Default Risk and Diversification: Theory and Applications*

Robert A. Jarrow
Cornell University
David Lando
University of Copenhagen
Fan Yu
University of California, Irvine

First Draft: April 1999
Current Version: August 16, 2001

---

*This paper is an expanded version of a previous paper titled "Diversification and Default Risk: An Equivalence Theorem for Martingale and Empirical Default Intensities." We thank Jaksa Cvitanic, Darrell Duffie, Jean Jacod, Lionel Martellini, Fuvio Ortu, Fernando Zapatero, and seminar participants at USC, UC-Irvine, PIMCO, the 11th Derivative Securities Conference, the 2000 Risk Management Conference at the NYU Salomon Center, and the 1999 Frank Batten Young Scholars Conference at the College of William and Mary for helpful comments. David Lando acknowledges the partial support of the Danish Social Science Foundation. Address all correspondence to Robert A. Jarrow, Johnson Graduate School of Management, Cornell University, Ithaca, NY 14853-6201, Phone: (607)255-4729, Fax: (607)254-4590, E-mail: raj15@cornell.edu.

Default Risk and Diversification: Theory and Applications

Abstract

Recent advances in the theory of credit risk allow the use of standard term structure machinery for default risk modeling and estimation. The empirical literature in this area often interprets the drift adjustments of the default intensity's diffusion state variables as the only default risk premium. We show that this interpretation implies a restriction on the form of possible default risk premia, which can be justified through exact and approximate notions of "diversifiable default risk". The equivalence between the empirical and martingale default intensity that follows from diversifiable default risk greatly facilitates the pricing and management of credit risk. The consistency of this diversifiable risk assumption is illustrated in the context of existing studies on corporate bonds, interest rate swaps, and mortgage-backed securities.

Reduced-form models of defaultable securities, which view the default of corporate bond issuers as an unpredictable event, have become a popular tool in credit risk modeling. A key advantage of this approach is that it brings into play the machinery of classical term structure modeling techniques. This is convenient for the econometric specification of models for credit risky bonds as well as for the pricing of credit derivatives.

The strong analogy with ordinary term structure modeling - which will be briefly recalled in the next section - allows for specifications of default intensities and short rates using for example the affine term structure machinery of which the models by Cox, Ingersoll and Ross (1985) and Vasicek (1977) are the classical examples.[1] Pricing bonds and derivatives in this framework requires only the evolution of the state variables under an equivalent martingale measure. However, in order to understand the factor risk premia in bond markets (and to utilize time-series information in the empirical estimation), a joint specification of the evolution of the state variables under the "physical measure" and the equivalent martingale measure is required. The structure of these risk premia is well understood in - for example - the affine models of the term structure.

A key concern in our understanding of the corporate bond market is the form (and size) of the risk premia for default risk. Since the reduced-form approach allows us to model default risk using standard term structure machinery, it is natural to use the same structure for the risk premia of the intensity processes as we would use for the short rate process in ordinary term structure models. This choice has led to an interpretation of the drift adjustment on the state variables underlying the martingale default intensity as a "default risk premium" or "price of default risk".[2] Recent examples of this approach are the empirical works by Duffie and Singleton (1997), Duffee (1999), and Liu, Longstaff and Mandell (2001). The last authors proceed a step further along the risk premium interpretation by computing the expected returns of defaultable bonds using these drift adjustments.

We show in this paper, that this specification for the default risk premia implies a strong restriction on the set of possible risk premia. The fact that the intensity process is not just an affine function of diffusion state variables but is also the compensator of a jump process allows for a much richer class of risk premia. The critical distinction is really whether agents only price variations in the default intensity, which then must be pervasive, or they also price the default event itself.

This insight can be derived from existing works such as Back (1991) and Jarrow and Madan (1995). Through the well-known connection between the state-price density and the marginal utility of (for example) a representative investor or a single optimizing agent, it is easy to see that the

[1] For more general works on affine models, see for example Duffie and Kan (1996) and Dai and Singleton (2000).

[2] To be precise, the martingale intensity is usually assumed to depend on short rate factors. This is to capture the systematic dependence of credit spreads on the default-free term structure. The drift adjustments on these short rate factors are appropriately interpreted as interest rate risk premia. However, usually one more state variable is included for the intensity and its risk adjustment is given the interpretation of a default risk premium.

structure of the default risk premia used in the current empirical literature implies that there can be no jumps in endowments or aggregate consumption at a default date. We will return to this argument below. It is useful, however, to state more generally and explicitly what the structure of default risk premia is in reduced-form credit risk models. We do this with an explicit description of the possible risk premia using the work of Jacod and Mémin (1976). This explicit analysis gives insights essential to understanding the economic content of different risk premium specifications in default modeling. Using a conditional diversification argument similar to that used for the original APT, we demonstrate another sense (in addition to the equilibrium characterization) in which one can view the "change in drift risk premium specification" as corresponding to a notion of diversifiable default risk.

Our results indicate that for diversifiable default risk, there is an equivalence between the martingale and empirical default intensity functions. In this context, the drift change in the intensity is a sufficient description of the default risk premium. A corollary is that if one is concerned with a systematic jump event carrying a non-zero risk premium, then the drift change in the intensity as specified above is not the appropriate specification. A systematic jump risk premium will - contrary to the change of drift for diffusion state variables - imply a larger instantaneous intensity (and hence a larger spread as maturity approaches zero). It will also generate a higher volatility in the intensity process suggesting larger fluctuations in yield spreads than what can be explained from fluctuations of observed default intensities alone.

With the necessary theory in place for the structure of default risk premia in reduced form intensity models, we next turn to potential applications involving conditionally diversifiable default risk. First, if default intensities are specified as functions of observable state variables, diversifiable risk connects the empirically estimated intensity function obtained from default data with prices observed in the market. Despite the use of an empirical intensity function for pricing, we stress that this is *not* a risk-neutrality result. Indeed, we show that in the setting of diversifiable default risk it is possible to have both a downward-sloping yield curve for credit spreads assuming risk neutrality and an upward-sloping curve under the pricing measure, consistent with the empirical evidence supplied by Helwege and Turner (1999).

Second, in the other direction, if we specify default intensities as functions of latent state variables, diversifiable risk establishes a link between the martingale intensities obtained from market prices and actual default probabilities. This link is potentially useful when trying to extract risk measures such as credit VaR from observed market prices - a key concern in modern credit risk management. To illustrate this approach, we take the estimated martingale intensities and the associated drift adjustments from Duffee (1999) (using corporate bonds) and Liu, Longstaff and Mandell (2001) (using interest rate swaps) to compute the term structure of default probabilities. We show that the assumption of diversifiable default risk produces estimates that are in reasonable agreement with numbers derived from Moody's rating migration data. This is especially true for the long-end of the term structure, and, after adjustments for liquidity and tax effects, for the short-end as well. Given the limited precision in the current estimates and the absence of a careful quantification for the non-default parts of the spread, one cannot reject the hypothesis of diversifiable default risk. Clearly, this calls for further empirical analysis with attention to the general structure of default risk premia as explained in this paper.

As a final observation on potential applications, we argue that the concept of diversifiable default risk is not limited to credit risk modeling. In the pricing of mortgage-backed securities, it is common to price prepayment risk using empirically estimated prepayment functions which depend on systematic variables such as the level of interest rates. We explain this connection by examining the model of Stanton (1995).

The structure of the paper is as follows. In Section 1 we provide an intuitive illustration of different forms of default risk premia. In Section 2 we formally introduce the concept of conditionally diversifiable default risk using the framework of Lando (1994). In Section 3, we first establish an exact equivalence between empirical and martingale default intensities using equilibrium-based arguments, then prove a more general asymptotic equivalence using the limit economy specified in Section 2. In Section 4 we discuss potential applications of diversifiable default risk. We conclude with Section 5.

1 Variations in Default Risk vs. Event Risk

In the next section we will explicitly construct a reduced-form credit risk model with several issuers. However, before giving this construction, it is helpful to explain in a very simple setting the critical distinction between the two types of default risk premia that we are trying to understand.

1.1 Comparisons with Ordinary Term Structure Modeling

Consider an economy indexed by the time interval $[0, T^*]$ on which we have a short rate process r and a collection of Treasury securities. In this economy there is a single issuer of a defaultable bond which has a default time τ. This default process is assumed to have an intensity λ under P, the "physical" measure. The intensity of the default process provides the local default probability in the sense that the probability of the issuer defaulting over a small interval $(t, t + \Delta t)$ is equal to $\lambda_t \Delta t$. This intensity may depend on the short rate r.

In an arbitrage-free market, we have the existence of an equivalent martingale measure Q. Hence the price of a zero-coupon Treasury bond is given as

$$p(t, T) = E_t^Q \exp(-\int_t^T r_u du). \tag{1}$$

It is shown in Artzner and Delbaen (1995) that under Q, τ has a default intensity also, and we label this intensity $\widetilde{\lambda}$. Using this intensity, the price of a defaultable bond with maturity T and zero recovery in default is, under weak regularity conditions, given by

$$v(t,T) = E_t^Q \exp(-\int_t^T \left(r_u + \widetilde{\lambda}_u\right) du). \tag{2}$$

Lando (1994, 1998) extends this formula into pricing building blocks for contingent claims and Duffie and Singleton (1999) show that with a fractional recovery rate one obtains the same formula except that $\widetilde{\lambda}$ is interpreted as the fractional loss rate multiplied by the default intensity. The common theme here is that we have reduced the problem of pricing defaultable securities to evaluating the same expectation used in ordinary term structure modeling. This analogy becomes very compelling if we model the intensity using stochastic processes for which we know explicit solutions.

For example, consider a CIR model for the default intensity of an issuer. Using the analogy with the theory of short rate models, we specify the behavior of the P-intensity λ_t as

$$d\lambda_t = \kappa(\theta - \lambda_t)dt + \sigma\sqrt{\lambda_t}dW_t^P, \tag{3}$$

and the behavior of λ_t under the equivalent measure Q as

$$d\lambda_t = (\kappa + \nu)(\frac{\kappa\theta}{\kappa + \nu} - \lambda_t)dt + \sigma\sqrt{\lambda_t}dW_t^Q, \tag{4}$$

where the processes W^P and W^Q are Brownian motions under P and Q, respectively. κ, θ, ν, and σ are constants (chosen so that λ stays positive under both P and Q). ν is then interpreted as the risk premium for default risk. To facilitate further reference, we will refer to this as a "drift change in the intensity". For example, Duffee (1999), Duffie and Singleton (1997), and Liu, Longstaff and Mandell (2001) all have parameters playing the role of a "drift change in the intensity".

When trying to quantify risk premia in default markets, it is critical to note that this drift change in the intensity only captures the compensation of taking on default risk which arises from systematic factors changing the intensity. As we will see later, if the default event itself (the point process) carries a risk premium, then the Q-intensity could, for some positive constant μ, be equal to $\widetilde{\lambda}_t = \mu\lambda_t$, with a dynamics given by

$$d\widetilde{\lambda}_t = (\kappa + \nu)\,(\frac{\mu\kappa\theta}{\kappa + \nu} - \widetilde{\lambda}_t)dt + \sqrt{\mu}\sigma\sqrt{\tilde{\lambda}_t}dW_t^Q. \tag{5}$$

The constant μ is the risk premium needed to represent compensation for the default event itself.

