

DEFAULT RISK AND DIVERSIFICATION: THEORY AND EMPIRICAL IMPLICATIONS

ROBERT A. JARROW

Cornell University

DAVID LANDO

University of Copenhagen

FAN YU

University of California, Irvine

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 intensities that follows from diversifiable default risk greatly facilitates the pricing and management of credit risk. We emphasize that this is not an equivalence in distribution, and illustrate its importance using credit spread dynamics estimated in Duffee (1999). We also argue that the assumption of diversifiability is implicitly used in certain existing models of mortgage-backed securities.

KEY WORDS: default risk, diversification, default risk premium, empirical and martingale default intensities

1. INTRODUCTION

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.

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, Miguel Ferreira, Michael Gordy, Jean Jacod, Lionel Martellini, Fuvio Ortu, Ken Singleton, Fernando Zapatero, an anonymous referee, and seminar participants at Stanford, USC, UC-Irvine, PIMCO, the 2002 Credit Risk Summit, the 11th Derivative Securities Conference, the 2001 QMF Conference in Sydney, the 2001 Meeting of the European Finance Association, 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. We also thank Jens Christensen for excellent research assistance. David Lando acknowledges the partial support of the Danish Social Science Foundation.

Manuscript received December 2002; final revision received November 2003.

Address correspondence to Prof. Fan Yu, Graduate School of Management, University of California Irvine, Irvine, CA 92697-3125; e-mail: fanyu@uci.edu.

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; hereafter CIR) and Vasicek (1977) are the classic 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, for example, in 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 paper proceeds a step further along the risk premium interpretation by computing the expected returns on 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 a representative investor or a single optimizing agent, it is easy to see that the 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 Mémén (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 Arbitrage Pricing Theory (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.

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.

Our results show 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 nonzero risk premium, then the drift change in the intensity as specified above is not the appropriate specification. Contrary to the change of drift for diffusion state variables, a systematic jump risk premium will 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 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 empirical implications of 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 using 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 one is 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) 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 for the long end of the term structure, but that the short end is more problematic. Adjusting for liquidity and tax effects here may provide a partial explanation of the deviation, but a more detailed empirical study along the lines of Driessen (2002) is needed to test the hypothesis formally.

As a final observation on empirical implications, 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 that 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 2 we provide an intuitive illustration of different forms of default risk premia. In Section 3 we formally introduce the concept of conditionally diversifiable default risk using the framework of Lando (1994). In Section 4, 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 3. In Section 5 we discuss the empirical implications of diversifiable default risk. We conclude with Section 6.

2. VARIATIONS IN DEFAULT RISK VERSUS EVENT RISK

In Section 3 we will explicitly construct a reduced-form credit risk model with several issuers. However, before giving this construction, we will explain in a very simple setting

the critical distinction between the two types of default risk premia that we are trying to understand.

2.1. 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 that 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

$$(2.1) \quad p(t, T) = E_t^Q \exp \left(- \int_t^T r_u du \right).$$

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

$$(2.2) \quad v(t, T) = E_t^Q \exp \left(- \int_t^T (r_u + \tilde{\lambda}_u) du \right).$$

Lando (1994, 1998) extended this formula into pricing building blocks for contingent claims, and Duffie and Singleton (1999) showed that with a fractional recovery rate one obtains the same formula except that $\tilde{\lambda}$ 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

$$(2.3) \quad d\lambda_t = \kappa(\theta - \lambda_t) dt + \sigma \sqrt{\lambda_t} dW_t^P$$

and the behavior of λ_t under the equivalent measure Q as

$$(2.4) \quad d\lambda_t = (\kappa + \nu) \left(\frac{\kappa\theta}{\kappa + \nu} - \lambda_t \right) dt + \sigma \sqrt{\lambda_t} dW_t^Q,$$

where the processes W^P and W^Q are Brownian motions under P and Q , respectively.³ Constants κ , θ , ν , and σ are 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, Duffie (1999), Duffie and Singleton (1997), and Liu et al. (2001) all have parameters playing the role of a drift change in the intensity.

³ W^P and W^Q are related through $dW_t^P = dW_t^Q - \frac{\nu}{\sigma} \sqrt{\lambda_t} dt$.

When trying to quantify risk premia in default markets, one must realize that this drift change in the intensity only captures the compensation of taking on default risk that 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 $\tilde{\lambda}_t = \mu\lambda_t$, with a dynamics given by

$$(2.5) \quad d\tilde{\lambda}_t = (\kappa + \nu) \left(\frac{\mu\kappa\theta}{\kappa + \nu} - \tilde{\lambda}_t \right) dt + \sqrt{\mu}\sigma\sqrt{\tilde{\lambda}_t} dW_t^Q.$$

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 Mémén (1976) is that, in contrast to that of Artzner and Delbaen (1995), it provides an explicit characterization of the possible risk premia. Jacod and Mémén's results will prove particularly useful when we consider 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.

2.2. 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 that 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. Let us 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 (2.5) with (2.4)). Using the results in Lando (1994, 1998), we see that the price of this claim is

$$(2.6) \quad v(0, T) = E_0^Q \int_0^T (r_t + \lambda_t) \exp\left(-\int_0^t (r_u + \lambda_u) du\right) dt \\ + E_0^Q \exp\left(-\int_0^T (r_u + \lambda_u) du\right) = 1,$$

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

$$(2.7) \quad v(0, T) = E_0^Q \int_0^T (r_t + \lambda_t) \exp\left(-\int_0^t (r_u + \mu\lambda_u) du\right) dt \\ + E_0^Q \exp\left(-\int_0^T (r_u + \mu\lambda_u) du\right),$$

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

2.3. 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 (2.3) and (2.4), with parameters specified as $\kappa = 0.186$, $\theta = 0.00499$, $\sigma = 0.074$, and $\nu = -0.216$. These values correspond to Duffee's (1999) 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 (2.2), assuming of course zero recovery and $\tilde{\lambda} = \mu\lambda$. We then consider the case where μ is taken to be equal to 1. This corresponds to the default risk premium arising solely from a "drift change in the intensity" as explained above. In this case, one naturally assumes that the dynamics of the Q -intensity $\tilde{\lambda}$ is given by

$$(2.8) \quad d\tilde{\lambda}_t = (\kappa + \nu') \left(\frac{\kappa\theta}{\kappa + \nu'} - \tilde{\lambda}_t \right) dt + \sigma \sqrt{\tilde{\lambda}_t} dW_t^Q.$$

However, in order to match the correct yield spread, we infer 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 $\mu > 1$.

