huge number of underlying processes prevent such a "brute force" approach. Note that this difficulty already appears with most synthetic corporate CDOs, ie, in the case of default risk only and without amortisation features. Using Zhou (2001), we can check the complexity of pricing CDOs by dealing with multidimensional asset value processes. Thus, the consideration of even a single process per name which would generate the minimum of the latent default time and of the latent prepayment time is not yet a solution that can be used in practice.

Therefore, simplifying assumptions have to be made. The first tempting assumption is related to the definition of the default/prepayment events themselves. As in the market standard "base correlation" framework, we could consider that the default (respectively, prepayment) of i occurs before T when X_{iT} is below some threshold d_{iT} (respectively, e_{iT}), for every maturity T. It is well known that such a strategy allows relatively simple closed-form pricing formulas for synthetic CDOs. Nonetheless, in general, the "correlations" are different from one time horizon to another and it is impossible to simulate all the default events consistently at several times T. In other words, we need to assume that some defaulted names can be resurrected, to be calibrated to every portfolio loss distribution at every time T. This weakness of the base correlation technique is well known (Brigo *et al* 2010). In the case of ABSs, it is particularly tricky, because the effect of time is even stronger than for (bullet) bond/CDS exposures. Moreover, the possibility of pricing by simulation is even more desirable, because it is most often the only way of integrating path-dependent features, or even to deal conveniently with amortisation schedules. Indeed, time after time, current notionals have to be calculated depending on some pre-specified amortisation schemes. Moreover, the latter can depend on many random events (triggers, tests, etc).

A shortcut would be to assume parametric profiles of amortisation and/or independence between amortization/prepayment and defaults. This is the solution proposed by Garcia and Goossens (2008a,b), for instance. Clearly, this approach is highly questionable for standard ABSs, even if it could be accepted by default for CDOs of ABS, due to the number of underlyings and the complexity of exact cashflow calculations for the latter structures.

Since there are often thousands of names in many ABS-type products, the "infinitely granular assumption" seems to be natural: the proportion of an individual current balance is negligible with respect to the total current notional, as in Hull and White (2009). In this simplified framework, closed-form pricing formulas can be obtained, but only for bullet exposures, and the previous intertemporal inconsistency still remains in general. Consequently, this practical trick does not provide the magical tool we are looking for.

To summarise, only very approximate pricing models of ABSs can be based on Merton's structural approach in practice, given the current state of the art.

STOCHASTIC INTENSITIES AND ABS-TYPE PRODUCTS

In addition to the structural approach, there exists another important stream of CDO pricing models, called "reduced-form" or "intensity-based" models. These are still bottom-up models, but, unlike Merton-style models, they concentrate only on the timing of defaults, particularly through stochastic intensities. There are no longer (relatively artificial) economic arguments that aim to explain firm behaviours as they can appear in most structural models. This difference is also standard in economics, between those authors who want to understand the fundamental underlying behaviours and those who concentrate on descriptions/predictions based on econometrics.

Bottom-up reduced-form models are now standard for pricing corporate CDOs; see Duffie and Garleanu (2001), Mortensen (2006), among others. These authors have proposed to associate a stochastic intensity process $(\lambda_{it})_{t>0}$ with every name i in a portfolio. At time t, the intensity λ_{it} can be seen as the "instantaneous" default likelihood of i, conditional on the information available at time t. In other words

$$P(\tau_i \in [t, t + \Delta t] \mid \tau_i > t, \mathcal{F}_t) \simeq \lambda_{it}\Delta t$$

where τ_i denotes i's (random) default time, and $\mathcal{F}_t$ denotes some amount of information that is available at time t (to be specified for each particular model).

Dependence between default events is obtained by building dependence between the default-intensity processes above, for instance, through factor models. As in the seminal paper by Duffie and Garleanu (2001), we can assume that there exists a common systemic (respectively, idiosyncratic) stochastic intensity $(\lambda_t)_{t>0}$ (respectively, $(\lambda^*_{it})_{t>0}$) such that

$$\lambda_{it} = a_i\lambda_t + b_i\lambda^*_{it}$$

for some positive constants a_i and b_i. In other words, conditionally on the systemic process $(\lambda_t)_{t>0}$, the laws of default (and default events themselves) are independent.

Despite the large number of parameters and practical issues in terms of calibration, the merits of these models are certain and the fit to the market is acceptable most of the time (Eckner 2009; Chapovsky *et al* 2006). A natural idea would be to extend such reduced-form models towards ABS structures. Indeed, such a framework is more in line with the standard approach of ABS pricing than Merton-style models. The econometric point of view has become common practice in pricing ABSs/MBSs. Huge amounts of historical data about prepayment and now default are available and potentially allow the estimation of sophisticated duration models (but under the physical measure). Therefore, many authors have explained the duration of mortgages directly in reduced-form approaches.[2]

Note that standard corporate CDOs pricing models are calibrated to market data most of the time, especially those based on ITraxx and CDX. For such standardised baskets, there is a one-to-one mapping between default times/recovery rates and realised cashflows. Unfortunately, this is not the case for most ABSs, and practitioners must rely on historical/econometrical procedures. For the latter, no sufficiently liquid and reliable instruments are available in the market to find (risk-neutral) estimates of the underlying model parameters. Moreover, it is very tempting to use the high level of information that is available in the pool descriptions for every loan and borrower: the Fair Isaac Corporation (FICO) score; the loan documentation level; the loan-to-value; the loan balance; the margin levels; whether the rate is fixed/floating; the US state of origination; the delinquency history, etc. Therefore, statistics and econometrics provide key tools for ABSs, when stochastic calculus is suitable for synthetic corporate CDOs, even if both frameworks are reduced form (or intensity based).

Nonetheless, the development of credit risk models and stochastic intensities approaches have influenced the academic literature on ABSs. Indeed, as Goncharov (2002) pointed out,

> [a]fter all, from a mathematical point of view nothing precludes one from interpreting prepayment as a "default" in the intensity-based approach to pricing credit risk.

Thus, it is tempting to introduce stochastic default intensities and prepayment intensities for every name in a given ABS pool. The techniques should be similar to those for corporate CDOs.