In general, this multiplicative risk premium need not be constant, just as the drift change in the intensity could be time varying and random as well. The advantage of the work of Jacod and Mémin (1976) is that, in contrast with Artzner and Delbaen (1995), it provides an explicit characterization of the possible risk premia. Their results will prove particularly useful when considering the infinite economy needed to prove our diversifiability result.

To further explain this distinction between the concepts of pricing variations in default risk versus pricing the jump event itself, we finish this informal introduction by considering the following two examples.

1.2 Floating Rate Note with Step-Up Provision

Consider a firm whose P-intensity is given by λ_t. Assume a riskless rate of r_t, and assume that the firm issues a short rate note promising to pay a continuous coupon flow equal to $r_t + \lambda_t$, up to a maturity date T and a lump sum payment of 1 at maturity. This is a bond with a continuously adjusted step-up provision which adjusts the coupon to reflect the instantaneous default intensity under the "physical measure". For simplicity, we assume that there is no recovery payment in default. Consider the pricing of this claim under the different possible measure changes.

With a measure change corresponding to a drift change of the P-intensity, the dynamics for the Q-intensity has a drift adjustment, but the intensity process is the same (set μ equal to one and compare equation (5) with (4)). Using the results in Lando (1994, 1998), we see that the price of this claim is

$$v(0,T) = E_0^Q \int_0^T (r_t + \lambda_t) \exp(-\int_0^t (r_u + \lambda_u) du) dt + E_0^Q \exp(-\int_0^T (r_u + \lambda_u)\, du) = 1, \tag{6}$$

regardless of how the drift is changed in the Q-intensity. In other words the payment of the instantaneous, objective default intensity is enough to compensate for the default risk, no matter how risk-averse the agents are with respect to changes in default risk.

Assume instead that there is a risk premium for the default event of the firm. This corresponds to the Q-intensity of the jump being a different process. Assume for simplicity that this intensity at time t is equal to $\mu\lambda_t$, where one should think of the constant $\mu > 1$ if agents are risk-averse. Using the same approach as before, the claim issued by the firm has a price equal to

$$v(0,T) = E_0^Q \int_0^T (r_t + \lambda_t) \exp(-\int_0^t (r_u + \mu\lambda_u) du) dt + E_0^Q \exp(-\int_0^T (r_u + \mu\lambda_u)\, du), \tag{7}$$

which is clearly decreasing in μ. Hence the more risk-averse the agents are towards the default event, the less are they willing to pay for a claim which steps up the coupon payment by an amount equal to the physical default intensity.

1.3 Short-Term Bonds

The above is an idealized example, but the insight carries over to more standard contracts. For example, if we use the P-intensity to price a short-term bond when the true risk adjustment contains compensation for the default event, the drift corrections to the state variables will have to be very large to produce the desired level of spreads.

To illustrate this claim, we take the dynamics of the P-intensity λ under both P and Q given in equations (3) and (4), with parameters specified as: $\kappa = 0.186$, $\theta = 0.00499$, $\sigma = 0.074$, and $\nu = -0.216$. These values correspond to Duffee's estimates for the martingale intensity of a generic Aa-rated issuer. For simplicity, the short rate is assumed to be independent of the default intensity. Hence we ignore short rate related factors in Duffee's framework. We also assume that there is a compensation for the default event in the form of a constant $\mu = 1.1$. This value is taken for illustrative purposes only since the empirical literature does not provide any guidance on the "reasonableness" of such a parameter.

Based on this setup, we compute the yield spread of a one-year zero-coupon bond using equation (2), assuming of course zero recovery and $\widetilde{\lambda} = \mu\lambda$. We then consider the case where one takes μ to be equal to 1. This corresponds to the default risk premium arising solely from a "drift change in the intensity" explained above. In this case, one naturally assumes that the dynamics of the Q-intensity $\widetilde{\lambda}$ is given by

$$d\widetilde{\lambda}_t = \left(\kappa + \nu'\right)\left(\frac{\kappa\theta}{\kappa + \nu'} - \widetilde{\lambda}_t\right)dt + \sigma\sqrt{\widetilde{\lambda}_t}dW_t^Q. \tag{8}$$

However, in order to match the correct yield spread, we observe that ν' would have to be set equal to -0.44, a value much larger in magnitude than the "true" drift adjustment $\nu = -0.216$. This problem worsens when one examines bonds with shorter maturities, or when the compensation for default event risk μ becomes greater. For example, to match the spread on a six-month bond, ν' would have to be equal to -0.60. In the limit as maturity approaches zero, the required drift will have to be infinite if one is to match the spread produced assuming that $\mu > 1$.

2 An Intensity-Based Model with Conditionally Diversifiable Default Risk

We now move on to the formal construction of our model following Lando (1994). Consider an economy indexed by the time interval $[0, T^*]$. In this economy, there is a d-dimensional vector of state variables X, which we think of as the systematic risk factors.

Following the standard literature on large markets, we assume that there exists a countably infinite number of firms in this economy.[3] Each firm is subject to default risk, and the default time of firm i is τ^i. It is convenient to consider the one-jump process associated with firm i, i.e.

[3] There are two ways to work with a large economy ("large" in the sense of the number of firms). In the first approach, one constructs a sequence of finite sub-economies which in the limit becomes a large market. This allows concepts such as the absence of arbitrage to be defined in a rigorous way, for example, see Kabanov and Kramkov (1998) and Klein and Schachermayer (1997). However, it is difficult to generate analytical results with this approach because the structure of risk premia for existing assets changes with the addition of each new asset. It implies that the sub-economies are not nested within each other. Therefore, we use the alternative approach of starting directly with a large market. More discussions of our methodology and its relation with the first approach can be found in Section 3.2.

$N_t^i = 1_{\{\tau^i \le t\}}$. This process is assumed to have an intensity process λ^i, which depends on the state variables X. The precise meaning of this intensity is given below. The intuition is that at time t the probability of defaulting over a small interval $(t, t+\Delta t)$ for firm i is equal to $\lambda_t^i \Delta t$.

The notion of conditional diversifiability imposed in our model requires that conditioning on the evolution of X, the default processes are independent of each other. This captures the idea that once the systematic parts of default risk have been isolated, the residual parts represent idiosyncratic, or firm-specific shocks that are uncorrelated across firms. Examples of idiosyncratic shocks may include lawsuits, technological advances and managerial incompetence. We could instead work with a limited number of contagious defaults and the approximate APT would still apply, but we prefer to keep the story as simple as possible at the start. Generalizations are apparent.

Formally, we start with a filtered probability space $\left(\Omega, \mathcal{F}_{T^*}^X, \{\mathcal{F}_t^X\}_{t=0}^{T^*}, P^X\right)$ where $\mathcal{F}_t^X$ is the filtration generated by the process X_t. Here, the probability measure P^X is the empirical measure describing the properties of the state variables observed in the real world.

On this space there are also a countably infinite number of nonnegative processes, $\left\{\lambda_t^i, i = 1, 2, \dots\right\}$ which are predictable with respect to $\mathcal{F}_t^X$.[4] To construct the default processes, first augment the probability space with a sequence of i.i.d. unit exponential random variables $\left\{E^i, i = 1, 2, \dots\right\}$ that are independent of the process X_t. Then for each i, define a stopping time $\tau^i = \inf\left\{t : \int_0^t \lambda_u^i du \ge E^i\right\}$. The ith default process can be defined as $N_t^i = 1_{\{\tau^i \le t\}}$, which can only take two values, 0 or 1. With this construction, the compensated point process $M_t^i = N_t^i - \int_0^t \lambda_u^i du$ is a (local) martingale and hence λ^i is indeed an intensity process for N^i. The default process given above is called a Cox process, a doubly stochastic Poisson process, or a conditional Poisson process.

The uncertainty in this economy is then summarized by the filtered probability space $\left(\Omega, \mathcal{F}_{T^*}, \{\mathcal{F}_t\}_{t=0}^{T^*}, P\right)$ where the augmented filtration $\mathcal{F}_t = \mathcal{F}_t^X \vee \mathcal{G}_t^1 \vee \mathcal{G}_t^2 \vee \cdots$, and $\mathcal{G}_t^i$ is the filtration generated by the ith default process. Here the probability measure P is the extension of P^X to $\mathcal{F}_{T^*}$. Note that by construction, conditioning on the history $\{X_t\}_{t=0}^{T^*}$, the default processes are independent of each other. This independence captures the essence of conditional diversifiability.

The conditional distribution of the ith default time is (assuming no default before t)

$$P_t\left(\tau^i > s \mid \mathcal{F}_{T^*}^X\right) = \exp\left(-\int_t^s \lambda_u^i du\right), \ s \in [t, T^*], \tag{9}$$

and consequently the unconditional distribution is

$$P_t(\tau^i > s) = E_t^P\left(\exp\left(-\int_t^s \lambda_u^i du\right)\right), \ s \in [t, T^*], \tag{10}$$

where $E_t^P(\cdot)$ denotes the expectation under P conditional on $\mathcal{F}_t$. This completes the specification

[4]We will refer to the notion of predictability repeatedly. Unless explicitly stated otherwise, it is safe to think of predictable processes in our context as left continuous and adapted to the filtration generated by the state variables *and* the default processes.

of the default processes under the empirical measure. We now turn to the pricing of defaultable bonds issued by the firms modeled above.

Let the time-t price of a zero-coupon bond issued by firm i with maturity T be denoted by $v^i(t,T)$ where $0 \leq t \leq T \leq T^*$. When firm i defaults, a fraction $0 \leq \delta^i < 1$ of the face value of its bond will be payable at the maturity date of the bond. This is the "recovery of Treasury" assumption used for example in Jarrow and Turnbull (1995) and Jarrow, Lando and Turnbull (1997).[5]

In addition, in this economy there is a collection of default-free zero-coupon bonds trading, whose prices are given by $p(t,T)$. There is a money market in the economy defined through a short rate process r, which is adapted to the filtration $\mathcal{F}_t^X$ generated by the state variables. We assume that the market is complete in the filtration generated by the state variables, so that there is a unique measure Q^X equivalent to P^X on $\mathcal{F}_{T^*}^X$, which satisfies

$$p(t,T) = E_t^{Q^X}\left(\exp\left(-\int_t^T r_u du\right)\right). \tag{11}$$

We will not assume completeness of the defaultable bond market. We denote by Q an extension of Q^X to $\mathcal{F}_{T^*}$, which prices all defaultable bonds by discounted expectation:

$$v^i(t,T) = E_t^Q\left(\exp\left(-\int_t^T r_u du\right)\left(\delta^i 1_{\{\tau^i \leq T\}} + 1_{\{\tau^i > T\}}\right)\right). \tag{12}$$

The goal is to show that when markets become "large," the fact that default risk is conditionally diversifiable implies a restriction on the set of pricing measures which leaves "most" of the default intensities almost invariant, i.e. the goal is to say something about the properties of Q on the filtration generated by the state variables *and* the default processes.

We end this section with a useful formula that relates the defaultable bond price to a firm's default intensity. Suppose that under Q the default intensity of each firm is given as $\widetilde{\lambda}^i$. In this case the price of a defaultable bond issued by firm i is

$$v^i(t,T) = \delta^i p(t,T) + 1_{\{\tau^i > t\}}\left(1-\delta^i\right) E_t^Q \exp\left(-\int_t^T \left(r_u + \widetilde{\lambda}_u^i\right) du\right). \tag{13}$$

The proof of this can be found in Lando (1994).