3. 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.

3.1. Default Processes

Following the standard literature on large markets, we assume that there exist a countably infinite number of firms in this economy.⁴ 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 -that is, $N_t^i = 1_{\{\tau^i \leq t\}}$. This process is assumed to have an intensity process λ^i that 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, conditional 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.

Formally, we start with a filtered probability space $(\Omega, \mathcal{F}_{T^*}^X, \{\mathcal{F}_t^X\}_{t=0}^{T^*}, P^X)$, 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 $\{\lambda_u^i, i = 1, 2, \dots\}$ that are predictable with respect to \mathcal{F}_t^X .⁵ To construct the default processes, we first augment the probability space with a sequence of i.i.d. unit exponential random variables $\{E^i, i = 1, 2, \dots\}$ that are independent of the process X_t . Then for each i we define a stopping time $\tau^i = \inf\{t : \int_0^t \lambda_u^i du \geq E^i\}$. The i th default process can be defined as $N_t^i = 1_{\{\tau^i \leq t\}}$, which can only take one of two values: 0 or 1. With this construction, the compensated point process $M_t^i = N_t^i - \int_0^{t \wedge \tau^i} \lambda_u^i du$ 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 $(\Omega, \mathcal{F}_{T^*}, \{\mathcal{F}_t\}_{t=0}^{T^*}, P)$, where the augmented filtration $\mathcal{F}_t = \mathcal{F}_t^X \vee \mathcal{G}_t^1 \vee \mathcal{G}_t^2 \vee \dots$, and \mathcal{G}_t^i is the filtration generated by the i th 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 $\mathcal{F}_{T^*}^X$, the default processes are independent of each other. This independence captures the essence of conditional diversifiability.

With this construction, the conditional distribution of the i th default time is (assuming no default before t)

$$(3.1) \quad P_t(\tau^i > s \mid \mathcal{F}_{T^*}^X) = \exp\left(-\int_t^s \lambda_u^i du\right), \quad s \in [t, T^*],$$

⁴ 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 subeconomies 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. This approach implies that the subeconomies 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 4.2.

⁵ 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.

and consequently the unconditional distribution is

$$(3.2) \quad P_i(\tau^i > s) = E_t^P \left(\exp \left(- \int_t^s \lambda_u^i du \right) \right), \quad s \in [t, T^*],$$

where $E_t^P(\cdot)$ denotes the expectation under P conditional on \mathcal{F}_t . This completes the specification of the default processes under the empirical measure. We now turn to the pricing of defaultable bonds issued by the firms.

3.2. Valuation of Defaultable Bonds

Let the time- t price of a zero-coupon bond with maturity T issued by firm i 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).⁶

In addition, in this economy there is trading of a collection of default-free zero-coupon bonds 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

$$(3.3) \quad p(t, T) = E_t^{Q^X} \left(\exp \left(- \int_t^T r_u du \right) \right).$$

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

$$(3.4) \quad v^i(t, T) = E_t^Q \left(\exp \left(- \int_t^T r_u du \right) (\delta^i 1_{\{\tau^i \leq T\}} + 1_{\{\tau^i > T\}}) \right).$$

3.2.1. Default Processes under the Equivalent Martingale Measure. At this point it is useful to note the properties of the default processes under an equivalent change of measure from P to Q . We are interested in knowing when the defaultable bond prices in equation (3.4) can be expressed in terms of $\tilde{\lambda}_t^i$, the martingale default intensities under Q . This is an important first step because we will attempt to derive a relation between the physical intensity and the martingale intensity using the process of diversification, forming larger and better diversified portfolios. To do this, we need to identify conditions under which the pricing formula from Lando (1998) holds:

$$(3.5) \quad v^i(t, T) = \delta^i p(t, T) + 1_{\{\tau^i > t\}} (1 - \delta^i) E_t^Q \exp \left(- \int_t^T (r_u + \tilde{\lambda}_u^i) du \right).$$

First, as we mentioned earlier, Artzner and Delbaen (1995) showed that it is no restriction to assume the existence of an intensity under an equivalent measure. Therefore, the notion of an alternative intensity $\tilde{\lambda}_t^i$ under Q is well justified. Second, under a change of measure the intensity will not stay invariant. Generally one has $\tilde{\lambda}_t^i = \mu_t^i \lambda_t^i$

⁶ Other recovery rate assumptions are possible, including the “recovery of market value” assumption of Duffie and Singleton (1999).

for some strictly positive \mathcal{F}_t -predictable process μ_t , with the Radon-Nikodym density martingale $Z_t = E_t^P(\frac{dQ}{dP})$ represented by (ignoring the part of the measure change corresponding to the state variables)

$$(3.6) \quad Z_t = 1 + \sum_{i=1}^I \int_0^t Z_{s-} (\mu_s^i - 1) dM_s^i,$$

for finite I . Recall that M_t^i is the compensated jump martingale from our previous construction of the default processes.

Apart from the obvious consequence of $\tilde{\lambda}_t^i \neq \lambda_t^i$, this has an additional implication for our model. Although we have assumed that λ_t^i is adapted to \mathcal{F}_t^X , there is no reason to expect that such a property will be preserved for $\tilde{\lambda}_t^i$. The general dependency on the filtration \mathcal{F}_t can be interpreted as due to counterparty risk (see Jarrow and Yu 2001), or changing perception of default risk due to specific events (see Collin-Dufresne, Goldstein, and Helwege 2002).

When the intensities are adapted to \mathcal{F}_t , generally the default processes are no longer independent conditional on the history of X_t and a potentially recursive structure results.⁷ Kusuoka (1999) constructed explicit examples using the measure change (3.6) and demonstrated that the counterparts to (3.1) and (3.2) do not hold under the new measure. This causes the explicit link between prices and intensities in (3.5) to fail. Collin-Dufresne, Goldstein, and Hugonnier (2003) showed that the interpretation of defaultable bond prices as promised cash flow discounted at a default risk-adjusted short rate can be preserved if the expectation in (3.5) is taken under an alternative measure that assigns zero probability to firm i defaulting before time T . However, their measure change is firm-specific and generally it is not possible to construct a single alternative measure that preserves the structure of (3.5) for bonds issued by all firms.