Actually, this is not entirely the case. A major difference is due to the competing risks feature of mortgages: the life of any loan is ended by default or prepayment, but only one of these two alternatives is observed. Possibly, right censoring can occur in practice, inducing another competing risk in models. When dealing with default and prepayment risks simultaneously and in the explicit framework of competing risk models, see, for example, Deng *et al* (2000) or Kau *et al* (2006).

As a consequence, the competing risk feature of MBS histories has to be taken into account, at least during the inference step. This apparently technical point changes the perspective, and has induced another significant gap between the two families of reduced-form models. The default intensities of corporate CDOs are associated explicitly with the underlying default times and there is a one-to-one relation between them. This is no longer the case with MBSs in general. Indeed, a well-known result of non-identifiability (Tsiatis 1975) states:

> Assume every individual is under the threat of two risks, whose associated underlying default times are denoted by τ and $\tilde{\tau}$. If we observe only realizations of the minimum of both times, ie, $Y = \min(\tau, \tilde{\tau})$ and of the nature of this event, ie, $\delta = \mathbf{1}(\tau \leqslant \tilde{\tau})$, then it is impossible to identify the joint law of τ and $\tilde{\tau}$.

Here, imagine that τ (respectively, $\tilde{\tau}$) is the "latent" default (respectively, prepayment) time of any borrower. In databases of loan histories, we observe only the liquidation or termination by prepayment of every loan, but never both events. Thus, in practice, either a parametric model for $(\tau, \tilde{\tau})$ is assumed and we recover identifiability, or a richer desirable semi-parametric approach is assumed. In the latter case (the more advanced one), the goal is no longer to identify stochastic intensity processes associated with an underlying default time τ and an underlying prepayment time $\tilde{\tau}$, but rather to be able to calculate quantities like

$$\lambda_{it}^{*} := P(\tau_i \in [t, t + \Delta t],\ \delta_i = k \mid \tau_i > t,\ \tilde{\tau}_i > t)\Delta t$$

and

$$\tilde{\lambda}_{it}^{*} := P(\tilde{\tau}_i \in [t, t + \Delta t],\ \delta_i = k \mid \tau_i > t,\ \tilde{\tau}_i > t)\Delta t$$

for every name i and $k \in \{0, 1\}$. The previous quantities λ_{it}^{*} and $\tilde{\lambda}_{it}^{*}$ are called case-specific rates (see, for example, Kalbfleisch and Prentice 1980). They are directly linked to data, but they no longer describe the law of (hypothetical) underlying default and prepayment times. Nevertheless, they are sufficient for inference and prediction.

Note that by putting the competing risks issue aside and trusting a risk-neutral calibration procedure it is theoretically possible to price ABS structures (even CDOs of ABSs) in the formalism of stochastic intensities, exactly as with synthetic CDOs. In particular, this was proposed in Jäckel (2008), where the notional of any loan is reduced by default and/or prepayment continuously through fractional loss rates. Both of these rates follow Gamma processes. Notably, there is independence between fractional loss intensities due to default and those due to prepayment, even if dependence across the names in the pool is obtained by common factor processes. The substance of this bottom-up model is close to the ideas we develop in the next section.

A TOP-DOWN APPROACH FOR ABSS AND CDOS OF ABSS

In contrast to “bottom-up” models that try to model every underlying individually, some authors have argued that it is not necessary to work name by name to price the most standard CDOs. Indeed, the payouts and the prices of most standard CDOs depend only on the behaviour of the aggregated loss process in the portfolio. Thus, the so-called “top-down” approach was developed for the pricing of corporate structured products (see, for example, Andersen *et al* 2008; Bennani 2005; Schönbucher 2005). These authors concentrate on the dynamics of a single continuous-time loss process. The idea is to evaluate options of different maturities that would be written on this loss. This method seems to be particularly relevant in the case of mortgage pools that put together thousands of underlying loans. The latter cannot be risk-managed individually, removing one of the main arguments in favour of "bottom-up" models. To the best of our knowledge, a "top-down" ABS pricing model was first proposed in Fermanian (2010). He built a random process around an "information summary", the mean amortisation profile. Similarly, the underlying default risk was tackled by the expected loss random process of the whole portfolio. Both processes were correlated with each other, and with the yield curve, which was considered as the main systemic driver. Independently, Lou's (2007) paper built up similar ideas, but mainly at the loan level and on terms too general to obtain analytical pricing formulas.

In this section, we describe in detail the model of Fermanian (2010). In practical terms, this approach is a lot cheaper than managing thousands of individual loan descriptions and their interdependencies inside micro-econometric models. The author assumes implicitly that the diversification in the underlying pool is large. Thus, a few macro-factors (mainly the moves of the interest rate curve) are sufficient to price and risk-manage ABS-type structures like CMOs. By exhibiting closed-form formulas, simulations are avoided. Therefore, it is a parsimonious and consistent way of dealing with the main underlying risks together. Note that, contrary to a large part of the literature that deals with conforming mortgages, default events are not neglected. In the structures we consider, they will constitute one of the main sources of risk in addition to prepayment.

Without loss of generality, we consider that the total notional amount of our pool of assets is 1. The assumed tranching process is related to several detachment points $K_0 < K_1 < \cdots < K_p$. We set $K_0 = 0$ and $K_p = 1$. At every time t, the outstanding notional of the whole portfolio will be denoted by $O(t)$ and the outstanding notional of the tranche $[0, K]$ by $O_K(t)$. Obviously, these quantities are random.

The notional amounts of these tranches can be reduced due to three different effects.

1. The "natural" amortisation process, which is deterministic for every underlying name and deduced from contractual terms. The loans are amortised from the most senior tranches to the most junior ones ("top-down" amortisation).

2. The prepayment process. This also reduces the most senior tranches first. It can be seen as a randomisation of the previous amortisation profile.
3. The default process: failure to pay remaining coupons or nominals. This concerns the most junior tranches first ("bottom-up" notional reduction).

These features cover the most frequent specification of simple ABS structures (standard CMOs, for instance). Potentially, all these effects can apply to a given tranche simultaneously, at least from a certain time on.

