Forecasting Default in the Face of Uncertainty

Kay Giesecke and Lisa R. Goldberg

Kay Giesecke
is a visiting assistant professor at the School of Operations Research and Industrial Engineering at Cornell University in Ithaca, NY.
giesecke@orie.cornell.edu

Lisa R. Goldberg
is vice president of Credit Research at MSCI Barra, Inc., Berkeley, CA.
lrg@barra.com

In our structural credit model based on incomplete information, investors cannot observe a firm's default barrier. As a consequence, such a model has both the economic appeal of a structural model and the tractable pricing formulas and empirical plausibility of a reduced-form model.

A comparison of default probability and credit spread forecasts generated by this model and two well-known structural models indicates that it reacts more quickly to new information and, unlike the other two models, it forecasts positive short-term credit spreads.

Quantitative credit risk models have become central to the investment process in today's very large and complex credit markets. There are two main types of quantitative credit models. The *structural* or *cause and effect* approach is based on a model definition of default. An alternative *intensity-based* or *reduced-form* model directly addresses the difficulty of fitting short-term credit spreads.

The first and best-known structural model is introduced in Merton [1974]. Here, firm debt is a single pure discount bond, and default occurs only if firm value falls below the face value of the bond at maturity.

The fundamental insight here is that firm debt can be viewed as a portfolio composed of a risk-free bond and a short put option on the value of the firm. Model credit spreads can be estimated with the European option formulas in Black and Scholes [1973].

Attempts to calibrate the Merton model led naturally to modifications and extensions. By empirical standards, Merton model credit spreads are too low. This issue is addressed by Black and Cox [1976], who argue that default can happen at any time. Bondholders may have the right to force Chapter 11 as soon as firm value falls to some prescribed lower threshold, or insiders may try to restructure a firm when they deem it optimal.

This line of reasoning led to the development of a *first-passage model*, where a firm defaults when its value falls to a barrier. Spreads forecast by a classic first-passage model are higher than in the Merton model but still an imperfect fit to empirical data.

In both the Merton and the Black-Cox models, the credit spread tends rapidly to zero as the term declines to zero. This is an unintended consequence of an assumption underlying the Merton, Black-Cox, and many other structural models. In these models, default is a predictable event. Here, *predictable* is a precise mathematical term; it means the event is foreshadowed by an observable phenomenon such as the value of the firm falling close to its default boundary.[1]

An immutable consequence of predictability is that the model forecasts zero short-term credit spreads. In other words, it forecasts that over a short term, a junk bond is no riskier than a U.S. Treasury.

In reality, default often comes without warning. This is reflected in the credit market by the prevalence of positive short-term credit spreads. If default were truly predictable, the empirical term structure of credit spreads would drop to zero as term shortens to zero.

An alternative *intensity-based* or *reduced-form* approach directly addresses the difficulty of fitting short spreads. It is considered by Artzner and Delbaen [1995], Duffie and Singleton [1999], and Jarrow and Turnbull [1995], among others. The reduced-form approach assumes that default hits the market by surprise. Each firm has an *intensity* or *conditional rate of default*.

The intensity $\lambda(\omega,t)$ expresses the probability that a firm will default in the next instant, given that a firm escapes default until time t and the state of the world is ω. From the intensity, default probabilities can be easily calculated. Simple extensions of reduced-form models that include loss given default can be used to price credit-sensitive securities.

It seems natural to integrate both approaches so as to gain both the economic appeal of the structural approach and the empirical plausibility and the tractability of the intensity-based approach. Yet integration is not a straightforward exercise. Most structural models posit that default is a predictable event. Reduced-form models, however, assume default is *unpredictable*.

A successful integration of structural and reduced-form models comes from the observation that most investment decisions are based on incomplete information. The first incomplete information model is introduced in Duffie and Lando [2001]. They retain the first-passage time definition of default but assume investors observe the true value of a firm imperfectly. This makes the default event unpredictable and leads to a hybrid structural-reduced-form model with an endogenously defined intensity.