Given these difficulties in modeling the default processes under the pricing measure, we can proceed with two different approaches. First, we 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; that is, the goal is to say something about the properties of Q on the filtration generated by the state variables *and* the default processes. This constitutes our main result to be found in Section 4.3 below. This approach requires no additional structure beyond those stated above. The downside, however, is that the equivalence between the physical intensity and the martingale intensity is only “asymptotic.”

In comparison, our second approach requires the additional assumption of conditionally independent defaults under the pricing measure Q , which affords an explicit application of the pricing formula (3.5) in the context of well-diversified portfolios. With the use of (3.5), we are essentially proposing a factor structure on prices similar to that of the APT on asset returns. At the expense of being more restrictive, this approach allows us to use diversification and utility-based arguments to motivate exact equivalence between the intensities. These results are contained in Sections 4.1 and 4.2.

To understand what it takes to preserve conditional independence, we note the following sufficient condition.

PROPOSITION 3.1. *Assume that μ_t^i in the Radon-Nikodym density (3.6) are \mathcal{F}_t^X -adapted. Then the default processes N_t^i are independent conditional on \mathcal{F}_T^X under the pricing measure Q .*

⁷ A “total hazard” construction can be used to build the default processes from i.i.d. unit exponential random variables. See Yu (2003) for details.

Proof. First, we note that when μ_t^i are adapted to \mathcal{F}_t^X , the density Z_T can be written as a product of terms that are conditionally independent given $\mathcal{F}_{T^*}^X$ under P :

$$(3.7) \quad Z_T = \prod_{i=1}^I Z_T^i,$$

where

$$(3.8) \quad Z_T^i = \exp \left(\int_0^T \ln(\mu_s^i) dN_s^i - \int_0^T (\mu_s^i - 1) \lambda_s^i ds \right).$$

This is because by construction N_t^i are mutually independent under P given $\mathcal{F}_{T^*}^X$.

We then note that, for $i \neq j$, $T \leq T^*$, and $T' \leq T^*$:

$$(3.9) \quad \begin{aligned} Q(\tau^i > T, \tau^j > T' | \mathcal{F}_{T^*}^X) &= E^Q(1_{\{\tau^i > T\}} 1_{\{\tau^j > T'\}} | \mathcal{F}_{T^*}^X) \\ &= E^P(1_{\{\tau^i > T\}} 1_{\{\tau^j > T'\}} Z_{T^*} | \mathcal{F}_{T^*}^X) \\ &= E^P \left(1_{\{\tau^i > T\}} 1_{\{\tau^j > T'\}} Z_{T^*}^i Z_{T^*}^j \prod_{k \neq i, j} Z_{T^*}^k | \mathcal{F}_{T^*}^X \right) \\ &= E^P(1_{\{\tau^i > T\}} Z_{T^*}^i | \mathcal{F}_{T^*}^X) E^P(1_{\{\tau^j > T'\}} Z_{T^*}^j | \mathcal{F}_{T^*}^X) \\ &= E^Q(1_{\{\tau^i > T\}} | \mathcal{F}_{T^*}^X) E^Q(1_{\{\tau^j > T'\}} | \mathcal{F}_{T^*}^X) \\ &= Q(\tau^i > T | \mathcal{F}_{T^*}^X) Q(\tau^j > T' | \mathcal{F}_{T^*}^X), \end{aligned}$$

as desired. \square

Intuitively, when μ_t^i is adapted to \mathcal{F}_t^X , so is the martingale intensity $\tilde{\lambda}_t^i$. This suggests that we can again use the conditional independent construction to define the default processes under Q . We also note that under the assumption of Proposition 3.1, individual defaults can still command an event risk premium. However, they do not directly affect the prices of bonds issued by other firms. Therefore we have essentially assumed away counterparty risk as defined by Jarrow and Yu (2001).

4. INVARIANCE OF THE DEFAULT INTENSITY

In this section we present three sets of results on the invariance of the default intensity under the equivalent change of measure. In Section 4.1 we use the notion of L^2 -convergence to examine the pricing of well-diversified portfolios. Assuming conditional independent defaults under Q , we further argue that only “approximate” equivalence should be expected and only under very special circumstances would one obtain exact equivalence. In Section 4.2 we use utility-based arguments to motivate exact equivalence, which can be thought of as a natural consequence of further restrictions on the state price density. The analogy with equilibrium-based exact APT is also noted. Finally, in Section 4.3 we present necessary restrictions on the martingale intensities in a large economy as a result of the conditionally independent construction under P and the absence of arbitrage. This provides the “asymptotic equivalence” mentioned in the previous section.

4.1. The Pricing of Well-Diversified Portfolios

First, we study the implications of diversification by examining the pricing of large, diversified portfolios. Recall that \mathcal{F}_T contains the information of the state variables and of an infinite collection of single jump processes, which are still assumed to be conditionally independent under P given \mathcal{F}_T^X . In this section we assume the existence of a martingale measure Q on (Ω, \mathcal{F}_T) such that the pricing functional induced by Q ,

$$(4.1) \quad \Phi(X) = E^Q \left(\exp \left(- \int_0^T r_s ds \right) X \right),$$

is defined on a domain M that contains $L^2(\Omega, \mathcal{F}_T, P)$. Recall that pricing functionals are defined to be strictly positive.

For simplicity, we let each firm i issue infinitely divisible claims C_T^i , payable at T , all bounded by a constant K and \mathcal{F}_T^X -measurable. Now consider the terminal payoff

$$(4.2) \quad Y_T^I = \sum_{i=1}^I w_I^i C_T^i 1_{\{\tau^i > T\}}.$$

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 $\lim_{I \rightarrow \infty} \sum_{i=1}^I (w_I^i)^2 = 0$ is imposed so that in the limit as $I \rightarrow \infty$ we will end up with a “well-diversified” portfolio.⁸

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

$$(4.3) \quad S_T^I = \sum_{i=1}^I w_I^i \exp \left(- \int_0^T \lambda_u^i du \right) C_T^i.$$

According to (3.1), S_T^I is the expected value of Y_T^I conditional on the filtration \mathcal{F}_T^X . The proof of the following proposition establishes that the distance between Y_T^I and S_T^I in the $L^2(P)$ norm converges to zero, and this implies convergence in price as well.