For the sake of simplicity, we shall consider an ABS synthetic structure here. In this case, we do not have to take care of coupon payments even if complex notional repayment schedules can occur. Typically, we can keep in mind a synthetic CDO of an ABS, whose underlyings are some CDSs on ABS tranches (the so-called ABCDSs). In this case, no initial fund is necessary to invest in such a structure. The cashflows come only from notional repayments and defaults. The main price driver here is default risk, as for the usual synthetic corporate CDOs. The case of the usual "cash" (coupon-bearing) structures like CMOs can be deduced relatively easily (Fermanian 2010).

Let $\mathrm{RA}_{t,K}$ and $\mathrm{DL}_{t,K}$ be respectively the time-t risky annuity and the default leg that are associated with the tranche $[0, K]$. The default leg is related to default event losses only. The spread associated with the tranche $[K_{j-1}, K_j]$ is denoted by $s_{t,j}$ and it satisfies, by definition

$$s_{t,j}\{\mathrm{RA}_{t,K_j} - \mathrm{RA}_{t,K_{j-1}}\} = \mathrm{DL}_{t,K_j} - \mathrm{DL}_{t,K_{j-1}} \tag{18.1}$$

for every time t and every $j = 1, \dots, p$. This identity is standard in credit derivatives, for the pricing of credit default swaps and CDOs. The first goal of our model will be to evaluate these risky annuities and these default legs. This can be done in closed form under reasonable assumptions.

Let us denote by T^* the maturity of our structure. It may be seen as the largest maturity date of all the underlyings, or as a contractual call date of the structure. By definition

$$\mathrm{RA}_{t,K} = E\bigg[\int_t^{T^*} \exp\bigg(-\int_t^s r_u \,\mathrm{d}u\bigg) O_K(s)\,\mathrm{d}s \;\bigg|\; \mathcal{F}_t\bigg] \tag{18.2}$$

and denoting by r_s the usual short interest rate process. We denote by $E_t[\cdot]$ expectations conditional on the market information $\mathcal{F}_t$ at time t

and under a risk-neutral measure Q. The filtration $(\mathcal{F}_t)$ records all the past and current relevant information concerning the description of the cashflows and the underlyings: past payments, contractual features, interest rates, recorded losses, etc. Note that our risky annuity definition is homogeneous with a duration multiplied by a notional amount.

To fix these ideas, let us denote by $A(s)$ the portfolio amortised amount at time s. Moreover, let $A_K(s)$ be the same amount but related to the tranche $[0, K]$, ie, $A_K(s) = [A(s) - (1-K)]^+$. The latter quantity is the amount of money by which the tranche $[0, K]$ has been reduced from above, due to the amortisation process only. Actually, since this tranche is also reduced potentially from below by the default events, we have $O_K(s) = [K - L(s) - A_K(s)]^+$. We have introduced $L(s)$, the loss of the whole portfolio at time s; this is simply the accumulated amount that is due to default events. The same quantity, but related to the tranche $[0, K]$, is denoted by $L_K(s)$. Note that the latter quantity depends on the outstanding notional process of this tranche, and also on the amortisation profile. This feature complicates the asset pricing formulas significantly. Moreover, note that the portfolio outstanding notional $O(s)$ is related to the other quantities by the relation $O(s) = 1 - L(s) - A(s)$.

The loss process that refers to the tranche $[0, K]$ can be rewritten as

$$L_K(s) = L(s) \cdot \mathbf{1}(L(s) \leqslant K - A_K(s))$$

when the tranche $[0, K]$ has not been fully paid down. Otherwise, the loss amount is fixed, and keeps its last value (just before this tranche has been fully paid down). Consequently, we can write the default leg of the tranche $[0, K]$ as seen at time t as follows

$$\begin{aligned} \mathrm{DL}_{t,K} &= E_t\left[\int_t^{T^*} \exp\left(-\int_t^s r_u \,\mathrm{d}u\right) L_K(\mathrm{d}s)\right] \\ &= E_t\left[\int_t^{T^*} \exp\left(-\int_t^s r_u \,\mathrm{d}u\right) \mathbf{1}(L(s) + A_K(s) \leqslant K) L(\mathrm{d}s)\right] \end{aligned} \tag{18.3}$$

In practice, we need to evaluate the latter integral with some matrix of dates $T_0 = t, T_1, \dots, T_p = T^*$. Therefore, by neglecting the accrued payments due to defaults between two successive dates, we consider that

$$\mathrm{DL}_{t,K} \simeq \sum_{i=1}^{p} E_t\biggl[\exp\biggl(-\int_t^{T_i} r_u\,\mathrm{d}u\biggr) \times \mathbf{1}(L(T_i) + A_K(T_i) \leqslant K)(L(T_i) - L(T_{i-1}))\biggr]$$

with a reasonable accuracy. To evaluate the functions $\mathrm{RA}_{t,K}$ and $\mathrm{DL}_{t,K}$ (and due to some elementary algebraic operations) it is sufficient to calculate the expectations

$$\mathcal{E}_1(s) = E_t\biggl[\exp\biggl(-\int_t^{s} r_u\,\mathrm{d}u\biggr)[K - L(s) - [A(s) - 1 + K]^+]^+\biggr] \tag{18.4}$$

$$\mathcal{E}_2(s,\bar{s}) = E_t\biggl[\exp\biggl(-\int_t^{s} r_u\,\mathrm{d}u\biggr)\mathbf{1}\{L(s) + [A(s) - 1 + K]^+ \leqslant K\}L(\bar{s})\biggr] \tag{18.5}$$

for every couple $(s, \bar{s})$, $t \leqslant \bar{s} \leqslant s \leqslant T^*$.

Thus, the evaluation of the previous expectations is sufficient to price such synthetic structures. Tractable formulas, possibly closed-form formulas, are highly desirable. Unfortunately, these expressions involve some tricky double indicator functions in general. To simplify the analysis, we could make the following assumption.

Assumption 18.1. The amortisation process and the prepayment process will reduce only the most senior tranche.

In other words, under Assumption 18.1, we consider only trajectories where the amortisation process is stopped (or the structure is repaid) before the most senior tranche is fully paid down from above. This assumption implies that $A_{K_j}(t, T_i) = 0$ for all dates $T_i \leqslant T^*$ and for all detachment points $K_j < 1$.