A few subtleties are noted here about the above formula. First, as we mentioned earlier, Artzner and Delbaen (1995) show that it is no restriction to assume the existence of an intensity under an equivalent measure. Therefore, the notion of an alternative intensity $\widetilde{\lambda}_t^i$ under Q is well justified. Second, under a change of measure, in general the intensity will not stay invariant. Apart from the obvious consequence of $\lambda_t^i \neq \widetilde{\lambda}_t^i$, this has an additional implication for our model. That is, although we assume that λ_t^i is adapted to $\mathcal{F}_t^X$, there is no reason to expect that such a property will be

[5] Other recovery rate assumptions are possible, including the "recovery of market value" assumption of Duffie and Singleton (1999).

preserved for $\widetilde{\lambda}_t^i$. Indeed, Kusuoka (1999) provides examples in which the point process N_t^i is no longer conditionally Poisson given the state variables, and a similar formula for the conditional default probability under Q,

$$Q_t\left(\tau^i > s \mid \mathcal{F}_{T^*}^X\right) = \exp\left(-\int_t^s \widetilde{\lambda}_u^i du\right), \; s \in [t, T^*], \tag{14}$$

simply breaks down. Furthermore, the default processes are no longer conditionally independent given the state variables and the intensities alone may be insufficient to characterize the joint distribution of jump times. The economic intuition of this is that firms may have the sort of counterparty dependence posited in Jarrow and Yu (1999).[6] In an equilibrium setting, this counterparty dependence can be more precisely attributed to changing risk premia in view of a default event.

While our framework does permit the existence of counterparty risk premia, what the following sections show is that they can only be "local" if default risk is conditionally diversifiable and no (asymptotic) arbitrage is allowed—notions that are clarified shortly.

3 Invariant Default Intensity

We first present several arguments to illustrate our notion of conditional diversification and relate this concept to the exact invariance of the default intensity. We then prove a more general asymptotic invariance using only the absence of arbitrage and conditional diversification. For simplicity, we assume that all recovery rates are equal to zero. We also assume the existence of a pricing measure Q in the limit economy as we did in Section 2. We show below some properties of Q which must hold due to diversification or continuity.

3.1 Exact Results

First, consider the diversification argument used in Lando (1994). The following proposition can be stated.

Proposition 1 *Assume that the following conditions hold: 1) All firms are identical, with the same empirical default intensity λ_t and the same martingale default intensity $\widetilde{\lambda}_t$. 2) The short rate r_t, λ_t, and $\widetilde{\lambda}_t$ all have left-continuous sample paths. Then λ_t and $\widetilde{\lambda}_t$ are equal in the limit economy.*

Proof: See the appendix.

The idea in the proof is the following. Since all firms are identical in this setup, we can form an equally weighted portfolio of all bonds with only systematic risk (as represented by the state

[6] In this case, the default process can still be made conditionally Poisson, although the conditioning has to be with respect to an expanded filtration that includes not only the state variables, but also the default status of relevant counterparties. This is fully exposed in Jarrow and Yu (1999).

variables X_t) in the limit. We then use the symmetry assumption to conclude that the empirical and martingale intensities must be equal.

Another way to see this diversification argument in a more general setting is to let each firm issue (infinitely divisible) claims C_T^i, all bounded by a constant K and $\mathcal{F}_T^X$-measurable. The market is assumed to be complete in these claims so that they are priced by discounted expectation under Q as well.[7] Now consider the terminal payoff

$$Y_T^I = \sum_{i=1}^{I} w_I^i C_T^i 1_{\{\tau^i > T\}}. \tag{15}$$

This represents a portfolio of defaultable claims with weights given by w_I^i where $\sum_{i=1}^{I} w_I^i = 1$. The condition $\lim_{I\to\infty} \sum_{i=1}^{I} \left(w_I^i\right)^2 = 0$ is also imposed to ensure that we will have a "diversified" portfolio.

Define the $\mathcal{F}_T^X$-measurable random variable

$$S_T^I = \sum_{i=1}^{I} w_I^i \exp\left(-\int_0^T \lambda_u^i du\right) C_T^i. \tag{16}$$

Note that S_T^I is the expected value of Y_T^I conditional on the filtration $\mathcal{F}_T^X$. The following proposition shows that in the limit, all the default risk in the portfolio Y_T^I is "priced-out" at zero with the exception of that generated by the state variables X.

Proposition 2 *Suppose that the pricing functional given by Q is continuous in L^2. Then the difference in price between the claims Y_T^I and S_T^I converges to 0 as $I \to \infty$.*

Proof: See the appendix.

Under the assumptions given above, the proof of the proposition shows that the portfolio payoff converges to the $\mathcal{F}_t^X$-measurable random variable S_T^I in L^2. Therefore if the pricing functional is continuous in L^2, the prices must converge as well.[8] This means roughly that in the usual APT sense not all jump risks can have large premia. We will show in Section 3.2 that this also means that the individual bonds are priced approximately using the intensity preserving measure.

To informally see the implication of the above result on the default intensity, we use (13) to rewrite the price of the portfolio as

$$p_0\left(Y_T^I\right) = \sum_{i=1}^{I} w_I^i E^Q\left(\exp\left(-\int_0^T \left(r_u + \widetilde{\lambda}_u^i\right) du\right) C_T^i\right). \tag{17}$$

[7]This is not necessary for our results below. We make this assumption in order to generate results for non-trivial contingent claims. If instead we work with constant claims, diversification is applied to portfolios of defaultable bonds instead.

[8]Very weak assumptions on preferences can ensure the continuity of prices. Jarrow (1988) shows that in complete markets, prices are continuous if there exists an investor with a strictly increasing preference, holding an optimal portfolio. For incomplete markets, the preference of this investor also has to be continuous.

On the other hand, the claim S_T^I is default-free and has a price of

$$p_0\left(S_T^I\right)=\sum_{i=1}^{I} w_I^i E^Q\left(\exp\left(-\int_0^T\left(r_u+\lambda_u^i\right)du\right)C_T^i\right). \tag{18}$$

The fact that the two are equal in the limit for all diversified portfolios suggests the equality between the two intensities, although it is difficult to formalize the argument. From this we can also see that the best we can hope for are approximate results, as there can be a finite number of violations of invariance which still preserves the equality between (17) and (18) for well-diversified portfolios.

One way to establish the exact invariance of the default intensity formally is to assume that all individual assets (defaultable bonds) can be fully diversified.[9] This is in the spirit of the concept of insurable factor economy in Connor (1984), which leads to exact APT results. To see what this means in our context, we slightly depart from the previous framework and assume that bond price processes under the empirical measure P are exogenously specified as follows:

$$\frac{dv^i\left(t,T\right)}{v^i\left(t-,T\right)}=\left(r_t+b^i\left(t,T\right)\right)dt+a^i\left(t,T\right)dX_t-dM_t^i,\ i\in Z_{++}, \tag{19}$$

where $M_t^i=N_t^i-\int_0^t\lambda_s^i ds$ is the compensated martingale associated with N_t^i.[10] The state variables X_t is assumed to be a semimartingale. The coefficients $a^i\left(t,T\right)$ and $b^i\left(t,T\right)$ are $\mathcal{F}_t^X$-adapted predictable processes. They can be interpreted as the volatility of the state variables and their market prices of risk, respectively. The last term represents default risk, and drops by -1 at default, consistent with the previously assumed zero recovery rate.

In this limit economy, for each bond we can form a corresponding bond portfolio with the following price dynamics:

$$\frac{dq^i\left(t,T\right)}{q^i\left(t-,T\right)}=\left(r_t+b^i\left(t,T\right)\right)dt+a^i\left(t,T\right)dX_t,\ i\in Z_{++}. \tag{20}$$

This portfolio has the same exposure to systematic risk as bond i but without the idiosyncratic default risk. It is therefore well-diversified. Note that this assumption is very restrictive and is unlikely to hold in any finite economy.[11] However, it does lead to a nice invariance result.

Proposition 3 *Suppose that the assumptions in equations (19) and (20) hold. Then* $\widetilde{\lambda}_t^i$ *is equal to* λ_t^i *for all* i.

[9] This is true in the case of Proposition 1, since all bonds are identical and each can be fully diversified by holding the equally weighted bond portfolio.

[10] If bond prices are described by semimartingales, this is without any loss of generality in the cases where the underlying uncertainties are Wiener and/or Poisson. This is due to well-known representation theorems for Poisson and Brownian martingales. See Brémaud (1981, p.64) and Protter (1990, p.155) for detail. Equation (19) can be considered as a generalization of Heath, Jarrow and Morton (1992) to account for default risk.

[11] Hence it is more appropriate to interpret this as a condition imposed on the infinite economy.

Proof: See the appendix.

Under the assumption that there exists a pricing measure Q in the limit economy, both the bond price v^i and the price of the corresponding diversified portfolio, q^i, are Q-martingales after discounting. This implies that M_t^i is also a Q-martingale. The equality between the intensities follows.

We also present a utility-based argument for the invariance of the default intensity.

Proposition 4 *Let $Z_t = E_t^P(dQ/dP)$ be the density process of Q with respect to P. If Z_t is adapted to $\mathcal{F}_t^X$, then $\widetilde{\lambda}_t^i$ is equal to λ_t^i for all i.*

Proof: See the appendix.

Note that the density process required to change the drift of a diffusion intensity is a martingale adapted to the filtration generated by the Brownian motion underlying the intensity. Hence it is a continuous process. Recall, that under technical conditions presented for example in Back (1991), the marginal utility u_c of consumption for each investor in a CCAPM setting is proportional to the state price density. That is, there exists for each investor i constants α_i such that

$$Z_t \exp(-\int_0^t r_u \, du) = \alpha_i u_c^i(t, c^i(t)) \tag{21}$$

holds for the optimal consumption choice $c^i(t)$. This proposition shows that in order to get an exact result, one would require that the optimal consumption and hence marginal utility of agents be independent of the jump event (or that at least one agent have such a preference). Only then can the right hand side of (21) be continuous. This is closely related to Proposition 3 since in an equilibrium model with risk-averse agents, the condition in Proposition 4 is satisfied as long as the individual assets can be fully diversified. The measurability condition given above is mentioned by Jarrow and Madan (1995) in their discussion of diversifiable jump risk in term structure modeling. Its implication for an intensity-based model of default risk is presented here.

3.2 Asymptotic Results

The previous subsection establishes the exact equivalence between martingale and empirical default intensities. However, the conditions leading to this result are quite restrictive. We had to use one of the following: 1) the symmetry assumption that all firms are identical; 2) individual assets that are fully diversifiable; 3) marginal rates of substitution that are independent of pure jump risk.

In this subsection we give a more precise description of the sense in which the empirical and martingale intensities must be approximately the same for "most" assets. A number of papers have addressed the issue of asymptotic arbitrage in dynamic models rigorously, including Kabanov and

Kramkov (1998), Björk and Näslund (1998) and Klein and Schachermayer (1997). Our approach is similar to that of Björk and Näslund (1998) in that we work directly on the economy with infinitely many assets. However, we allow for more general dynamics. We will relate our results to the other two papers below.