PROPOSITION 4.1. *Assume that the pricing functional $\Phi : M \rightarrow R$ defined in (4.1) is strictly positive on its domain M and that this domain contains $L^2(\Omega, \mathcal{F}_T, P)$. Then the difference in price between Y_T^I and S_T^I converges to 0 as $I \rightarrow \infty$.*

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

$$(4.4) \quad V(Y_T^I - S_T^I) = E^P(V(Y_T^I - S_T^I | \mathcal{F}_T^X)) + V(E^P(Y_T^I - S_T^I | \mathcal{F}_T^X)).$$

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

$$(4.5) \quad V(Y_T^I - S_T^I | \mathcal{F}_T^X) = \sum_{i=1}^I (w_I^i C_T^i)^2 \exp \left(- \int_0^T \lambda_u^i du \right) \left(1 - \exp \left(- \int_0^T \lambda_u^i du \right) \right),$$

and so

$$(4.6) \quad E^P(V(Y_T^I - S_T^I | \mathcal{F}_T^X)) \leq \sum_{i=1}^I E^P \left((w_I^i C_T^i)^2 \right) \leq K^2 \sum_{i=1}^I (w_I^i)^2 \rightarrow 0 \quad \text{as } I \rightarrow \infty.$$

The second term in equation (4.4) is zero because

$$(4.7) \quad E^P(Y_T^I - S_T^I | \mathcal{F}_T^X) = 0.$$

⁸ Gordy (2003) used a similar condition.

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

$$(4.8) \quad \|Y_T^I - S_T^I\| \rightarrow 0 \text{ in } L_P^2.$$

Since the pricing functional Φ is (strictly) positive on the restriction to the complete space $L^2(\Omega, \mathcal{F}_T, P)$, it is continuous on this restriction and therefore prices must converge to each other as well. \square

This proposition shows that the pricing of well-diversified portfolios of defaultable claims can be reduced to the pricing of \mathcal{F}_t^X -adapted claims. It provides a sense in which diversification has completely eliminated the default event risk component from valuation.⁹

For individual default intensities, however, this argument does not go far enough. To see this informally, we invoke the assumption of conditional independent defaults under Q , which enables us to use (3.5) to rewrite the price of the portfolio as

$$(4.9) \quad p_0(Y_T^I) = \sum_{i=1}^I w_i^I E^Q \left(\exp \left(- \int_0^T (r_u + \tilde{\lambda}_u^i) du \right) C_T^i \right).$$

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

$$(4.10) \quad p_0(S_T^I) = \sum_{i=1}^I w_i^I E^Q \left(\exp \left(- \int_0^T (r_u + \lambda_u^i) du \right) C_T^i \right).$$

The fact that the two are equal in the limit for all well-diversified portfolios suggests a link between the two intensities, although it is difficult to formalize the relationship in this framework. In particular, we see that the best we can hope for are approximate results, as there can be a finite number of violations of invariance that still preserve the equality between (4.9) and (4.10) for well-diversified portfolios.

4.2. Utility-Based Arguments and Exact Equivalence

To obtain exact equivalence with an explicit application of diversification, much stronger assumptions are needed. For example, to the extent that there are a large number of firms within the same industry and credit rating category, issuing bonds with similar characteristics (e.g., maturity and coupon) and trading at similar spreads, one may form a homogeneous portfolio to diversify away the default risk in each bond. In the extreme, in addition to conditional independence under both P and Q , we will simply assume that $\lambda_t^i = \lambda_t$ and $\tilde{\lambda}_t^i = \tilde{\lambda}_t$ for all i . An application of Proposition 4.1 to bonds that have not defaulted by time t yields

$$(4.11) \quad E_t^Q \left(\exp \left(- \int_t^T (r_u + \tilde{\lambda}_u) du \right) \right) = E_t^Q \left(\exp \left(- \int_t^T (r_u + \lambda_u) du \right) \right).$$

In other words, individual bonds can be priced using the dynamics of the physical intensity under the pricing measure, precisely the “drift change of the intensity” discussed in Section 2.1. Assuming left-continuity, the above can be differentiated with respect to T . Setting $T = t$, we obtain $\tilde{\lambda}_t = \lambda_t$.

⁹ Jarrow (1988) explored sufficient conditions on preferences for the pricing operator given in expression (4.1) to be strictly positive. In incomplete markets, when M does not contain $L^2(P)$, he also provided sufficient conditions on preferences for the continuity of the pricing operator.

In this special case, the payoff of the well-diversified portfolio is equal to the conditional expected payoff of the individual bonds for all possible realizations of the state variables. Consequently any risk-averse investor would never place a finite proportion of her wealth in any individual bond. This logic is similar to Connor's (1984) derivation of exact APT using equilibrium-based arguments, where he assumed, among other things, a linear factor structure for asset returns, risk-averse agents, and an insurable factor economy in which any allocation has a well-diversified factor equivalent. Together these assumptions imply an equilibrium allocation that is always well-diversified and, as a result, firm-specific risks are not priced in equilibrium.

These observations suggest that a result of exact equivalence is closely related to restrictions on the equilibrium marginal utility of investors. Recall that under technical conditions presented in Back (1991), the marginal utility of consumption for each investor in a CCAPM setting is proportional to the state price density. That is, there exists for each investor k a constant α_k such that

$$(4.12) \quad Z_t \exp\left(-\int_0^t r_u du\right) = \alpha_k u_c^k(t, c^k(t))$$

holds for the optimal consumption choice $c^k(t)$. If investors hold only well-diversified portfolios in equilibrium, their consumption bundles, and hence the state price density, are insulated from individual default event risks. Indeed, if Z_t is adapted to \mathcal{F}_t^X , we see immediately from equation (3.6) that we must have $\mu_t^i = 1$ and $\tilde{\lambda}_t^i = \lambda_t^i$ for all i ; otherwise Z_t would depend on the compensated jump martingales.