Other information scenarios for first-passage models can be envisioned. For example, we can assume that investors observe firm value but not the default barrier. Alternatively, we can assume that neither firm value nor the default barrier can be observed. These examples can be analyzed in the framework described in Giesecke [2001]. This framework also allows for a wide range of default definitions, such as firm value crossing a barrier multiple times, or spending a sufficiently long time below a barrier.

Incomplete information models can incorporate the following features:

- Structure plus short-term uncertainty: Incomplete information provides a framework to coherently integrate a cause-and-effect model of default with the short-term uncertainty that surrounds default events.
- Economic reasonability and flexibility: Incomplete information models account for the fact that investors may not have an accurate picture of a firm's value or liabilities.
- Tractable pricing formulas and empirical plausibility: Many incomplete information models have reduced-form pricing formulas and are easily calibrated to market data.
- Unified perspective: The incomplete information framework gives a common perspective on intensity-based models and traditional structural models. Hence, these models can be directly compared.
- New possibilities: The incomplete information framework generates previously unseen models.

We have two purposes in writing this article. The first is to elucidate the theory of incomplete information credit models that is developed in Giesecke [2001]. We show how to construct such a model based on a simplifying set of assumptions. Our first two sections are for the mathematically inclined reader who wants to experiment with incomplete information models.

The second goal is to document and empirically assess I^2, which is an incomplete information model based on a first-passage definition of default. In I^2, investors cannot observe a firm's default barrier so that, prior to default, the distance to default is unknown in all states of the world.

Some first-passage models such as Black and Cox [1976] and Longstaff and Schwartz [1995] assume the default barrier is a deterministic function of time. Others, including Nielsen, Saá-Requejô, and Santa-Clara [1993], are based on a random barrier. Implicit in all these models is the assumption that the barrier, whether it is fixed or random, can be observed by investors. In each case, it follows that the distance to default is observable and default is predictable.

The I^2 default barrier by contrast is both random and unobserved. This is consistent with recent experiences at Enron, WorldCom, and Tyco, which surely show that investors cannot observe the barrier. In each of these cases, the true level of liabilities was not disclosed to the public.

At first glance, the "CreditGrades" first-passage model, described in the RiskMetrics technical document [2002], seems similar to I^2. It is predicated on an observable firm value process and an unobservable random default barrier. "CreditGrades," however, approximates the unobservable distance to default with an observable process, resulting in a predictable default.

EXHIBIT 1
Default Indicator for Four States

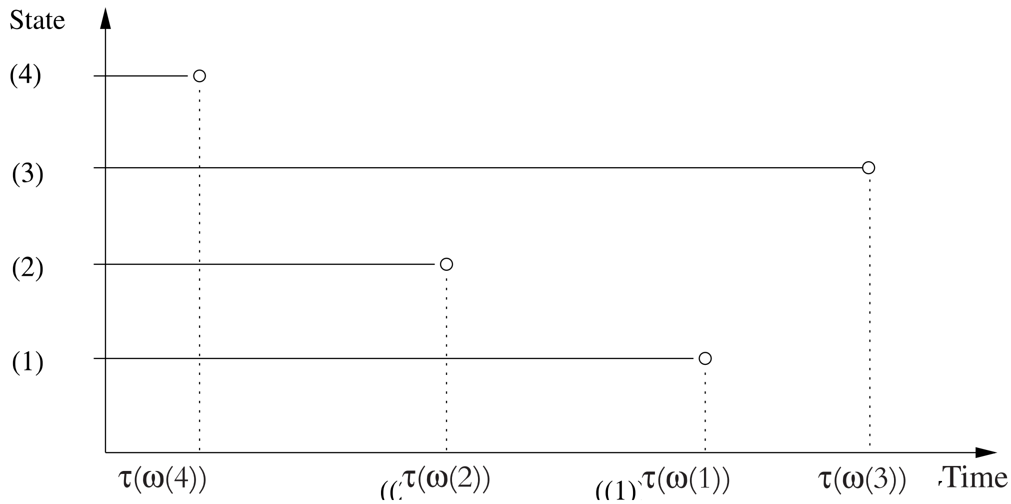

The incomplete information model in Duffie and Lando [2001] implicitly defines an intensity, which generates reduced-form formulas for default probabilities and prices of credit-sensitive securities. As we demonstrate, I^2 does not have an intensity, but does have reduced-form pricing formulas that are analogous to the intensity-based formulas.