Consider the economy formally defined in Section 2 and assume that there is a single state variable given as an Ito process:

$$dX_t = \mu(t)\,dt + \sigma(t)\,dW_t, \tag{22}$$

where W_t is a Wiener process under P, and $\mu(t)$ and $\sigma(t)$ are stochastic processes adapted to the filtration generated by W and regular enough to ensure a unique strong solution. It is trivial to generalize the following analysis to a multivariate setting. We specialize X_t to an Ito process in order to simplify the presentation of our main results below, but the argument works for jump-diffusions and more general classes of semimartingales as well.

As in Section 2, we define a pricing measure for the economy to be a measure Q equivalent to P such that under Q all bonds are priced as discounted expected values, as in equations (11) and (12). In a finite economy, the existence of such a measure precludes arbitrage. In the setup here with an infinite collection of assets it clearly excludes arbitrage in any finite sub-economy, but as we shall see, it also rules out asymptotic arbitrage as defined in Kabanov and Kramkov (1998).

Our assumption of a complete and arbitrage-free market in claims depending only on X implies the existence of an $\mathcal{F}_t^X$-predictable process g such that

$$dX_t = (\mu(t) + g(t)\,\sigma(t))\,dt + \sigma(t)\,d\widetilde{W}_t, \tag{23}$$

where $\widetilde{W}_t = W_t - \int_0^t g(u)\,du$ is a Wiener process under Q. The process $g(t)$ is assumed to satisfy

$$E^P \exp\left(\int_0^{T^*} g(u)\,dW_u - \frac{1}{2}\int_0^{T^*} g^2(u)\,du\right) = 1. \tag{24}$$

We are concerned with the form of the intensities under an equivalent measure Q. We make no assumption of a complete market for defaultable claims, but the presence of infinitely many assets still imposes an "asymptotic" structure on the intensities under Q as we shall see.

To characterize the ways in which the intensities can be modified under the equivalent measure, we need the concept of a predictable function. The predictable field on $\Omega \times [0, T^*]$ is the field $\mathcal{P}$ generated by sets of the form $A \times \{0\}$ with $A \in \mathcal{F}_0$ and $A \times (s, t]$ with $A \in \mathcal{F}_s$. A function $Y : [0, T^*] \times \Omega \times \mathbb{N} \to \mathbb{R}$ is called a predictable function if it is measurable with respect to the sigma field $\mathcal{P} \times \mathcal{E}$ where $\mathcal{E}$ is the set of all subsets of positive integers $\mathbb{N}$. We are now able to state our result which is an application of Jacod and Mémin (1976).

In the setup of our economy, we have the following:

Proposition 5 *Under an equivalent measure Q, the intensities of the one-jump processes are given as*

$$\widetilde{\lambda}_t^i = Y(t, \omega, i)\, \lambda_t^i, \tag{25}$$

for a strictly positive, predictable function Y which satisfies

$$\int_0^{T^*} \sum_{i=1}^{\infty} \left(1 - \sqrt{Y(u, \omega, i)}\right)^2 \lambda_u^i du < \infty. \tag{26}$$

Proof: See the appendix.

One should think of condition (26) as follows. If the intensities λ^i are uniformly bounded away from zero by a positive constant, then for the above condition to hold, only a finite number of default processes can have martingale intensities that deviate by more than a factor of $1 + \epsilon$ from the empirical intensities. In a finite sub-economy there can be perturbations in risk premia due to defaults of other firms [as in Jarrow and Yu (1999) and Kusuoka (1998)] since the process $Y(\cdot, \cdot, i)$ may depend on the jump times of firms other than the ith. However, such a counterparty dependence must "die out" asymptotically in the infinite economy if default risk is conditionally diversifiable under the empirical measure.

The conditional diversification construction is used for two purposes. First, the conditional independence assumption facilitates the construction of jump processes which are driven by exogenous state variables. Second, in our proof it ensures that there are no simultaneous jumps under P of the infinite collection of jump times. While the first element facilitates empirical work, it is not a necessary ingredient for Proposition 5. The second element, however, is critical. To assess and illustrate their proper roles in the above result, consider the following three examples.

First, consider the simple example of a default process N^1 which can trigger the simultaneous default of every other firm in the economy. This is the most severe incidence of "contagious default". In this case, we could have a change in the default intensity of N^1 by a factor different from 1 and this would still be consistent with the equivalence of measures. This in turn would imply that an infinite collection of default intensities deviate under Q by a fixed factor from their values under P. Hence we see that when the no simultaneous jumps condition is violated, the asymptotic equivalence result does not hold.

Second, consider the milder case of a default process N^1 which can trigger a finite change in the default intensity of other firms. In this case, the no simultaneous jumps condition could still hold. In terms of conditional independence, we can think of N^1 as representing a pervasive default risk factor and this factor would then naturally have to be included as one of the state variables with its own risk premium. If all (or many) other default processes depend on this factor under the measure P, one could then ask whether they satisfy the conditional diversification condition given the expanded set of state variables. In any case, asymptotic equivalence can be preserved for this economy.

Finally, take the counterparty default risk framework considered in Jarrow and Yu (1999). This is an economy in which defaults are correlated due to a complex structure of counterparty relations. Assuming that default processes can be constructed (recursively) and intensities are well-defined, the no simultaneous jumps condition will hold and the implication of Proposition 5 will still be valid. However, the convenience of working with conditional independence is lost as demonstrated by the looping default example in Jarrow and Yu (1999).

We have chosen to work with the pricing measure directly on a space with infinitely many firms. The definition of asymptotic arbitrage proposed in Kabanov and Kramkov (1998) uses sub-economies constructed on a sequence of filtered probability spaces to define notions of asymptotic arbitrage. In "asymptotic arbitrage of the first kind," a sequence of trading strategies is constructed such that the initial cost of the strategies approaches zero while the gains process is always non-negative and in fact becomes strictly greater than 1 at the terminal date with positive probability.[12] In their setting, the absence of asymptotic arbitrage is then linked to the notion of *contiguity* of a sequence of measures. They show that if each sub-economy has a unique equivalent martingale measure Q^n, then the absence of asymptotic arbitrage of the first kind is equivalent to the condition that the sequence (P^n) of empirical measures be contiguous to the sequence (Q^n). Klein and Schachermayer (1997) extend this result to the incomplete market case where Q^n is not necessarily unique.

Note that in the general setting, there is not necessarily any connection between the individual sub-economies in the sequence. Indeed, they can be constructed on different probability spaces altogether. However, working with this construction in our setting does not produce intuitive results unless very special structures are imposed, since the structure of risk premia in the $(n+1)$th economy can be completely unrelated to that of the nth economy. Nevertheless, one could link our construction with the theory of asymptotic arbitrage by including the first n default processes in the nth economy and assuming that there is a large N with the following property: for $n > N$, the predictable function Y^{n+1} defined on $\Omega \times [0, T^*] \times \{1, 2, \ldots, n+1\}$ is equal to Y^n on the restriction to $\Omega \times [0, T^*] \times \{1, 2, \ldots, n\}$. Using this construction, we are able to view element n of the sequence of sub-economies as a restriction of the infinite economy used in our proof to the economy generated by the first n default processes and the state variable process. This simplifies the proof of no asymptotic arbitrage considerably, since the critical condition of contiguity merely becomes a condition of absolute continuity of the unrestricted measures. We are then back to the condition in Proposition 5.

[12] The significance of the value being strictly greater than "one" is unimportant. Due to rescaling the terminal value by an arbitrary constant, the essence of this condition is that the value is bounded above zero by a strictly positive, albeit small constant.

4 Applications

Turning to applications of diversifiable default risk, we focus on the modeling of corporate bonds. We use numerical examples and estimates given in the recent empirical literature to illustrate the empirical consistency of conditionally diversifiable default risk. Other applications including mortgage-backed securities and credit derivatives are also discussed.

4.1 Corporate Bonds

4.1.1 From Empirical Intensity to Bond Prices

The equivalence result established in the previous section creates a link between the empirically estimated intensity from historical default data and the prices of defaultable securities. This is the first important application of diversifiable default risk. This underscores, for example, the pricing of the short rate note with a step-up provision discussed informally in Section 1. We use the framework of Duffee (1999) to illustrate this procedure.

In Duffee (1999), the martingale default intensity is assumed to be

$$h_t = \alpha + h_t^* + \beta_1 s_{1t} + \beta_2 s_{2t} \tag{27}$$

where α, β_1, and β_2 are constants, s_{1t} and s_{2t} are factors driving the short rate (the former is related to the term structure slope and the latter to its level), and h_t^* is a square-root diffusion

$$dh_t^* = \kappa\left(\theta - h_t^*\right)dt + \sigma\sqrt{h_t^*}dZ_t \tag{28}$$

with κ, θ, and σ constants and Z_t a Wiener process under the physical measure P. To complete the specification, Duffee assumes that under the equivalent measure Q, the process for h_t^* is

$$dh_t^* = \left(\kappa\theta - (\kappa + \lambda)\right)h_t^* dt + \sigma\sqrt{h_t^*}d\widetilde{Z}_t, \tag{29}$$

where λ is a constant and $\widetilde{Z}_t = Z_t + \int_0^t \frac{\lambda}{\sigma}\sqrt{h_u^*}du$ is a Wiener process under Q.

The constant parameter λ is then given the interpretation of a default risk premium. The analysis in the preceding section shows that this is a case of diversifiable default risk where the martingale intensity h_t can be interpreted as the empirical intensity. Along with the risk premia for the term structure factors s_{1t} and s_{2t}, λ determines the way the market prices systematic variations in the default intensity. Viewed in this light, the factor h_t^* that Duffee (1999) alluded to as a firm-specific variation of default risk probably has its origin in some common risk factor.[13] For example,

[13]The way Duffee (1999) conducts the estimation (firm by firm) suggests that the h_t^* terms are assumed to be independent across firms. While a robustness check is not done to check whether this is indeed the case, we suspect (based on current empirical evidence on credit spreads) that they contain a large common component if, say, we apply a principal component decomposition to the estimated values of h_t^*. Another way to see that it is unreasonable to assume independence is through a diversification argument - with a portfolio consisting of 100 bonds with identical credit quality, for example, the variance of the spread attributed to h_t^* has to decrease 100-fold. This would imply an incredibly large reward-to-risk ratio for the portfolio.

Elton et. al. (2001) show that a large part of credit spreads can be explained by factors considered systematic in stock markets. Similarly, Pedrosa and Roll (1998) and Collin-Dufresne, Goldstein and Martin (2001) show that movements in credit spreads have a dominant common source related to a "market spread factor" that perhaps proxies for credit market conditions. In the examples below we follow a "naive" interpretation that h_t^* and its risk premium arise from a linear relation with this factor.

The practical implication of having conditionally diversifiable default risk is now easily illustrated. Given historical data on default rates, Treasury yields, and (say) the spread index of Aa to Treasury, one may estimate, using well-established procedures in survival analysis, an affine empirical default intensity with these macroeconomic factors as time-varying covariates. We can then price corporate bonds using this estimated affine intensity function along with information on the factor evolution under the equivalent martingale measure, which can be obtained from the prices of Treasury securities and (for example) interest rate swaps.

We conclude this discussion with an example emphasizing the importance of the risk adjustments for the common factors determining the empirical intensity. In a world with diversifiable default risk and in the absence of other market imperfections for the spread, the drift change in the intensity is the sole reason why there is a gap between empirical and martingale (implied) default probabilities. This default risk premium could also provide an explanation for the well-documented empirical feature that lower-grade issuers have downward-sloping conditional default probabilities while their yield spreads may be upward-sloping.