We can draw an analogy with Connor's arguments here by assuming the existence of diversified portfolios that mimic the systematic risk exposure of individual bonds. Motivated by (3.5), we let instantaneous bond returns be jointly driven by common state variable and firm-specific default risk:¹⁰

$$(4.13) \quad \frac{dv^i(t, T)}{v^i(t-, T)} = (r_t + b^i(t, T)) dt + a^i(t, T) dX_t - L^i dM_t^i, \quad i \in Z_{++},$$

where $M_t^i = N_t^i - \int_0^{t \wedge T} \lambda_s^i ds$ is the compensated martingale associated with N_t^i . The state variables X_t is assumed to be a semimartingale. The coefficients $a^i(t, T)$ and $b^i(t, T)$ 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 L^i at default, consistent with an assumption of constant recovery of pre-default market value $1 - L^i$. Assuming that the bond then becomes risk-free and still trades, equation (4.13) describes the dynamics of bond prices before as well as after default.

Following Connor's definition of insurability, we assume the existence of diversified portfolios with the dynamics

$$(4.14) \quad \frac{dq^i(t, T)}{q^i(t-, T)} = (r_t + b^i(t, T)) dt + a^i(t, T) dX_t, \quad i \in Z_{++}.$$

These portfolios have the same exposure to systematic risk as individual bonds but without the firm-specific default event risks. As a result, investors would always hold diversified portfolios and individual default risk would not be compensated. Mathematically, under the assumption that there exists a pricing measure Q in an economy with an infinite

¹⁰ Without the conditional independence assumption we would have to assume that the bond return is affected not only by its own default but by the default event of other issuers.

collection of assets specified by (4.13) and (4.14), 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. This argument is made precise in the proof of the following proposition.

PROPOSITION 4.2. *Assume that the economy consists of a money market account with short rate r_t and an infinite collection of traded securities with dynamics specified in (4.13) and (4.14). Then the Q -intensity $\tilde{\lambda}_t^i$ is equal to the P -intensity λ_t^i for all i .*

Proof. Without loss of generality, we assume $L^i = 1$. Let $Y_t^i = \int_0^t b^i(u, T) du + \int_0^t a^i(u, T) dX_u$. From (4.14),

$$(4.15) \quad \frac{q^i(t, T)}{B(t)} = \mathcal{E}(Y_t^i),$$

where $B(t) = \exp(\int_0^t r_u du)$ is the money market account and $\mathcal{E}(\cdot)$ is the Doléans-Dade exponential operator. Similarly, from (4.13),

$$(4.16) \quad \begin{aligned} \frac{v^i(t, T)}{B(t)} &= \mathcal{E}(Y_t^i - M_t^i) \\ &= \mathcal{E}\left(Y_t^i - \tilde{M}_t^i + \int_0^{t \wedge T} (\lambda_u^i - \tilde{\lambda}_u^i) du\right) \\ &= \mathcal{E}(Y_t^i) \mathcal{E}(-\tilde{M}_t^i) \exp\left(\int_0^{t \wedge T} (\lambda_u^i - \tilde{\lambda}_u^i) du\right), \end{aligned}$$

where $\tilde{M}_t^i = N_t^i - \int_0^{t \wedge T} \tilde{\lambda}_u^i du$ is a Q -martingale due to the existence of a Q -intensity $\tilde{\lambda}_t^i$ associated with N_t^i .

Since $q^i(t, T)/B(t)$ is a Q -martingale, $\mathcal{E}(Y_t^i)$ is a Q -martingale due to (4.15). On the other hand, $\mathcal{E}(-\tilde{M}_t^i)$ is also a Q -martingale. Let $A_t = \mathcal{E}(Y_t^i)$ and $B_t = \mathcal{E}(-\tilde{M}_t^i)$. Their product $A_t B_t$ is a Q -local martingale since

$$(4.17) \quad \begin{aligned} [A_t, B_t] &= \left[1 + \int_0^t A_{u-} dY_u^i, 1 - \int_0^t B_{u-} d\tilde{M}_u^i\right] \\ &= - \int_0^t A_{u-} B_{u-} d[Y_u^i, \tilde{M}_u^i] \\ &= 0. \end{aligned}$$

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

$$(4.18) \quad U_t V_t = U_0 V_0 + \int_0^t U_{u-} dV_u + W_t,$$

where W_t is a Q -local martingale with $W_0 = 0$.¹¹ The decomposition (4.18) should be unique because the second term, the FV component, is continuous and hence predictable. However, because $U_t V_t = v^i(t, T)/B(t)$ is itself a Q -martingale, the FV component in the

¹¹ See Dellacherie and Meyer (1982, p. 223) for details.

decomposition must vanish. This implies that V_t is a constant, and subsequently equal to one. Hence $\tilde{\lambda}_t^i = \lambda_t^i$. \square

Apparently, this notion of diversifiability is much more restrictive than our assumption of conditionally diversifiable default risk. However, it underscores our earlier statement that much stronger assumptions are needed to obtain exact equivalence.

A special example further illustrates the previous proposition. Assuming the absence of state variables, zero interest rate, and zero recovery, a differentiation of (3.5) shows

$$(4.19) \quad \frac{dv^i(t, T)}{v^i(t-, T)} = (\tilde{\lambda}^i - \lambda^i) dt - dM_t^i.$$

According to the assumption of (4.14), a diversified portfolio q^i must provide the same return but without the default event risk. Hence,

$$(4.20) \quad \frac{dq^i(t, T)}{q^i(t-, T)} = (\tilde{\lambda}^i - \lambda^i) dt.$$

Clearly, individual bond prices grow over time at rate $\tilde{\lambda}^i$ but suffer from occasional -100% returns from default. Meanwhile, q^i represents a diversified portfolio not affected by individual default events. Its rate of return must equal the risk-free rate of zero, implying that $\tilde{\lambda}^i = \lambda^i$. Its value stays constant over time because the growth in bond value is exactly offset by average losses in the portfolio.

4.3. Necessary Conditions in a Large Economy and Asymptotic Equivalence

The previous subsection discusses the sufficient conditions for an exact equivalence between the martingale and empirical default intensities. We acknowledge that these conditions are quite restrictive. In a finite economy, it would be a coincidence if the default of one bond could be perfectly hedged by a portfolio of other bonds. Even in a limit economy, one still may not have the ability to form the type of diversified portfolio in (4.14).

We present a more precise description of the sense in which the empirical and martingale intensities must be approximately the same for “most” assets. Since no additional structure is imposed apart from that of Section 3.1, we consider this subsection the main result of the paper.