Typically, in such structures, the most junior tranches are very thin with respect to the most senior tranches. It is not unusual for the latter to be related to more than 90% of total initial portfolio nominals. Moreover, in practice, these structures are "called" when the amortisation process has reduced a large part of the pool (typically 90%). Therefore, Assumption 18.1 was particularly reasonable before the 2007–9 crisis and it is still relevant for most structures. Under Assumption 18.1, Fermanian (2010) has provided closed-form formulas for $\mathcal{E}_1(s)$ and $\mathcal{E}_2(s, \bar{s})$. In his paper, general semi-analytical formulas have been proved when this assumption is removed.

Note that, under Assumption 18.1, $\mathcal{E}_1(s)$ and $\mathcal{E}_2(s, \bar{s})$ are significantly simplified. For instance, for all the tranches except the most senior one, we now have

$$\mathcal{E}_1(s) = E_t\left[\exp\left(-\int_t^s r_u \, \mathrm{d}u\right)[K - L(s)]^+\right]$$

and

$$\mathcal{E}_2(s, \bar{s}) = E_t\left[\exp\left(-\int_t^s r_u \, \mathrm{d}u\right)\mathbf{1}\{L(s) \leqslant K\}L(\bar{s})\right]$$

The expectations $\mathcal{E}_1$ and $\mathcal{E}_2$ above can be deduced from the value of some options that are written on the loss process $L(\cdot)$. Basically, it is more relevant to work in terms of the (non-discounted) expected loss process itself, which is defined by

$$\mathrm{EL}(t, T) := E[L(T) \mid \mathcal{F}_t] = E_t[L(T)]$$

Indeed, this process takes into account all the forecasts in the market continuously. Therefore, it is more acceptable to set some diffusion processes on the expected losses than on the credit losses directly. Moreover, it allows us to specify a large variety of behaviours depending on the remaining time to maturity in a consistent way. The effect of time and maturities is strong for MBSs: at the beginning of the pool, default/prepayment rates are low, due to borrowers' selection processes (the so-called "seasoning effect"). Then, typically, these rates are upward sloping, and after some years they start to decrease. Indeed, the borrowers that have "bad" individual characteristics have already defaulted/prepaid, and only the relatively safest borrowers stay in the pool (the "burnout effect"). Finally, when the process L is essentially increasing, there is no such constraint on EL, which is a desirable property in terms of model specification.

Thus, these previous expectations can be rewritten as functions of the expected losses themselves by noting that $L(s) = \mathrm{EL}(s, s)$. For instance

$$E_1(s) = E_t\left[\exp\left(-\int_t^s r_u \, \mathrm{d}u\right)[K - \mathrm{EL}(s, s)]^+\right]$$

and

$$E_2(s, \bar{s}) = E_t\left[\exp\left(-\int_t^s r_u \, \mathrm{d}u\right)\mathbf{1}\{\mathrm{EL}(s, s) \leqslant K\}\, \mathrm{EL}(\bar{s}, \bar{s})\right]$$

These expectations can be calculated as functions of the spot curve $\mathrm{EL}(t,\cdot)$ and the model parameters only. From now on, we will consider only the expected loss process, which can be seen as our underlying. Similarly, we define the expected amortised amount process $A(t,T)$ by $A(t,T) := E_t[A(T)]$.

Moreover, to evaluate the functions $\mathrm{RA}_{t,K}$ and $\mathrm{DL}_{t,K}$, we have to take into account the randomness of interest rates. Indeed, it has long been observed that the price of mortgage-backed securities depends strongly on the interest rate curves (see, for example, Boudoukh *et al* 1997), at least through prepayment. But we go further. It is intuitively clear that default rates themselves depend on interest rates. All other things being equal, the higher the rates, the more numerous the default events should be. Indeed, such increases induce financial difficulties for the debtors that are involved in floating rate loans. It should be stressed that, due to central bank policies, interest rates tend to fall when default rates skyrocket during strong recessions. This could create an opposite trend. As a consequence, the sign of the correlation between interest rate and default rates moves may be ambiguous.

Now, we assume that the expected loss process is lognormal. Since $\mathrm{EL}(\cdot,T)$ is a Q-martingale by construction, its drift is zero. That is, we make the following assumptions.

Assumption 18.2. For all times t and T

$$\mathrm{EL}(\mathrm{d}t,T) = \mathrm{EL}(t,T)\sigma(t,T)\,\mathrm{d}W_t$$

Obviously, $(W_t)_{t\in[0,T]}$ is an $\mathcal{F}$-adapted Brownian motion under Q. For the moment, we assume that we know the quantities $\mathrm{EL}(t,\cdot)$ at the current time t, as they can be observed in the market. Now, let us deal with the interest rate process.

Assumption 18.3. Every discount factor $B(t,T)$ follows the dynamics

$$\frac{B(\mathrm{d}t,T)}{B(t,T)} = (\cdots)\,\mathrm{d}t + \bar{\sigma}(t,T)\,\mathrm{d}\overline{W}_t$$

for all T and $t \in [0,T]$. Here, $(\overline{W}_t)$ is an $\mathcal{F}$-adapted Brownian motion under Q, and $E[\mathrm{d}W_t \cdot \mathrm{d}\overline{W}_t] = \rho\,\mathrm{d}t$.

Similarly, we can tackle the amortisation process $A(t,T)$, which will be assumed to be lognormal.

Assumption 18.4. For all times t and T

$$A(\mathrm{d}t, T) = A(t, T)\tau(t, T)\,\mathrm{d}\overline{W}_t$$

All the previous volatility functions are assumed to be deterministic, in order to simplify calculations. But Assumption 18.4 could be weakened. Moreover, a perfect correlation between the amortisation process and the interest rate process has been assumed, essentially for convenience. This could be weakened at the cost of additional technicalities. The lognormal specification does not prohibit some strange features like $A(t, s) > 1$ or $\mathrm{EL}(t, s) > 1$ for some dates (s, t) and with some (small) probability. Theoretically, such events are unlikely. This is the price to be paid in order to have a simple specification under Assumption 18.1, and such phenomena are common in asset pricing (think about Gaussian interest rate models). Actually, this issue will be significantly reduced with reasonable parameter values and without pricing very high tranches of the capital structure.