We consider the example of a generic Baa-rated issuer with an empirical intensity given in Table 4 of Duffee (1999).[14] The parameters specified according to equations (27)-(29) are: $\alpha = 0.00961$, $\beta_1 = -0.171$, $\beta_2 = -0.006$, $\kappa = 0.212$, $\theta = 0.00628$, $\sigma = 0.059$. We take the risk adjustment parameter for the h_t^* process as a free parameter in a range bracketing the estimated value of $\lambda = -0.307$. The parameters for the two short rate factors can be found in Duffee's Table 2. Furthermore, we assume that the current values for the short rate factors are set to their long-run mean values, and the current value for h_t^* is set equal to the mean fitted value of Duffee's sample, 0.00864.

Figure 1 shows the term structure of martingale default probabilities for maturities up to 30 years. The $\lambda = 0$ series represents a case with no risk adjustment (risk-neutrality) on the h_t^* factor. Since the effect on default risk of the risk adjustments on the interest rate factors is small, this series can be interpreted as representing the empirical default probabilities. The $\lambda = -0.307$ series is the actual one estimated from data, and the $\lambda = -0.5$ series is a case with roughly a one standard deviation change in the risk premium parameter. It is apparent that the risk adjustment produces

[14]Note that the estimates in Table 4 of Duffee (1999) are given for the martingale intensity. However, under the maintained diversifiability assumption, they also are the parameters for the empirical intensity.

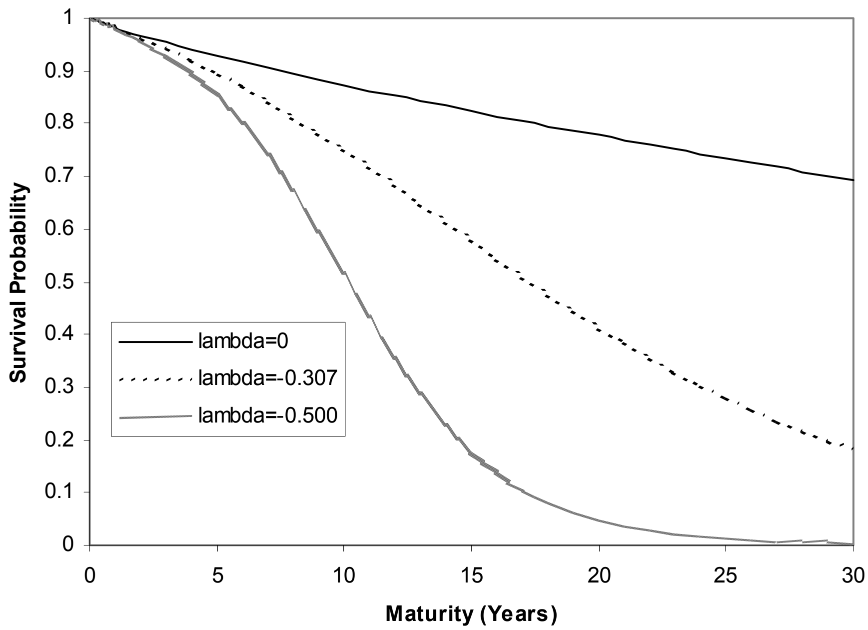

Figure 1: Term Structure of Martingale Default Probabilities for Different Default Risk Premium. The $\lambda = 0$ series represents empirical default probabilities. The $\lambda = -0.307$ series is the actual one estimated by Duffee (1999). The $\lambda = -0.5$ series represents approximately a one standard deviation change in the default risk premium parameter.

large differences in long-term actual and implied default probabilities. For the extreme short-end there are no such differences, but this is due to the structure of the intensity that we assume. This is discussed further in Section 4.1.2.

What is perhaps a more profound insight can be gleaned from Figure 2, in which we plot the term structure of yield spreads given different values of λ. Following Duffee's assumptions, we use a constant recovery rate $\delta = 0.44$ with equation (13) to compute the spreads. In the case of risk neutrality ($\lambda = 0$), we have a downward-sloping yield spread curve consistent with the fact that given survival up to time t, the conditional probability of default is decreasing as t becomes larger.[15] This is consistent with the pattern observed in Jarrow, Lando and Turnbull (1997) for lower-rated firms under a risk neutrality assumption. However, Figure 2 also shows that for the other two cases of non-trivial default risk premium, we obtain either an upward-sloping or a hump-shaped yield spread curve, consistent with the evidence reported in Helwege and Turner (1999).[16] While it is

[15]Disregarding the short rate factors for the moment, since the initial value of h_t^* is higher than its long-run mean, under the physical measure spreads will become narrower over time. Thus under risk-neutrality we would obtain a downward sloping credit spread curve for Baa issuers according to Duffee's estimates.

[16]The reason for this, mathematically, is that given our parameter values the martingale intensity becomes an explosive process after the drift adjustment. The interpretation is that investors seem to consider the conditional default probability as increasing over time whereas it actually has exactly the opposite behavior. This feature, which manifests itself in a negative but close to zero mean reversion parameter for the martingale intensity, is confirmed

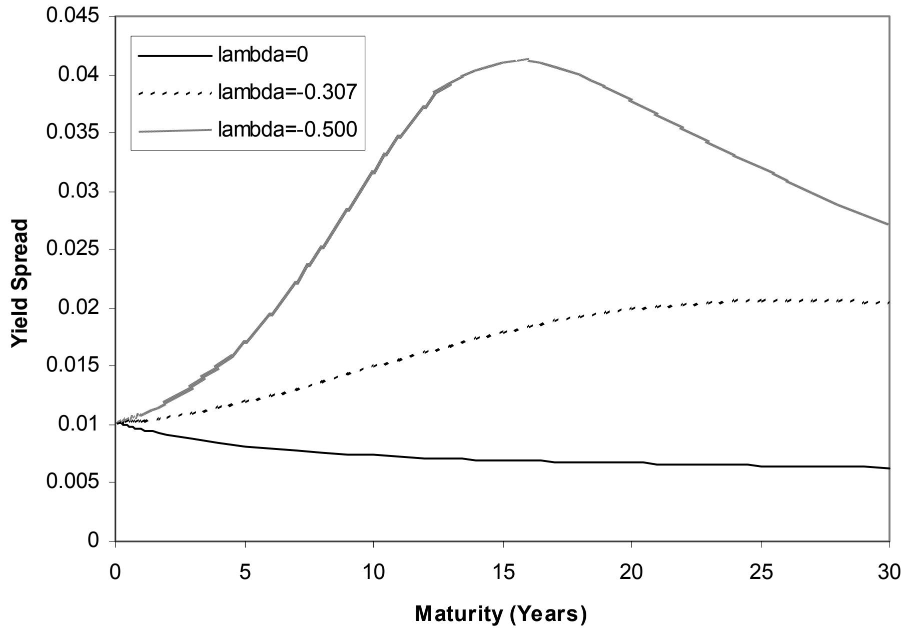

Figure 2: Term Structure of Yield Spreads for Different Default Risk Premium. The default risk premium parameter λ is specified as in Figure 1.

still a controversial issue whether the curve for lower rated issuers is truly upward-sloping, what we see here is that the existing evidence is consistent with the assumption of diversifiable default risk.

4.1.2 From Martingale Intensity to Credit VaR

The second important application of diversifiable default risk is the link between the martingale intensity estimated from market prices and the credit VaR measures needed in risk management. To illustrate this computation, we again compute the term structure of default probabilities based on the estimates from Duffee, but this time for different rating classes (Aa, A, and Baa).[17] The interpretation is that the estimated martingale intensity is equivalent to the empirical intensity, which we then integrate over time to obtain the default probabilities. The conditional default probability q_n that we use is the one-year default rate given that the issuer has survived for the first $n-1$ years. If the survival probability for the first n years is p_n, the conditional default probability is then equal to $1 - p_n/p_{n-1}$.

Before the presentation of the results, it is beneficial to discuss features of the data and the estimation procedure that may potentially limit the usefulness of such an exercise. This will assist

by other studies such as Liu, Longstaff and Mandell (2001). In terms of data, this is dictated by the need to fit a gradually increasing yield spread curve for investment-grade issuers.

[17]We drop the Aaa estimates since they seem to suffer instability problems and are unable to match the triple-A credit spread curves according to Duffee (1999).

in the interpretation of our results.

First, the martingale intensity is estimated using corporate bonds with maturities typically in the mid-range. For Duffee's sample the median of the mean number of years to maturity for fitted bonds is 7.22. This suggests that our computed default probabilities will be most accurate within this range as well. It also suggests that our computed default probabilities for the very short-term will have to be based on extrapolations from the empirical data and that their accuracy is doubtful.

Second, we expect that non-default related reasons for the spread will have a proportionally greater impact for higher quality debt and for the short-end of the term structure. These may include liquidity differences between on-the-run and off-the-run Treasury securities, liquidity and tax differentials across Treasury and corporate securities, a "non-transparency spread" due to Duffie and Lando (2001), and so on.[18] Unfortunately, there is very limited empirical evidence on the size of these market imperfections and this reduces our ability to interpret the short-end, of course.

Third, we note that with diversifiable default risk and affine diffusion state variables, spreads must be exactly equal to the empirical expected loss rate in the short-end since the drift adjustments do not have enough time to take effect. This is the sense in which observed spreads are "too large" for the empirical hazard rates of default. Instead of assuming a risk premium for the default event - which implies the breakdown of the diversification assumption - an alternative solution to this puzzle that also preserves the invariance of the intensity is to have state variables with jumps. Conditional on these (affine) jump diffusions, default risks could still be diversifiable and no risk premium for the default event needs to be assumed. Obviously, a detailed empirical analysis using both prices and default data is the only way to determine the "correct" framework.

There are several other considerations that might affect the interpretation of our results. Duffee's estimation, for example, tracks the time-series of bond prices for a median period of 96 months. Within this period, the credit rating of some issuers (investment-grade) surely will have declined. Since Duffee shows that his intensity parameters change systematically across rating categories, the term structure of default probabilities computed from his estimates will be upward biased even after accounting for the non-default part of the spread (compared to, say, Moody's average default rates).[19] Another reason why computed probabilities may be inflated is that recovery rates are assumed to be constant in Duffee's estimation, whereas empirical evidence shows that they are procyclical. Under the pricing measure, recovery rates may have a lower mean value than the historical estimates given by Moody's. Dividing by the historical loss rate will then bias the estimated

[18]Evaluated as an impact on spreads, tax differentials may have a constant effect and liquidity differences will have a declining effect as maturity increases (if, say, liquidity generates a proportional discount on the price of the bond). Duffie and Lando (2001) show, in a structural framework, that accounting imprecision may generate a non-zero short-term spread while having negligible impact for long-term spreads.

[19]In other words, the empirical default probabilities thus computed will be unbiased if the intensity parameters stays the same across ratings, i.e. if rating changes are entirely captured through changes in the intensity's state variables.