A number of studies have rigorously addressed the issue of asymptotic arbitrage in dynamic models, including Kabanov and Kramkov (1998), Björk and Näslund (1998), and Klein and Schachermayer (1997). Our approach is similar to that of Björk and Näslund in that we work directly in an economy with an infinite number of assets. However, we allow for more general dynamics. In the following discussion we will relate our results to the other two papers.

Consider the economy formally defined in Section 3.1 and assume that there is a single state variable given as an Ito process:

$$(4.21) \quad dX_t = \mu(t) dt + \sigma(t) dW_t,$$

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 3, 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 (3.3) and (3.4). 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 subeconomy, 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

$$(4.22) \quad dX_t = (\mu(t) + g(t) \sigma(t)) dt + \sigma(t) d\tilde{W}_t,$$

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

$$(4.23) \quad E^P \left(\exp \left(\int_0^{T^*} g(u) dW_u - \frac{1}{2} \int_0^{T^*} g^2(u) du \right) \right) = 1.$$

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 .

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 $Y: \Omega \times [0, T^*] \times \mathbb{N} \rightarrow \mathbb{R}$ 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 Mémmin (1976).

In the setup of our economy, we have the following proposition.

PROPOSITION 4.3. *Under an equivalent measure Q , the intensities of the one-jump processes are given as*

$$(4.24) \quad \tilde{\lambda}_t^i = Y(t, \omega, i) \lambda_t^i,$$

for a strictly positive, predictable function Y that satisfies

$$(4.25) \quad \int_0^{T^*} \sum_{i=1}^{\infty} \left(1 - \sqrt{Y(u, \omega, i)} \right)^2 \lambda_u^i du < \infty.$$

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

$$(4.26) \quad S_t = X_t + \sum_{i=1}^{\infty} a_i 1_{\{\tau^i \leq t\}}.$$

Let μ denote the (random) jump measure associated with this process; that is,

$$(4.27) \quad \mu([0, t] \times i) = 1_{\{\tau^i \leq t\}}, \quad i \geq 1,$$

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

$$(4.28) \quad \begin{aligned} d\alpha_t &= \mu(t) dt + \sum_{i=1}^{\infty} a_i \lambda_t^i dt, \\ d\beta_t &= \sigma^2(t) dt, \quad \text{and} \end{aligned}$$

$$\nu(dt, \{i\}) = a_i \lambda_t^i dt.$$

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

$$(4.29) \quad X_t = S_t - \sum_{u \leq t} \Delta S_u,$$

and

$$(4.30) \quad N_t^i = 1_{\{\Delta S_u = a_i \text{ for some } u \leq t\}}.$$

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 Mémmin 1976), we have that the characteristics under Q are given as

$$(4.31) \quad \begin{aligned} d\tilde{\alpha}_t &= d\alpha_t + g(t)\sigma(t) dt + \sum_{i=1}^{\infty} a_i (Y(t, i) - 1) \lambda_t^i dt, \\ d\tilde{\beta}_t &= d\beta_t, \quad \text{and} \end{aligned}$$

$$\tilde{\nu}(dt, \{i\}) = Y(t, i)\nu(dt, \{i\}).$$

This follows from Theorem 3.3 of Jacod and Mémmin (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 Mémmin (1976) is equivalent to the condition that

$$(4.32) \quad \sum_{i=1}^{\infty} \int_0^{T^*} |Y(u, i) - 1| 1_{\{Y > 2\}} \lambda_u^i du + \int_0^{T^*} (Y(u, i) - 1)^2 1_{\{Y \leq 2\}} \lambda_u^i du < \infty, \quad P\text{-a.s.}$$

But using the inequality

$$(4.33) \quad (1 - \sqrt{y})^2 \leq (y - 1)^2 1_{\{y \leq 2\}} + |y - 1| 1_{\{y > 2\}} \leq \frac{1}{(1 - \sqrt{2})^2} (1 - \sqrt{y})^2,$$

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

$$(4.34) \quad \sum_{i=1}^{\infty} \int_0^{T^*} \left(1 - \sqrt{Y(u, i)}\right)^2 \lambda_u^i du < \infty,$$

as was to be proved. \square

Expression (4.25) is similar, in spirit, to the sufficient condition for diversification given in the original Ross (1976) diversification result. One should think of it 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 subeconomy there can be perturbations in risk premia due to defaults of other firms (as in Jarrow and Yu 2001 and Kusuoka 1999) since the process $Y(\cdot, \cdot, i)$ may depend on the jump times of firms other than the i th. However, such a counterparty dependence must “die out” asymptotically in the infinite economy if default risk is conditionally diversifiable under the empirical measure. Stated differently, the change in risk premium induced by one firm’s default may only have an effect (over a certain value) on a finite segment of the economy.

The conditional diversification construction is used for two reasons: First, it facilitates the construction of and calculations related to jump processes, which are driven by exogenous state variables. Violating this condition quickly leads to complications as demonstrated by the looping default example in Jarrow and Yu (2001).

Second, it ensures that there are no simultaneous jumps under P of the infinite collection of jump times, thus playing a role similar to the assumption of independent noise terms in the classical formulation of APT. We could set up a model where finite clusters of defaults are interlinked under P , just as we can have nondiagonal covariance matrices in APT. But the model is then very messy to write out. We have chosen a simple starting point with conditional independence, and the equivalence result then shows us the degree of perturbation of this property that is possible under an equivalent change of measure. It is clear that under an equivalent change of measure we cannot introduce simultaneous defaults. Hence the direct contagion cannot exist under Q if it does not exist under P already. If we want a jump triggering an infinite collection of defaults under P , then we need to let that jump be part of the state variable process X , and it will be possible to change the intensity of this state variable (thus affecting an infinite collection of intensities simultaneously but through the effect of only one intensity).