It should come as no surprise that the expected loss process $\mathrm{EL}(\cdot, T)$ may decrease in time, for a given time horizon T. Indeed, this process is related partly to some expectations of future losses and partly to the current realised losses. Future expected losses could become smaller tomorrow if the market participants become more confident about the financial strength of the borrowers in the pool. Moreover, in the ABS world, it is even possible to recover some losses in the future. Indeed, recorded losses do not necessarily imply the closure of deals/tranches. These losses can be temporary, because they are based on some statistical models and projected cashflows. In other words, losses are not clearly identified with certainty before the formal termination of the distressed loans considered. Therefore, marked-to-model losses can be recorded one day and recovered at least partly afterwards. These temporary losses are clearly a source of difference with corporate CDOs.

Through some change-of-numéraire techniques, Fermanian (2010) has calculated the expectations $\mathcal{E}_1$ and $\mathcal{E}_2$, both under and without Assumption 18.1. In the former case, simple closed-form formulas were obtained. In the latter case, semi-analytical formulas were obtained: conditionally on the amortised amount at the relevant date, formulas of the "Assumption 18.1 case" can be applied and integrated with respect to the law of the random variable $A(s, s)$, for every s.

Now, assume the current time is t. To obtain t-prices of tranches, it remains to evaluate the t-spot expected losses $\mathrm{EL}(t,T)$ and, for the most senior tranche, the expected amortised amount $A(t,T)$. Since they are not observed directly in the market, additional assumptions concerning the shape of the t-current profile $T \mapsto E(t,T)$ have to be made. For instance, we could state that, with a constant rate $\theta_t > 0$, we have

$$\mathrm{EL}(t,\mathrm{d}T) = \theta_t \cdot E_t[O(T)]\,\mathrm{d}T$$

This constant θ_t has the status of a constant default rate, even if it is related to some expectations of losses. Note that the above relation induces a feedback of losses towards the amortisation process $A(\cdot,\cdot)$ through $O(\cdot,\cdot)$.

Formally, we could deal with the amortisation process as with the expected loss process itself. For the moment, we have merely to evaluate $A(t,T)$ knowing the information at time t. We will assume that

$$A(t,\mathrm{d}T) = [\xi_{t,T} + b_t] \cdot E_t[O(T)]\mathrm{d}T$$

where $\xi_{t,T}$ is the "theoretical" amortisation rate at time t for the T maturity, and b_t is a constant risk premium. The former quantity is the time-T rate of decrease in terms of the notional, assuming there will be no prepayments. This can be inferred from a description of the cashflows that are associated with the survival assets in the pool at time t. This rate $\xi_{t,T}$ is the global risk premium associated with both the prepayment process and the amortisation process. If $\xi_{t,\cdot}$ were a constant function, then $(\xi_{t,\cdot} + b_t)$ would be the t-amortisation rate and b_t could be seen as the so-called "constant prepayment rate" at time t. The constancy of b_t with respect to T has been made just for convenience. It is straightforward to extend the results to deal with (deterministic) term structures of prepayment rates. Since $E_t[O(T)] = 1 - A(t,T) - \mathrm{EL}(t,T)$, we deduce that

$$E_t[O(\mathrm{d}T)] = -A(t,\mathrm{d}T) - \mathrm{EL}(t,\mathrm{d}T) = -(\theta_t + b_t + \xi_{t,T})E_t[O(T)]\,\mathrm{d}T$$

and therefore

$$E_t[O(T)] = O(t)\exp\left(-(\theta_t + b_t)\cdot(T-t) - \int_t^T \xi_{t,u}\,\mathrm{d}u\right)$$

Finally, we get

$$\mathrm{EL}(t,T) = \mathrm{EL}(t,t) + O(t)\theta_t \int_t^T \exp\left(-(\theta_t + b_t)\cdot(u-t) - \int_t^u \xi_{t,v}\,\mathrm{d}v\right)\mathrm{d}u \tag{18.6}$$

Thus, at the current time t, the expected loss depends only on the "no-default, no-prepayment" amortisation profile and on some constants θ_t and b_t. We have replaced a whole unknown spot curve $\mathrm{EL}(t, \cdot)$ by a parameterisation of this curve, given by Equation 18.6. Obviously, other choices of spot EL curves are possible, even from a purely non-parametric point of view, if the data related to the pool are sufficiently rich. Note that Equation 18.6 does not contradict Assumption 18.2.

Similarly, we deduce the spot amortisation profile, ie, the curve $A(t, \cdot)$

$$A(t,T) = A(t,t) + O(t) \int_t^T (\xi_{t,u} + b_t) \times \exp\left(-(\theta_t + b_t) \cdot (u - t) - \int_t^u \xi_{t,v} \, \mathrm{d}v \right) \mathrm{d}u$$

Then we have obtained all the building blocks to price standard synthetic tranches.

DISCUSSION AND PRACTICAL ISSUES

The different approaches given above have their pros and cons. But besides theoretical arguments in favour of or against any of them, the practical implications are crucial, particularly in terms of calculation time. Indeed, any ABS bond contains at least hundreds (more often thousands) of individual loans, for which a lot of information is available: the history of past payments, the features of every loan and every borrower, etc. A first crucial decision is to know whether quants/modellers want to use all this information or just a relatively small part of it. In the former case, they will naturally build a "bottom-up" model. In the latter case, they will be attracted by "top-down" approaches.

Clearly, this decision has significant practical consequences. Actually, such a decision is related to the "infinitely granular" assumption, which was established by the famous "Basel II model" (Gordy 2003): any individual credit exposure in the pool is sufficiently small that some "law of large numbers" effect applies. In particular, all idiosyncratic behaviours are perfectly diversified, and the risk in the portfolio is driven entirely by the underlying systemic factors. Such a behaviour is likely in mortgages that include large pools of "not too dissimilar" loans. When most of the loans are of the same nature (fixed rate or ARM 2/28, for example) and when borrowers share similarities (all of them are poor quality subprime borrowers, or, alternatively, all of them are implicitly guaranteed by an agency), we accumulate arguments in favour of the "infinitely granular" assumption.