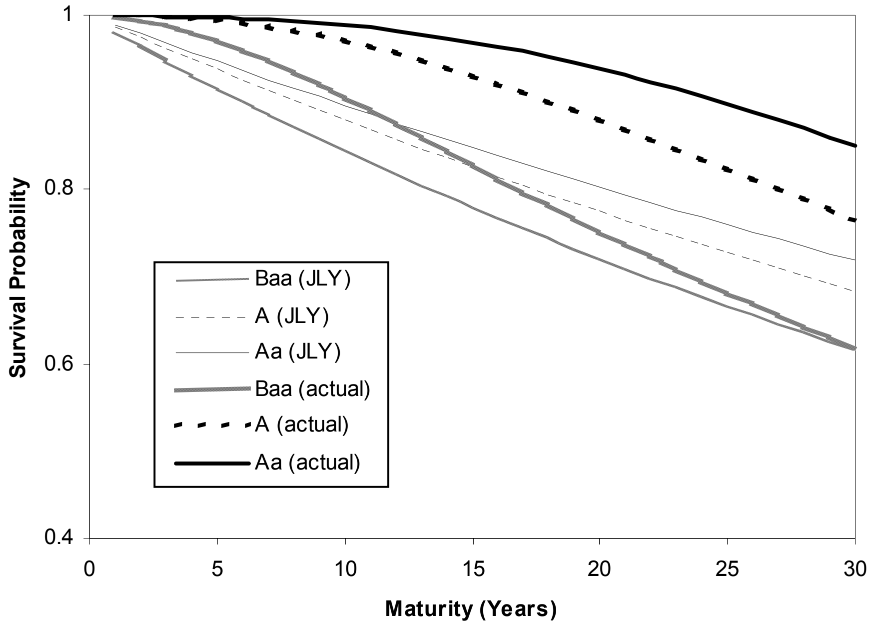

Figure 3: Term Structure of Survival Probabilities for Investment Grade Issuers. The "JLY" series are computed using Duffee's estimates. The "actual" series are computed from Moody's data.

intensity upwards.

In Figure 3, we contrast the term structures of survival probabilities for investment-grade issuers based on Duffee's estimates (the "JLY" series) with those calculated from Moody's data (the "actual" series).[20] We note several distinctive patterns. First, Moody's estimates are consistently higher than those implied from prices and the notion of diversifiable default risk. Second, the differences are more pronounced at the short-end of the term structure and decrease in relative terms as one moves to longer maturities. Third, our method seems to perform better for lower-rated issuers. The first observation is consistent with the presence of non-default reasons for the spread, and the last two with the behavior of liquidity and non-transparency premia as noted above. These differences are likely to remain in more rigorous empirical work.

Figure 4 plots the term structure of conditional default probabilities. Here we can also see the same issues. First, at the short-end our estimates are not as close to zero as are Moody's estimates. In order to fit the level of the spread with the given state variable exposure, the estimation forces a high value for the constant α in the default intensity, which then acts as a lower bound (ignoring short-rate factors for the moment) for the conditional default probabilities. This is indicative of the non-default parts of the spread contaminating the inferred default probabilities. In lieu of a

[20] We use the average one-year transition matrix between 1980 to 1999 to compute the survival probabilities as a function of maturity.

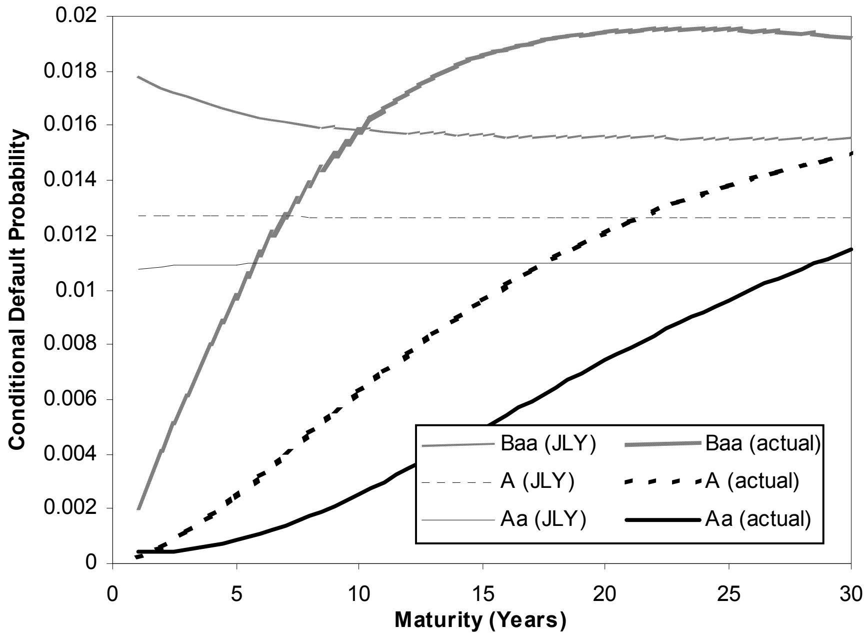

Figure 4: Term Structure of Conditional Default Probabilities for Investment Grade Issuers. The data series are explained as in Figure 3.

careful empirical study, anecdotal evidence suggests that short-term corporate spreads may not be too large to be justified by the empirical default rates. For example, Duffee (1999) shows that for Aa-rated issuers the short-term spread is about 60 bps. Of this, Elton et. al. (2000) estimate that state taxes may account for as much as 30 to 50 bps. The specialness premium (between on- and off-the-run Treasury securities) is on the order of 10 bps. Add to these the potential corporate bond liquidity and non-transparency premia, it is not clear that the actual and implied conditional default probabilities differ at the short-end.

A second issue is that unlike the "actual" series that are upward-sloping throughout the range of maturities, our "JLY" series are quite flat, eventually matching the "actual" series at the mid- to long-end. This flatness is partly due to the fact that we set the initial value of the h_t^* process equal to its mean fitted value over Duffee's sample period, which is quite close to the long-run mean value θ. However, even if we set the initial value close to zero, the upward-sloping shape cannot be reproduced. This is because under the empirical measure the state variables mean-revert "too quickly" - for example, for an Aa issuer the half life of the h_t^* factor is slightly less than 4 years. Accounting for the term structure of the non-default parts of the spread may help to resolve this issue as well.

The third issue is that Figure 4 is based on estimates with very high standard errors. For example, the α estimates alone have standard errors on the order of 30 bps (judging from the quartile figures in Duffee's Table 3). The mean-reversion parameter and the long-run mean values have similar if not larger estimation errors. These sampling errors suggest that it will be difficult to reject the diversifiability assumption.

These same issues can also be examined in the context of interest rate swaps. A priori, we know that the swap market has higher liquidity than the corporate bond market, and swap-Treasury spreads are less affected by tax considerations. Thus we expect liquidity and tax premia to play a lesser role. Furthermore, since the floating side of the swap is indexed to a refreshed LIBOR quality issuer (Aa or better), the upward bias associated with Duffee's estimation should instead become a downward bias. Consequently, our methodology should produce default probability estimates that are closer to the "true values" obtained from Moody's data.

We use Liu, Longstaff and Mandell (2001, LLM hereafter) to investigate these conjectures. In their paper, the loss rate is modeled as a four-factor affine diffusion which is estimated from Treasury and interest rate swap yields. Assuming diversifiable default risk, the empirical expected loss rate is

$$\lambda_t = \gamma r_t + Z_t, \tag{30}$$

where the short rate $r_t = W_t + X_t + Y_t$. The four factors, W, X, Y, and Z are independent Vasicek processes under the empirical measure.

Similar to the previous example, we compute survival probabilities and conditional default probabilities and compare them with Moody's numbers. To convert the loss rate to a default intensity, we use a constant recovery rate of 44 percent. The results are shown in Figure 5. Evidently, the level of conditional default probability is much lower than in Figure 4. However, despite a reduction in the effect of liquidity and tax differentials, a gap in the short-end of more than 60 bps is still unaccounted for.

While the patterns here are otherwise similar to those produced using Duffee (1999), there is a crucial difference regarding the structure of the default risk premium. Since LIBOR rates are polled from major banks and then averaged (dropping the outliers first), it does not represent a single Aa-rated issuer but rather the consensus of a collection of large, Aa-rated banks. If we think of this as a "representative" issuer, then the default of such an issuer is likely to have an economy-wide impact and thus would command its own risk premium. This could help to explain the short-end gap documented above. It would also imply that the risk premium structure of LLM may be misspecified and their subsequent calculation of the expected return is biased.

To summarize the evidence presented thus far, we have shown that the assumption of diversifiable default risk cannot be rejected based on the evidence contained in existing studies. Because this assumption greatly facilitates pricing and credit risk management, its empirical validation merits further investigation. A rigorous testing of this hypothesis is contingent upon a careful decompo-

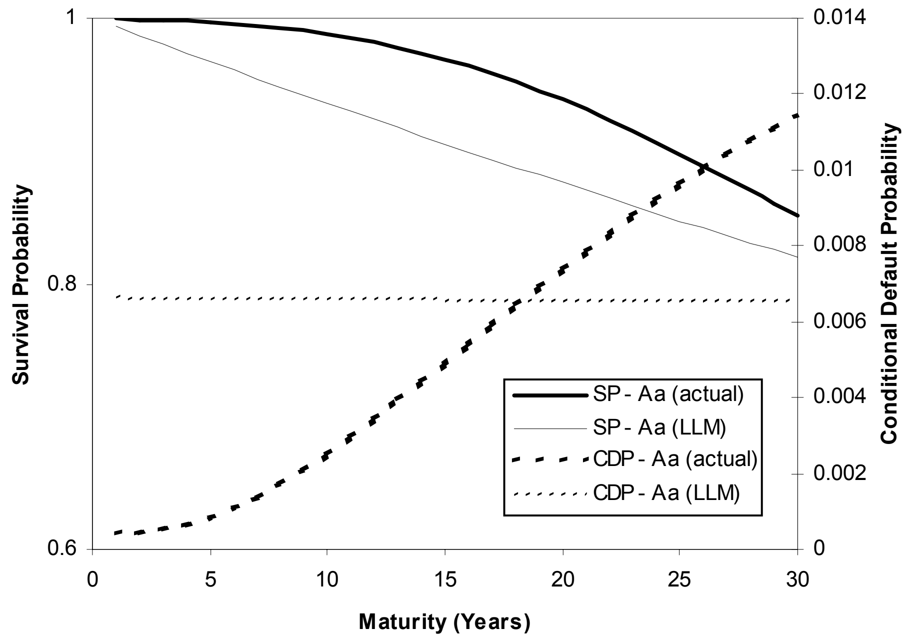

Figure 5: Term Structure of Survival Probabilities and Conditional Default Probabilities for a "Refreshed" LIBOR Quality Issuer. The "actual" series are computed from Moody's data. The "LLM" series are computed using the estimates from Liu, Longstaff and Mandell (2001).

sition of the components of the credit spread. In the absence of such a detailed empirical analysis, we expect that this assumption will produce more accurate default rate estimates for lower-rated issuers in the mid- to long-maturity range.

4.2 Mortgage-Backed Securities

In the literature on mortgage-backed securities, it is common to deviate from an assumption of "rational" prepayment (based on American bond option techniques) and include an empirically estimated prepayment function. This empirical prepayment function captures the stochastic nature of prepayments stemming from differences in transactions costs and individual circumstances. Since there is some (but not perfect) rationality in prepayment behavior, the specified empirical prepayment function often depends on both the level of interest rates and the history of interest rates to quantify a "burnout factor." In both Schwartz and Torous (1989) and Stanton (1995), an empirical prepayment function is specified and the functional form of the prepayment intensity is the same under both the "physical" and the "risk-neutral" measure.