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 subeconomies 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 and the gains process is always nonnegative 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 subeconomy 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 is 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 subeconomies in the sequence. Indeed, they can be constructed on different probability spaces altogether. However, working with this construction in our setting

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.

does not produce intuitive results unless very special structures are imposed because the structure of risk premia in the $(n + 1)$ th economy can be completely unrelated to that of the n th economy. Nevertheless, one could link our construction with the theory of asymptotic arbitrage by including the first n default processes in the n th economy and assuming that there is a large N with the following property: for $n > N$, the predictable function Y^{n+1} defined on $\Omega \times [0, T^*] \times \{1, 2, \dots, n + 1\}$ is equal to Y^n on the restriction to $\Omega \times [0, T^*] \times \{1, 2, \dots, n\}$. Using this construction, we are able to view element n of the sequence of subeconomies 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 considerably the proof of no asymptotic arbitrage, since the critical condition of contiguity becomes merely a condition of absolute continuity of the unrestricted measures. We are then back to the condition in Proposition 4.3.

5. SOME EMPIRICAL IMPLICATIONS OF CONDITIONALLY DIVERSIFIABLE DEFAULT RISK

In this section we use numerical examples to illustrate some empirical implications of conditionally diversifiable default risk. Notably, we show that the risk adjustment through the state variable is able to produce upward-sloping corporate bond yield spread curves even if the underlying credit class has decreasing conditional default probabilities.¹³ We also show how to infer the physical default rates from corporate bond prices. We complement this discussion by showing that the diversifiable risk argument is already implicitly used in empirical prepayment models for pricing mortgage-backed securities.

5.1. Corporate Bonds

5.1.1. 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. 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 interest rate swaps.

For the purpose of illustrating this methodology, we assume in this section that the evolution under P of the default intensity process estimated by Duffee (1999) from corporate bond prices is the empirical default intensity of the default event. This allows us to illustrate two points: First, 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 capable of producing a large gap between empirical and martingale (implied) default probabilities. Hence, the gap between these probabilities is not a reason for ruling out the relevance of estimating default intensities empirically and using them for pricing. Second, the default risk premium is capable of explaining the well-documented empirical feature that

¹³ Although this feature does not specifically rule out the existence of compensation for jump risk, it would be difficult to obtain this feature with jump risk premia as the primary risk adjustment.

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 are: $\alpha = 0.00961$, $\beta_1 = -0.171$, $\beta_2 = -0.006$, $\kappa = 0.212$, $\theta = 0.00628$, and $\sigma = 0.059$. 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 initial values for the short-rate factors and their averages ($\bar{s}_{1,t}$ and $\bar{s}_{2,t}$) are set to their long-run mean values under P , and the current value for h_t^* is set equal to the mean fitted value over Duffee's sample, 0.00864.

Figure 5.1 shows the term structure of default probabilities, under both the physical measure and under risk-adjusted measures, for maturities up to 30 years. The $\lambda = 0$ series represents a case with no risk adjustment (risk-neutrality) on the h_t^* factor and no risk adjustment on the Treasury rate factors s_{1t} and s_{2t} . This represents the empirical default probabilities if the intensity process under the physical measure P in Duffee (1999) is the intensity of the default event under P . The $\lambda = -0.307$ series is the intensity under the risk-neutral measure 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 large differences in long-term actual and implied default probabilities.

The second, and more important, insight can be gleaned from Figure 5.2, 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 (3.5) to compute the spreads. In the case of risk neutrality on the intensity factor (a case that closely approximates overall risk neutrality because the influence from the Treasury factors is very small), 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 et al. (1997) for lower-rated firms under risk neutrality and zero recovery assumptions. However, Figure 5.2 also shows that for the other two cases of nontrivial 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 Although it is still a controversial

¹⁴ In Duffee (1999), the default intensity is assumed to be

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

where α , β_1 , and β_2 are constants, and 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). Under P , h_t^* is a square-root diffusion

$$dh_t^* = \kappa (\theta - h_t^*) dt + \sigma \sqrt{h_t^*} dZ_t,$$

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

$$dh_t^* = (\kappa\theta - (\kappa + \lambda)h_t^*) dt + \sigma \sqrt{h_t^*} d\tilde{Z}_t,$$

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

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 the 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.




FIGURE 5.1. Survival probabilities are computed under the physical measure (labeled “ $\lambda = 0$ ”) obtained by setting all risk adjustment parameters equal to zero, both for the factors driving Treasury rates and the default intensity factor. Survival probabilities are also computed using the same process for the default intensity, but with risk adjustment in the factors driving the Treasury rates set as in Duffee (1999) and with two different values for the risk adjustment λ of the default intensity factor.

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.

5.1.2. From Martingale Intensity to Empirical Default Probabilities. The second important implication of diversifiable default risk is the link between the martingale intensity estimated from market prices and the default probabilities needed for computing VaR measures in risk management. To illustrate this procedure, 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 under the physical measure P to obtain the default probabilities.

Since any error in the survival probabilities in the short-end will manifest itself through the entire time horizon, to get a more precise picture we consider instead conditional

This feature, which manifests itself in a negative but close to zero mean reversion parameter for the martingale intensity, is confirmed by other studies such as that of Liu et al. (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 Aaa credit spread curve according to Duffee (1999).




FIGURE 5.2. Term structures of yield spreads for different values of the risk adjustment on the default intensity factor h^* . Here, $\lambda = 0$ corresponds to having risk neutrality with respect to the default intensity factor, and due to the small influence of riskless bond factors, this closely approximates overall risk neutrality. The downward sloping shape is a consequence of decreasing conditional default probabilities under the physical measure. The risk adjustment allows upward sloping yield curves and downward sloping conditional default probabilities to coexist.

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 $1 - p_n/p_{n-1}$. We control for liquidity and tax components of the credit spread by subtracting the sample average of the Aaa intensity, 0.00931, from the default intensities of other credit rating classes.¹⁸ For comparison purposes, we contrast these “implied” conditional default probabilities with those obtained from a one-year transition matrix from Moody’s (as reported on the CreditMetric home page, November 1999).

Figure 5.3 shows that unlike the “actual” series that are upward-sloping throughout the range of maturities, our “JLY” series is quite flat, eventually matching the “actual” series at longer maturities. This flatness is partly due to the fact that we set the initial value of the h^* 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,

¹⁸ We assume that the sample means of the two Treasury rate factors are equal to their long-run means under P . As a result, the sample mean of the Aaa intensity is the sum of Duffee’s α estimate, 0.00594, and his mean fitted h^* , 0.00337. We cannot use the intensity process for Aaa bonds because the estimated parameters were noted by Duffee to be unstable.