Actually, the distinction above between "bottom-up" and "top-down" approaches is not just a consequence of the potential granularity level. Indeed, in both frameworks we can find variable degrees of use of loan-per-loan static data and dynamic features: the current and past status of the borrowers, their time spent in delinquency or foreclosure states, and, more generally, all the individual loan histories. To be specific, a basic bottom-up model could not be calibrated precisely to all these features, and it will use only a few of them (for example, the default/prepayment probability conditional on the initial FICO score, or conditional on the current borrower state). It is particularly true for structural models. On the other hand, there is some space, at least in theory, for the building of very rich dynamic intensity-based processes for every loan, possibly with many explanatory variables. Apparently, a top-down model does not have to deal with numerous individual characteristics. But, in the previous section, we pointed out that these characteristics have to be managed in order to estimate the relevant mean default rates θ_t and mean prepayments rates b_t. In practice, under a historical calibration point of view, this may imply the estimation of individual default and prepayment rates followed by an aggregation stage over all the borrowers. Therefore, the calibration step of a top-down model often involves the management of loan-level information. Actually, this historical calibration procedure is the standard approach because of the lack of market instruments to calibrate top-down model parameters. Indeed, with the notable exception of ABX-type quotes, it is difficult to find sufficiently reliable quotes to lead a usual "risk-neutral" calibration with ABS structured products.

An important discrepancy between the corporate synthetic CDO world and the ABS/MBS world is hedging. Indeed, credit default swaps are available for a lot of corporate firms. They allow continuous (partial) hedging, and even perfect replication under some no-jump-to-default assumptions (Fermanian and Vigneron 2009). Even with top-down approaches, it may be possible to calculate hedge ratios (at least in theory) by the use of random thinning techniques

(Giesecke *et al* 2010). Unfortunately, micro-hedging is clearly a pipe dream for retail-based products (as it is for most ABSs), because no market instruments are available. Actually, management of loan-per-loan risk is virtually impossible with large pools (larger than several hundreds of names) in practice. This is particularly true for the securitisation of tranches, such as for CDOs of ABSs. In this case, the usual shortcut is to assume that every tranche behaves like a bond and we recover the usual CDO framework (see, for example, Jäckel 2008). It is clearly a strong assumption, because of the non-linearity of the credit risk of tranches. In theory, it would be necessary to model and price simultaneously and consistently every "first-level" tranche. But the IT and operational effort required is huge. Thus, it is not really a surprise that rating agency models and even most dealer models of CDOs of ABSs have been accused of using models/platforms that are too simplistic to manage such a level of complexity, especially during the stressed period of time we saw in 2007–9.

Another source of discrepancy is in the definition of default events themselves. Officially, a corporate default is realised when an announcement appears in a shortlist of financial newspapers. Even if the recovery process can take several weeks or even months, the default event is taken as certain and quoted recovery rates appear quickly in the market. Therefore, formally and in practice between dealers, a default event can be considered as a point in time with a slight approximation. This is clearly an advantage in terms of modelling. Unfortunately, for mortgages, it is not entirely the case. Default is progressive and there is no consensus in defining a particular point in time that would be called "default time": 30, 60 or 90 days of delinquency, or even more? What about foreclosure or real estate owned? In every case, the legal termination of a mortgage debt in trouble is the final stage of a relatively long process (several months at least), that often (but not always) finishes with the sale of a house, and during which cashflows continue due to servicers' activity. This uncertainty concerning default dates of mortgages is a source of noise and fragility in most pricing models, especially when we want them to mimic CDO models, which do not really suffer from this issue.

In the corporate world, it has been recognised for a long time that recovery rates cannot be considered to be constant over time.

Indeed, they vary significantly depending on the macro-economic environment. In particular, when the economy is plummeting and default rates surge, recovery rates tend to be smaller than their historical averages.[3] But, for a long time, corporate CDS and CDOs were priced with constant recoveries, essentially for the sake of simplicity. Recently, and as a result of the credit crisis in 2007–9, stochastic recoveries have been introduced in order to price synthetic CDOs. Particularly in the standard structural approach, recoveries are linked with the systemic factor.[4] Similar propositions have emerged in the same spirit for intensity-based CDO models.[5] This approach is clearly relevant for ABS products. In a certain sense, the link between recoveries (or severities) and some systemic factors is even stronger and more natural for ABSs than for corporate debt. Indeed, in the case of mortgages, the value of houses and their recovered values in case of liquidation are directly linked to the prices on the local housing market. Moreover, at least in the US, mortgage default likelihoods are closely related to the value of collaterals: when the market value of the collateral (the house) is larger then the value of the debt (the mortgage), the "negative equity" borrowers tend to default more frequently, all other things being equal; this is in line with rational behaviour. Therefore, when dealing with mortgages, the independence of recovery rates and default times is clearly not a reasonable assumption, when it is acceptable with corporate debt as a first approximation. Furthermore, the necessity for modelling jointly default, prepayment and severity adds another level of complexity to the pricing of ABSs and CMOs.

THE "PERFECT" WORLD

In theory, financial engineers can imagine a "perfect model" for pricing ABS-related structures, conditionally on the current historical and market information and on the existing information technology. Here we briefly propose some guidelines for pricing a CMO tranche, knowing that this programme can be extended to other ABSs including CDOs of ABSs.

(i) For every individual loan i in the pool, define a "latent" process $(X_{it})_{t>0}$ that drives i's default risk. Similarly, define a "latent" process $(Y_{it})_{t>0}$ that drives i's prepayment risk. These processes could be dealt with as classical asset values but they could also be stochastic intensities or other quantities that allow the calculation of default/prepayment durations.

(ii) Introduce dependency between all these processes. The most standard way of doing this is through factor models: split all these processes between systemic and idiosyncratic components; the former will be dependent, while the latter will be independent.

(iii) Generate a vector path of all the previous systemic factors. Draw individual trajectories of (X_{it}) and (Y_{it}) and then default and/or prepayment events.

(iv) Compute the associated projected cashflows, conditionally on the previous simulated scenario. Calculate the present value of these future cashflows.

(v) Repeat steps (iii) and (iv) multiple times. The price of the tranche is given by the average of all these present values.