To be concrete, let us consider for example the model of Stanton (1995). In this model, the (Treasury) bond market is controlled by a short rate CIR process r whose behavior differs under the physical measure and the pricing measure by a market price of risk parameter (which Stanton denotes q). From this risk-neutral evolution of r, Stanton obtains the value $M_u^l(r_t, t)$ at time t of the mortgage liability conditional on the prepayment option remaining unexercised. Denoting the exercise price of the mortgage (including transactions costs) $F_t(1 + X_t)$, Stanton then specifies the prepayment intensity as

$$\lambda + \rho 1_{\{M_u^l(r_t,t) \geq F_t(1+X_t)\}} \tag{31}$$

where λ is an intensity of prepayment which is unrelated to changes in interest rates and ρ is an increase in the prepayment intensity which kicks in when the value, conditional on no exercise, of the liability exceeds the cost of prepaying. This intensity specification is then used for pricing the mortgage-backed security and the parameters λ and ρ for assessing the empirical prepayment rates.

This is an implicit conditional diversification condition. Note, however, that in the model there are risk adjustments through the interest rate process. This implies that prepayment frequencies are actually different under the two measures, despite the invariant prepayment function. With systematic prepayment risk, a risk premium would multiply the empirical prepayment function by a positive factor, and in this case the prepayment function would not be the same under the two measures.

4.3 Credit Derivatives

To summarize the importance of Proposition 5 for the credit derivatives markets, note that conditional diversification implies that idiosyncratic risk (due to factors other than the state variables X) is diversifiable in large loan portfolios. Consequently, in the market's pricing of risky debt, there would be no risk premium for idiosyncratic default risk. Risk premia would only enter the loan's return process through the state variables and their influence on the return.[21]

Thus, only large loan portfolios would be "efficient" in the sense that their expected returns compensate for the risk borne. Small bond portfolios would be "inefficient," bearing idiosyncratic default risk that is not priced in higher expected returns. Unfortunately, holding a large diversified loan portfolio may require that the portfolio satisfy significant geographical and industry diversification that, because of market and origination frictions (discreteness in the face value of a bond, limited information, and limited supply), may be difficult to obtain.

This difficulty provides one rationale for the existence and use of credit derivatives (e.g., default swaps). For small and undiversified loan portfolios, credit derivatives provide the vehicle for obtaining this diversification without the direct purchase or sales of the individual loans themselves.

[21] This is evident, for example, if the assumptions underlying equations (19) and (20) hold.

5 Conclusion

In this paper we examine the general specification of default risk premium in the context of an intensity-based model. We argue that the "drift change of the intensity" used in the empirical literature constitutes a restriction on the set of possible default risk premia. We show that this restriction can be justified through a suitably defined notion of conditional diversifiable default risk, which leads to the equivalence between the empirical and martingale intensities, either exactly or in an asymptotic sense. We stress that this does not imply the equivalence between implied and actual default probabilities. Indeed, if the intensities are sensitive to the factors carrying a risk premium, the deviations in the long end between implied and actual default probabilities can be substantial. It does, however, imply the equivalence of actual and implied default probabilities in the very short end. If one believes that this equivalence does not hold, even after adjusting for taxes, liquidity risk or informational asymmetries, then a risk premium on the jump event itself should be included, introducing a risk adjustment which is different from the frequently used drift change of the default intensity.

An important application of our equivalence result is that it integrates pricing and risk management for defaultable securities. This has two meanings. First, we can estimate an empirical default intensity from historical default data and use it to price defaultable bonds. This provides a link from empirical default prediction models such as Altman (1968, 1993) and Shumway (2001) to pricing models. Second, we can imply out a martingale default intensity from defaultable bond prices and use it to construct actual default probabilities. A set of estimated systematic risk premia enables us to go back and forth between the two worlds. We demonstrate the consistency of this methodology in the context of some existing empirical studies of corporate bonds and interest rate swaps.

For further applications, we observe that exactly the same methodology can be applied to mortgage-backed securities. In this case, the relevant quantities are the prepayment functions and the prepayment frequencies. We also show that conditional diversification imparts a sense in which credit derivatives can be used to achieve a more "efficient" credit risk portfolio.

A Appendices

A.1 Proof of Proposition 1

Proof: The terminal payoff of an equally weighted portfolio of the first I bonds with maturity T is $\frac{1}{I}\sum_{i=1}^{I} 1_{\{\tau^i > T\}}$. Since the existence of an equivalent martingale measure Q is assumed, the price of this portfolio is

$$q^I(t,T) = E_t^Q\left(\exp\left(-\int_t^T r_u du\right)\frac{1}{I}\sum_{i=1}^{I} 1_{\{\tau^i > T\}}\right). \tag{32}$$

Under the empirical measure P, conditioning on $\mathcal{F}_{T^*}^X$ the default times are i.i.d. random variables, so by the strong law of large numbers, the average of the indicator functions will converge almost surely to its mean, which is just the probability distribution given in equation (9). Hence

$$P_t\left(\left\{\omega : \frac{1}{I}\sum_{i=1}^{I} 1_{\{\tau^i > T\}} \underset{I\to\infty}{\longrightarrow} \exp\left(-\int_t^T \lambda_u du\right)\right\} \mid \mathcal{F}_{T^*}^X\right) = 1, \tag{33}$$

which implies that

$$P_t\left(\left\{\omega : \frac{1}{I}\sum_{i=1}^{I} 1_{\{\tau^i > T\}} \underset{I\to\infty}{\longrightarrow} \exp\left(-\int_t^T \lambda_u du\right)\right\}\right) = 1. \tag{34}$$

The same probability under Q must be 1 as well due to the equivalence to P.

By dominated convergence, the price of this equally weighted portfolio in the limit will be

$$q^\infty(t,T) = \lim_{I\to\infty} q^I(t,T) = E_t^Q\left(\exp\left(-\int_t^T (r_u + \lambda_u)\, du\right)\right). \tag{35}$$

However, since all firms have the same martingale default intensity and recovery rates are all zero, all bond prices are the same. By price linearity,

$$v^i(t,T) = q^\infty(t,T),\ i = 1,\ldots,I, \tag{36}$$

which by equation (13) suggests that

$$E_t^Q\left(\exp\left(-\int_t^T \left(r_u + \widetilde{\lambda}_u\right) du\right)\right) = E_t^Q\left(\exp\left(-\int_t^T (r_u + \lambda_u)\, du\right)\right). \tag{37}$$

Due to left-continuity, the above can be differentiated with respect to T. Setting $T = t$, we obtain $\lambda_t = \widetilde{\lambda}_t$. ■

A.2 Proof of Proposition 2

Proof: Let $V(\cdot)$ be the variance operator under P. Then

$$V\left(Y_T^I - S_T^I\right) = E^Q\left(V\left(Y_T^I - S_T^I \mid \mathcal{F}_T^X\right)\right) + V\left(E\left(Y_T^I - S_T^I \mid \mathcal{F}_T^X\right)\right). \tag{38}$$

We now show that this variance goes to 0 as $I \to \infty$:

$$V\left(Y_T^I - S_T^I \mid \mathcal{F}_T^X\right) = \sum_{i=1}^{I} \left(w_I^i C_T^i\right)^2 \exp\left(-\int_0^T \lambda_u^i du\right)\left(1 - \exp\left(-\int_0^T \lambda_u^i du\right)\right), \tag{39}$$

and so

$$E^Q\left(V\left(Y_T^I - S_T^I \mid \mathcal{F}_T^X\right)\right) \leq \sum_{i=1}^{I} E\left(\left(w_I^i C_T^i\right)^2\right) \leq K^2 \sum_{i=1}^{I} \left(w_I^i\right)^2 \to 0 \text{ as } I \to \infty. \tag{40}$$

The second term in equation (38) is zero since

$$E\left(Y_T^I - S_T^I \mid \mathcal{F}_T^X\right) = 0. \tag{41}$$

Since the variance of the difference goes to 0, we have shown that

$$\left\|Y_T^I - S_T^I\right\| \to 0 \text{ in } L^2. \tag{42}$$

Therefore in a pricing functional that is continuous in L^2 the prices must converge to each other as well. ■

A.3 Proof of Proposition 3

Proof: Let $Y_t^i = \int_0^t b^i(u,T)\,du + \int_0^t a^i(u,T)\,dX_u$. From (20),

$$\frac{q^i(t,T)}{B(t)} = \mathcal{E}\left(Y_t^i\right), \tag{43}$$

where $B(t) = \exp\left(\int_0^t r_u du\right)$ is the money market account and $\mathcal{E}(\cdot)$ is the Doléans-Dade exponential operator. Similarly, from (19),

$$\begin{aligned}
\frac{v^i(t,T)}{B(t)} &= \mathcal{E}\left(Y_t^i - M_t^i\right) \\
&= \mathcal{E}\left(Y_t^i - \widetilde{M}_t^i + \int_0^t \left(\lambda_u^i - \widetilde{\lambda}_u^i\right) du\right) \\
&= \mathcal{E}\left(Y_t^i - \widetilde{M}_t^i\right) \exp\left(\int_0^t \left(\lambda_u^i - \widetilde{\lambda}_u^i\right) du\right) \\
&= \mathcal{E}\left(Y_t^i\right) \mathcal{E}\left(-\widetilde{M}_t^i\right) \exp\left(\int_0^t \left(\lambda_u^i - \widetilde{\lambda}_u^i\right) du\right),
\end{aligned} \tag{44}$$

where $\widetilde{M}_t^i = N_t^i - \int_0^t \widetilde{\lambda}_u^i du$ is a Q-martingale due to the existence of a Q-intensity $\widetilde{\lambda}_t^i$ associated with N_t^i. The above derivation uses the formula $\mathcal{E}(X)\mathcal{E}(Y) = \mathcal{E}(X + Y + [X,Y])$ and properties of the quadratic covariation $[X,Y]$ extensively. Specifically, the third equality is due to the fact that $\int_0^t \left(\lambda_u^i - \widetilde{\lambda}_u^i\right) du$ is a process of finite variation (FV)—it is the difference between two increasing processes, and that it is also continuous. The last equality results from $\left[Y_t^i, \widetilde{M}_t^i\right] = 0$ since $\widetilde{M}_t^i$ is a quadratic pure jump semimartingale that shares no jump with Y_t^i.[22]

Since $q^i(t,T)/B(t)$ is a Q-martingale, $\mathcal{E}\left(Y_t^i\right)$ is a Q-martingale due to (43). On the other hand, $\mathcal{E}\left(-\widetilde{M}_t^i\right)$ is also a Q-martingale. Let $A_t = \mathcal{E}\left(Y_t^i\right)$ and $B_t = \mathcal{E}\left(-\widetilde{M}_t^i\right)$. Their product $A_t B_t$ is a Q-local martingale since

$$\begin{aligned}
[A_t, B_t] &= \left[1 + \int_0^t A_{u-} dY_u^i, 1 - \int_0^t B_{u-} d\widetilde{M}_u^i\right] \\
&= -\int_0^t A_{u-} B_{u-} d\left[Y_u^i, \widetilde{M}_u^i\right] \\
&= 0.
\end{aligned} \tag{45}$$