FIGURE 5.3. Conditional default probabilities for investment-grade issuers computed by assuming conditional diversifiability (JLY-series) and using Moody’s one-year transition matrix (actual). Here we subtract 93 bps in the intensity, based on Duffee’s estimates for Aaa spreads, proxying for a pervasive component in corporate bond spreads due to nondefault related issues.

the upward-sloping shape cannot be reproduced. This is because under the empirical measure the state variables mean-revert “too quickly” (e.g., for an Aa issuer the half life of the h^* factor is less than 4 years).

Subtracting the sample average of Aaa spreads brings the actual and implied conditional default probabilities closer to each other in the short end for Baa bonds but less so for the higher ratings. Note, however, that this simple procedure still will not match the shape of the two series because a constant adjustment is likely to overcompensate for the effect of nondefault factors in the long-end. It is a first priority, then, for future empirical credit research to quantify the effect of nondefault factors on the term structure of yield spreads and to extract a default intensity function that is “uncontaminated” by these factors.

5.2. 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.

6. 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.

An important implication of our equivalence result is that it integrates pricing and risk management for defaultable securities. This has two meanings. First, when the diversifiability conditions hold, 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 those of 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 use of this methodology in the context of existing empirical studies on corporate bonds.

For further implications, 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.

REFERENCES

- ALTMAN, E. I. (1968): Financial Ratios, Discriminant Analysis, and the Prediction of Corporate Bankruptcy, *J. 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*. New York: Wiley.
- ARTZNER, P., and F. DELBAEN (1995): Default Risk Insurance and Incomplete Markets, *Math. Finance* 5, 187–195.
- BACK, K. (1991): Asset Pricing for General Processes, *J. Math. Econ.* 20, 371–395.
- BJÖRK, T., and B. NÄSLUND (1998): Diversified Portfolios in Continuous Time, *European Finance Rev.* 1, 361–387.
- COLLIN-DUFRESNE, P., R. GOLDSTEIN, and J. HELWEGE (2002): Are Jumps in Corporate Bond Yields Priced? Modeling Contagion via the Updating of Beliefs. Working paper, Carnegie Mellon University.
- COLLIN-DUFRESNE, P., R. GOLDSTEIN, and J. HUGONNIER (2003): A General Formula for Pricing Defaultable Securities. Working paper, Carnegie Mellon University.
- CONNOR, G. (1984): A Unified Beta Pricing Theory, *J. Econ. Theory* 34, 13–31.

- COX, J., J. INGERSOLL, and S. ROSS (1985): 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, *J. Finance* 55, 1943–1978.
- DELLACHERIE, C., and P. MEYER (1982): *Probabilities and Potential B: Theory of Martingales*. Amsterdam: North-Holland.
- DRIESSEN, J. (2002): Is Default Event Risk Priced in Corporate Bonds? Working paper, University of Amsterdam.
- DUFFEE, G. R. (1999): Estimating the Price of Default Risk, *Rev. Financial Stud.* 12, 197–226.
- DUFFIE, J. D., and R. KAN (1996): A Yield-Factor Model of Interest Rates, *Math. Finance* 6, 379–406.
- DUFFIE, J. D., and K. J. SINGLETON (1997): An Econometric Model of the Term Structure of Interest Rate Swap Yields, *J. Finance* 52, 1287–1321.
- DUFFIE, J. D., and K. J. SINGLETON (1999): Modeling Term Structures of Defaultable Bonds, *Rev. Financial Stud.* 12, 687–720.
- GORDY, M. (2003): A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules, *J. Financial Intermed.* 12, 199–232.
- HELWEGE, J., and C. M. TURNER (1999): The Slope of the Credit Yield Curve for Speculative Grade Issuers, *J. Finance* 54, 1869–1884.
- JACOD, J., and J. MÉMIN (1976): Caractéristiques locales et conditions de continuité absolue pour les semimartingales, *Z. Wahrsch. verw. Gebiete* 35, 1–37.
- JARROW, R. A. (1988): Preferences, Continuity, and the Arbitrage Pricing Theory, *Rev. Financial Stud.* 1, 159–172.
- JARROW, R. A., D. LANDO, and S. M. TURNBULL (1997): A Markov Model for the Term Structure of Credit Risk Spread, *Rev. Financial Stud.* 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, *Math. Finance* 5, 311–336.
- JARROW, R. A., and S. M. TURNBULL (1995): Pricing Derivatives on Financial Securities Subject to Credit Risk, *J. Finance* 50, 53–85.
- JARROW, R. A., and F. YU (2001): Counterparty Risk and the Pricing of Defaultable Securities, *J. Finance* 56, 1765–1799.
- KABANOV, Y. M., and D. O. KRAMKOV (1998): Asymptotic Arbitrage in Large Financial Markets, *Finance Stoch.* 2, 143–172.
- KLEIN, I., and W. SCHACHERMAYER (1997): Asymptotic Arbitrage in Noncomplete Large Markets, *Theory Prob. Appl.* 41, 780–788.
- KUSUOKA, S. (1999): A Remark on Default Risk Models, *Adv. Math. Econ.* 1, 69–82.
- LANDO, D. (1994): Three Essays on Contingent Claims Pricing. Ph.D. thesis, Cornell University.
- LANDO, D. (1998): On Cox Processes and Credit Risky Securities, *Rev. Derivat. Res.* 2, 99–120.
- LIU, J., F. LONGSTAFF, and R. MANDELL (2001): The Market Price of Credit Risk: An Empirical Analysis of Interest Rate Swap Spreads. Working paper, UCLA.
- ROSS, S. (1996): The Arbitrage Theory of Capital Asset Pricing, *J. Econ. Theory* 13, 341–360.
- SCHWARTZ, E., and W. TOROUS (1989): Prepayment and the Valuation of Mortgage-Backed Securities, *J. Finance* 44, 375–392.

- SHUMWAY, T. (2001): Forecasting Bankruptcy More Accurately: A Simple Hazard Model, *J. Business* 74, 101–124.
- STANTON, R. (1995): Rational Prepayment and the Valuation of Mortgage-Backed Securities, *Rev. Financial Stud.* 8, 677–708.
- VASICEK, O. (1977): An Equilibrium Characterization of the Term Structure, *J. Financial Econ.* 5, 177–188.
- YU, F. (2003): Dependent Default in Intensity-Based Models. Working paper, UC-Irvine.