Potentially, step (iii) is able to reflect all the features of any structure: changes of tranche ratings, moves on the house market, credit-trigger behaviour, etc. Obviously, such a program has to solve some technical issues. The first issue is simply the exact specification of all the above processes and their interdependency. Strangely, however, this point is not the most problematic, because an arsenal of potential models has been developed intensively by quants since the 1990s especially to price CDOs.

Calibration is more critical. Indeed, under this reduced-form point of view, financial engineers need an impressive amount of data to calibrate the thousands of underlying processes. Clearly, there are very few market quotes/prices, even if these few can provide valuable benchmarks. Thus, by default, historical calibration is mandatory, but under the implicit assumption that the borrower behaviours and the past explanatory variables will have the same effects in the future. Unfortunately, this assumption is highly questionable, particularly in the case of structural modifications in the market, as occurred in the US with the federal Home Affordable Modification Program (HAMP) during 2009–10.

Even if past history is relevant to explain the future behaviour of borrowers, the estimation of all the model coefficients is a statistical challenge. In theory, the literature on duration models (Kalbfleisch and Prentice 1980; Andersen *et al* 1993) and competing risks (Crowder 2001) is significant and we can rely on numerous estimation techniques. However, these techniques have been developed particularly in biology, where samples are small (at most a few hundred names). An even better source could be provided by the works of micro-econometricians who have studied individual strategies and behaviours on the labour market.[6] In every case, the identification and measurement of the underlying explanatory variable effects is far from straightforward. This task is complicated by the existence of unobservable latent variables, common factors, time-dependent covariates and by the treatment of right censoring and possibly left truncation.

Another significant hurdle is due to IT infrastructures and calculation times. Indeed, the computational power that is required to follow our "perfect model" program is impressive. To fix these ideas, imagine we want to price an MBS that contains 1,000 loans with maturities of around 30 years. It is usual practice to use monthly time steps, because of monthly coupon payments. Moreover, in the world of derivatives, a realistic pricing by simulation cannot usually be based on less than 1,000 random paths. Thus, we will have to calculate a borrower state and its current balance at least 360 million times, at least. Furthermore, in the case of intensity models, it is necessary to estimate default/prepayment monthly rates through the lines above, before simulating these events, ie, to estimate laws before simulating realisations under these laws. It is well known this "Monte Carlo of Monte Carlo" procedure is particularly costly. Even with the most up-to-date computers, the computational challenge is significant and only a few dealers would be able to develop a sufficiently powerful pricing platform to follow our "perfect world" agenda.

CONCLUSION

Corporate collateralised debt obligations and structures based on asset-backed securities are clearly different products. They are traded in different markets, as proved by the determinants of their spread behaviours (Vink and Thibeault 2008). Nonetheless, their significant similarities in terms of payouts justify our comparison of their respective pricing models. We have argued that these two worlds have not sufficiently interacted in the past. We have tried to explain why and to what extent some models that have been developed by the quants on the one side could be adopted by the other side. Particularly, we have detailed an original one-factor "top-down" approach for ABS structures, in the same spirit as the models proposed for synthetic corporate CDOs in recent years. Finally, the discrepancies and particularities of both worlds have been discussed, particularly from a practical point of view (eg, calibration, data sets, IT platforms). Today, in order to obtain a "pure" pricing of ABS-related products, as may be done with other underlyings (rates or stock derivatives), the hurdles are numerous even if a theoretically ideal pricing framework can be described.

1 A complete picture of the historical debate is given in Brigo *et al* (2010).

2 See Deng (1997), Kariya and Kobayashi (2000), Kariya *et al* (2002), among others.

3 See, for example, Altman *et al* (2004, 2005) and Altman (2006).

4 See Andersen and Sidenius (2004), Amraoui *et al* (2009), among others.

5 See Christensen (2005, 2007), Karoui (2007), Dobranszky (2008), etc.

6 See, for example, Lancaster (1990).

REFERENCES

Altman, E., 2006, "Default Recovery Rates and LGD in Credit Risk Modeling and Practice: An Updated Review of the Literature and Empirical Evidence", Working Paper, New York University.

Altman, E., A. Resti and A. Sironi, 2004, "Default Recovery Rates in Credit Risk Modelling: A Review of the Literature and Empirical Evidence", *Economic Notes* 33(2), pp. 183–208.

Altman, E., B. Brady, A. Resti and A. Sironi, 2005, "The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications", *Journal of Business* 78(6), pp. 2203–27.

Amraoui S., L. Cousot, S. Hitier and J.-P. Laurent, 2009, "Pricing CDOs with State Dependent Stochastic Recovery Rates", Working Paper.

Andersen L., and J. Sidenius, 2004, "Extension to the Gaussian Copula: Random Recovery and Random Factor Loadings", *Journal of Credit Risk* 1(1), pp. 29–70.

Andersen, L., V. Piterbarg and J. Sidenius, 2008, "A New Framework for Dynamic Credit Portfolio Loss Modeling", *Journal of Theoretical Applied Finance* 11, pp. 163–97.

Andersen, P. K., Ø. Borgan, R. D. Gill and N. Keiding, 1993, *Statistical Models Based on Counting Processes* (Springer).

Bennani, N., 2005, "The Forward Loss Model: A Dynamic Term Structure Approach for the Pricing of Portfolio Credit Derivatives", Working Paper, Royal Bank of Scotland.

Brigo, D., A. Pallavicini and R. Torresetti, 2010, "Credit Models and the Crisis: Or How I Learned to Stop Worrying and Love the CDOs", Working Paper, URL: http://ssrn.com/abstract=1529498.

Boudoukh, J., R. F. Whitelaw, M. Richardson and R. Stanton, 1997, "Pricing Mortgage-Backed Securities in a Multifactor Rate Environment: A Mutivariate Density Estimation Approach", *Review of Financial Studies* 10, pp. 405–46.

Chapovsky, A., A. Rennie and P. Tavares, 2006, "Stochastic Intensity Modelling for Structured Credit Exotics", Technical Report, Merrill Lynch, October.

Christensen, J. H. E., 2005, "Joint Estimation of Default and Recovery Risk: A Simulation Study", Working Paper, Department of Finance, Copenhagen Business School.