Therefore one can write $v^i(t,T)/B(t) = U_t V_t$, where $U_t = \mathcal{E}\left(Y_t^i\right)\mathcal{E}\left(-\widetilde{M}_t^i\right)$ is a Q-local martingale and $V_t = \exp\left(\int_0^t \left(\lambda_u^i - \widetilde{\lambda}_u^i\right) du\right)$ is an FV process because it is a monotonic transformation of an FV process. V_t is also predictable because it is pathwise continuous. It should then have the following semimartingale decomposition:

$$U_t V_t = U_0 V_0 + \int_0^t U_{u-} dV_u + W_t, \tag{46}$$

where W_t is a Q-local martingale with $W_0 = 0$.[23] The decomposition (46) should be unique, since the second term, the FV component, is continuous and hence predictable. However, since $U_t V_t = v^i(t,T)/B(t)$ is itself a Q-martingale, the FV component in the decomposition must vanish. This implies that V_t is a constant, and subsequently equal to one. Hence $\lambda_t^i = \widetilde{\lambda}_t^i$. ■

A.4 Proof of Proposition 4

Proof: A result in Artzner and Delbaen (1995) states that

$$\widetilde{\lambda}_t^i = \lambda_t^i \frac{K_t}{Z_{t-}}, \tag{47}$$

where the predictable process K_t is defined through

$$K_{\tau^i} = E^P\left(\frac{dQ}{dP} \mid \mathcal{F}_{\tau^i -}\right). \tag{48}$$

However,

$$\begin{aligned}
E^P\left(\frac{dQ}{dP} \mid \mathcal{F}_{\tau^i -}\right) &= E^P\left(E^P\left(\frac{dQ}{dP} \mid \mathcal{F}_{\tau^i}\right) \mid \mathcal{F}_{\tau^i -}\right) \\
&= E^P\left(Z_{\tau^i} \mid \mathcal{F}_{\tau^i -}\right) \\
&= Z_{\tau^i -},
\end{aligned} \tag{49}$$

since Z_t is adapted to the filtration generated by the state variables and by construction these processes jump at τ^i with probability zero. The equality between $\widetilde{\lambda}_t^i$ and λ_t^i follows. ■

[22] This is true if the jumps of X_t coincide with those of N_t^i with zero probability. The standard reference for the results on the quadratic covariation and the Doléans-Dade exponential is Protter (1990, II-6 and II-8).

[23] See Dellacherie and Meyer (1982, p.223) for detail.

A.5 Proof of Proposition 5

Proof: Let $a_i = 1/2^i$, $i \geq 1$. The infinite economy can be embedded into a one-dimensional semimartingale

$$S_t = X_t + \sum_{i=1}^{\infty} a_i 1_{\{\tau^i \leq t\}}. \tag{50}$$

Let μ denote the (random) jump measure associated with this process, i.e.,

$$\mu([0,t] \times i) = 1_{\{\tau^{\hat{\imath}} \leq t\}},\ i \geq 1, \tag{51}$$

and let ν be the compensating measure of μ (i.e., the third characteristic of S) as defined for example in Jacod and Mémin (1976). Clearly, S_t is a locally bounded (hence special) semimartingale with characteristics under P given as

$$\begin{aligned} d\alpha_t &= \mu(t)\,dt + \sum_{i=1}^{\infty} a_i \lambda_t^i dt, \\ d\beta_t &= \sigma^2(t)\,dt, \\ \text{and } \nu(dt, \{i\}) &= a_i \lambda_t^i dt. \end{aligned} \tag{52}$$

Here, and in what follows, we will omit ω from our notation. We can recover the jump processes and the state variable from S by defining

$$X_t = S_t - \sum_{u \leq t} \Delta S_u, \tag{53}$$

$$\text{and } N_t^i = 1_{\{\Delta S_u = a_i \text{ for some } u \leq t\}}. \tag{54}$$

Now assume that Q is equivalent to P. The semimartingale S also has bounded jumps under an equivalent measure and hence it is special under Q as well. Since S is also quasi-left continuous [see Jacod and Mémin (1976)], we have that the characteristics under Q are given as

$$\begin{aligned} d\widetilde{\alpha}_t &= d\alpha_t + g(t)\,\sigma(t)\,dt + \sum_{i=1}^{\infty} a_i \left(Y(t,i) - 1\right) \lambda_t^i dt, \\ d\widetilde{\beta}_t &= d\beta_t, \\ \text{and } \widetilde{\nu}(dt, \{i\}) &= Y(t,i)\,\nu(dt, \{i\}). \end{aligned} \tag{55}$$

This follows from Theorem 3.3 of Jacod and Mémin (1976).

Since we know exactly the form of the measure change on the diffusion part, we have automatically that $\int_0^{T^*} g^2(u)\,\sigma^2(u)\,du < \infty$, P-a.s., and hence we see that the condition in Theorem 4.1 of Jacod and Mémin (1976) is equivalent to the condition that

$$\sum_{i=1}^{\infty} \int_0^{T^*} |Y(u,i) - 1|\, 1_{\{Y>2\}} \lambda_u^i du + \int_0^{T^*} \left(Y(u,i) - 1\right)^2 1_{\{Y \leq 2\}} \lambda_u^i du < \infty,\ P\text{-}a.s. \tag{56}$$

But using the inequality

$$\left(1-\sqrt{y}\right)^2 \leq (y-1)^2 \, 1_{\{y\leq 2\}} + |y-1| \, 1_{\{y>2\}} \leq \frac{1}{\left(1-\sqrt{2}\right)^2} \left(1-\sqrt{y}\right)^2 \tag{57}$$

which holds for positive y we see that this is equivalent to

$$\sum_{i=1}^{\infty} \int_0^{T^*} \left(1-\sqrt{Y(u,i)}\right)^2 \lambda_u^i du < \infty, \tag{58}$$

as was to be proved. ■

References

Altman, E. I., 1968, "Financial Ratios, Discriminant Analysis, and the Prediction of Corporate Bankruptcy," *Journal of Finance*, 23, 589-609.

Altman, E. I., 1993, *Corporate Financial Distress and Bankruptcy: A Complete Guide to Predicting and Avoiding Distress and Profiting from Bankruptcy*, John Wiley & Sons, New York.

Artzner, P., and F. Delbaen, 1995, "Default Risk Insurance and Incomplete Markets," *Mathematical Finance*, 5, 187-195.

Back, K., 1991, "Asset Pricing for General Processes," *Journal of Mathematical Economics*, 20, 371-395.

Björk, T., and B. Näslund, 1998, "Diversified Portfolios in Continuous Time," *European Finance Review*, 1, 361-387.

Brémaud, P., 1981, *Point Processes and Queues: Martingale Dynamics*, Springer-Verlag, New York.

Collin-Dufresne, P., R. Goldstein, and S. Martin, 2000, "The Determinants of Credit Spread Changes," Working Paper, Carnegie Mellon University; forthcoming in *Journal of Finance.*

Connor, G., 1984, "A Unified Beta Pricing Theory," *Journal of Economic Theory*, 34, 13-31.

Cox, J., J. Ingersoll, and S. Ross, 1981, "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385-408.

Dai, Q., and K. Singleton, 2000, "Specification Analysis of Affine Term Structure Models," *Journal of Finance*, 55, 1943-1978.

Dellacherie, C., and P. Meyer, 1982, *Probabilities and Potential B: Theory of Martingales*, North-Holland, Amsterdam.

Duffee, G. R., 1999, "Estimating the Price of Default Risk," *Review of Financial Studies*, 12, 197-226.

Duffie, J. D., and R. Kan, 1996, "A Yield-Factor Model of Interest Rates," *Mathematical Finance,* 6, 379-406.

Duffie, D., and D. Lando, 2001, "Term Structures of Credit Spreads with Incomplete Accounting Information," *Econometrica*, 69, 633-664.

Duffie, J. D., and K. J. Singleton, 1997, "An Econometric Model of the Term Structure of Interest Rate Swap Yields," *Journal of Finance* 52, 1287-1321.

Duffie, J. D., and K. J. Singleton, 1999, "Modeling Term Structures of Defaultable Bonds," *Review of Financial Studies,* 12, 687-720.

Elton, E., M. Gruber, D. Agrawal, and C. Mann, 2001, "Explaining the Rate Spread on Corporate Bonds," *Journal of Finance* 56, 247-278.

Heath, D., R. Jarrow, and A. Morton, 1992, "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation," *Econometrica,* 60, 77-105.

Helwege, J., and C. M. Turner, 1999, "The Slope of the Credit Yield Curve for Speculative Grade Issuers," *Journal of Finance,* 54, 1869-1884.

Jacod, J., and J. Mémin, 1976, "Caractéristiques locales et conditions de continuité absolue pour les semimartingales," *Z. Wahrsch. Verw. Geb.,* 35, 1-37.

Jarrow, R. A., 1988, "Preferences, Continuity, and the Arbitrage Pricing Theory," *Review of Financial Studies*, 1, 159-172.

Jarrow, R. A., D. Lando, and S. M. Turnbull, 1997, "A Markov Model for the Term Structure of Credit Risk Spread," *Review of Financial Studies,* 10, 481-523.

Jarrow, R. A., and D. Madan, 1995, "Option Pricing Using the Term Structure of Interest Rates to Hedge Systematic Discontinuities in Asset Return," *Mathematical Finance,* 5, 311-336.

Jarrow, R. A., and S. M. Turnbull, 1995, "Pricing Derivatives on Financial Securities Subject to Credit Risk," *Journal of Finance,* 50, 53-85.

Jarrow, R. A., and F. Yu, 1999, "Counterparty Risk and the Pricing of Defaultable Securities," Working Paper, University of California at Irvine; forthcoming in *Journal of Finance.*

Kabanov, Y. M., and D. O. Kramkov, 1998, "Asymptotic Arbitrage in Large Financial Markets," *Finance and Stochastics*, 2, 143-172.

Klein, I., and W. Schachermayer, 1997, "Asymptotic Arbitrage in Noncomplete Large Markets," *Theory of Probability and Its Applications*, 41, 780-788.

Kusuoka, S., 1999, "A Remark on Default Risk Models," *Advances in Mathematical Economics*, 1, 69-82.

Lando, D., 1994, "Three Essays on Contingent Claims Pricing," Ph.D. Dissertation, Cornell University.

Lando, D., 1998, "On Cox Processes and Credit Risky Securities," *Review of Derivatives Research,* 2, 99-120.

Liu, J., F. Longstaff, and R. Mandell, 2001, "The Market Price of Credit Risk," Working Paper, University of California at Los Angeles.

Pedrosa, M., and R. Roll, 1998, "Systematic Risk in Corporate Bond Credit Spreads," *Journal of Fixed Income*, 8, 9-26.

Protter, P., 1990, *Stochastic Integration and Differential Equations: A New Approach,* Springer-Verlag, New York.

Schwartz, E., and W. Torous, 1989, "Prepayment and the Valuation of Mortgage-Backed Securities," *Journal of Finance*, 44, 375-392.

Shumway, T., 2001, "Forecasting Bankruptcy More Accurately: A Simple Hazard Model," *Journal of Business,* 74, 101-124.

Stanton, R., 1995, "Rational Prepayment and the Valuation of Mortgage-Backed Securities," *Review of Financial Studies*, 8, 677-708.

Vasicek, O., 1977, " An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177-188.