Christensen, J. H. E., 2007, "Joint Default and Recovery Risk Estimation: An Application to CDS Data", Working Paper, URL: http://www.fdic.gov/bank/.

Crowder, M., 2001, *Classical Competing Risks* (London: Chapman and Hall).

Deng, Y., 1997, "Mortgage Termination: An Empirical Hazard Model with a Stochastic Term Structure", *Journal of Real Estate and Economics* 14, pp. 309–31.

Deng, Y., J. M. Quigley and R. Van Order, 2000, "Mortgage Termination, Heterogeneity, and the Exercise of Mortgage Options", *Econometrica* 68, pp. 275–307.

Dobranszky, P., 2008, "Joint Modelling of CDS and LCDS Spreads with Correlated Default and Prepayment Intensities and with Stochastic Recovery Rate", Working Paper, URL: http://ssrn.com/abstract=1262217.

Duffie, D., and N. Garleanu, 2001, "Risk and Valuation of Collateralized Debt Obligations", *Financial Analysts Journal* 57(1), pp. 41–59.

Duffie, D., J. Pan and K. Singleton, 2000, "Transform Analysis and Asset Pricing for Affine Jump Diffusions", *Econometrica* 68, pp. 1343–76.

Eckner, T., 2009, "Computational Techniques for Basic Affine Models of Portfolio Credit", *Journal of Computational Finance* 13(1), Fall.

Fabozzi, F. J., and V. Kothari, 2008, *Introduction to Securitization* (Chichester: John Wiley & Sons).

Finger, C., 2004, "Issues in the Pricing of Synthetic CDOs", Working Paper 04-01, Riskmetrics.

Fermanian, J.-D., 2010, "A Top-Down Approach for Asset-Backed Securities: A Consistent Way of Managing Prepayment, Default and Interest Rate Risks", Working Paper, URL: http//www.ssrn.com.

Fermanian, J.-D., and O. Vigneron, 2009, "On Break-Even Correlation: The Way to Price Structured Credit Derivatives by Replication", Working Paper, URL: http://www.defaultrisk.com.

Fermanian, J.-D., and O. Vigneron, 2010, "Pricing and Hedging Basket Credit Derivatives in the Gaussian Copula", *Risk Magazine*, February, pp. 92–6.

Garcia, J., and S. Goossens, 2008a, "One Factor Models for the ABS Correlation Market: Pricing TABX Tranches", Working Paper, URL: http://ssrn.com/abstract=1274808.

Garcia, J., and S. Goossens, 2008b, "One Credit Event Models for CDO of ABS", Working Paper.

Giesecke, K, L. Goldberg and X. Ding, 2010, "A Top-Down Approach to Multi-Name Credit", Working Paper, URL: http://ssrn.com/abstract=1312275.

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

Goncharov, Y., 2002, "An Intensity-Based Approach to Valuation of Mortgage Contracts subject to Prepayment Risk", Working Paper, University of Illinois at Chicago.

Hull, J., and A. White, 2009, "The Risk of Tranches Created from Residential Mortgages", Working Paper.

Jäckel, P., 2008, "Gamma Loss and Prepayment", *Risk Magazine*, September.

Kalbfleisch, J. D., and R. L. Prentice, 1980, *The Statistical Analysis of Failure Time Data* (Chichester: John Wiley & Sons).

Kariya, T., and M. Kobayashi, 2000, "Pricing Mortgage Backed Securities (MBS): A Model Describing the Burnout Effect", *Asia–Pacific Financial Markets* 7, pp. 182–204.

Kariya, T., S. R. Pliska and F. Ushiyama, 2002, "A 3-Factor Valuation Model for Mortgage Backed Securities (MBS)", Working Paper, Kyoto Institute of Economic Research.

Karoui, L., 2007, "Modeling Defaultable Securities with Recovery Risk", Working Paper, McGill University.

Kau, J., D. Keenan and A. Smurov, 2006, "Reduced-Form Mortgage Pricing as an Alternative to Option-Pricing Models", *The Journal of Real Estate Finance Economics* 33(3), pp. 183–96.

Lancaster, T., 1990, *The Econometric Analysis of Transition Data* (Cambridge University Press).

Li, D., 2000, "On Default Correlation: A Copula Function Approach", *Journal of Fixed Income* 9(4), pp. 43–54.

Lipton, A., and A. Rennie (eds), 2007, *Credit Correlation: Life after Copulas* (World Scientific).

Lou, W., 2007, "Valuation of Subprime Mortgage-Backed Securities: A Portfolio Credit Derivatives Approach", Working Paper.

Merton, R., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates", *The Journal of Finance* 29(2), pp. 449–70.

Mortensen, A., 2006, "Semi-Analytical Valuation of Basket Credit Derivatives in Intensity-Based Models", Working Paper, URL: http://ssrn.com/abstract=663425.

Richard, F., and R. Roll, 1989, "Prepayment on Fixed-Rate Mortgage-Backed Securities" *Journal of Portfolio Management* 15, pp. 73–82.

Salmon, F., 2009, "Recipe for Disaster: The Formula That Killed Wall-Street", *Wired Magazine*, URL: http://www.wired.com/.

Schönbucher, P., 2005, "Portfolio Losses and the Term Structure of Loss Transition Rates: A New Methodology for the Pricing of Portfolio Credit Derivatives", Working Paper, ETH, Zurich.

Schwartz, E., and W. N. Torous, 1989, "Prepayment and the Valuation of Mortgage-Backed-Securities", *The Journal of Finance* 44, pp. 375–92.

Takavoli, J., 2008, *Structured Finance and Collateralized Debt Obligations: New Developments in Cash and Synthetic Securitization* (Chichester: John Wiley & Sons).

Tsiatis, A., 1975, "A Nonidentifiability Aspect of the Problem of Competing Risks", *Proceedings of the National Academy of Science of the United States of America* 72, pp. 20–2.

Vink, D., and A. E. Thibeault, 2008, "ABS, MBS and CDO Compared: An Empirical Analysis", *The Journal of Structured Finance* 14(2), pp. 27–45.

Zhou, C., 2001, "An Analysis of Default Correlations and Multiple Defaults", *The Review of Financial Studies* 14, pp. 555–76.