At first, this might seem impossible since the intensity seems central to the reduced form. In every instance, the intensity appears as the integrand of a time integral. In other words, only the cumulative intensity or *pricing trend* is required for model forecasts. From the perspective of credit modeling, the pricing trend does not need to be the integral of an intensity. This means that the trend does not need to have a meaningful time derivative.

If we broaden our scope to include trends that do not arise from intensities, we gain access to a class of generalized reduced-form models, of which I^2 is an example. If the trend happens to have a meaningful time derivative, the familiar reduced-form results are obtained.

We compare the I^2 and Merton and Black-Cox default probability forecasts. There are many similarities but also important differences including:

- The I^2 model reacts more quickly than the other models since it takes direct account of the entire history of public information rather than just current values.
- The I^2 model forecasts positive short spreads for firms whose leverage ratio is at an historic high. The Merton and Black–Cox models always predict that short spreads are zero.
- The I^2 default barrier parameters are intuitive and can be calibrated to the quality of information available to investors.

I. A COMMON PERSPECTIVE

To understand what is common to structural and reduced-form credit models, we need to start at the most rudimentary level. At the core of every credit model is a *default time*. This is a random variable $\tau(\omega)$ whose value depends on the state of the world ω.[2]

Of course, we do not know the default time. We use a credit model to estimate from observable information the distribution of the default time and prices of credit-sensitive securities. To explain how this works, it is helpful to recast the content of the default time as a process. Let $N(t, \omega)$ be an indicator that is zero until $t = \tau(\omega)$ at which time it jumps to one and stays there. Exhibit 1 shows a schematic view of the default process N. The exhibit shows the default indicator for four states, $\omega(1)$, $\omega(2)$, $\omega(3)$, $\omega(4)$, indexed along the vertical axis. The horizontal axis measures time and τ is the time of default. For example, if the state of the world is $\omega(2)$, then default occurs at time $\tau(\omega(2))$. State $\omega(4)$ could be considered the worst, since default occurs earliest.

We describe the default process in terms of observable information. Let $E_t[X](\omega)$ be the conditional expectation of the random variable X, given the information that can be observed at time t if the world is in state ω. We introduce the *conditional default probability*

$$F(t,\omega) = E_t[N(t)](\omega) \tag{1}$$

Note that the conditional expectation at time t depends on which events can be observed at t. For readability, we do not specify this in the notation, but it is important nevertheless. In some cases, the information set is so refined that the indicator $N(t)$ is observable at time t. In this case, the expectation in Equation (1) drops out and, and $F(t, \omega)$ simplifies to $N(t, \omega)$. This occurs in examples I.1 and I.2 below. If the information at time t is too coarse for $N(t)$ to be observable, the expectations operator does not drop out. This occurs in Examples I.3 and I.4.

Equation (1) defines a stochastic process whose value at time t in state ω is the probability that a firm will default before time t, conditional on information that can be observed at time t in state ω. Note that

$$E[F(t)] = E[N(t)] = P[\tau \leq t] \tag{2}$$

Therefore, the conditional default probability can be used to estimate the probability of default. We construct the conditional default probability in several different contexts.

Example I.1. In the Merton model, investors observe the continuous value V of the firm as it evolves. They also observe the face value D and maturity T of a pure discount bond that represents firm debt in its entirety. Default occurs only at T in states ω for which $V(T, \omega) \leq D$. Since both V and D can be observed by investors:

$$F(t,\omega) = \begin{cases} 0 & \text{if } t < T \\ 0 & \text{if } t = T \text{ and } V(T,\omega) > D \\ 1 & \text{otherwise.} \end{cases} \tag{3}$$

Example I.2. In the Black-Cox model, investors observe the continuous firm value once again. Default occurs when firm value first hits an observable barrier. Therefore, if $M(t, \omega) = \min_{0\leq s\leq t} V(s, \omega)$ denotes the historical low of firm values:

$$F(t,\omega) = \begin{cases} 0 & \text{if } M(t,\omega) > D \\ 1 & \text{otherwise.} \end{cases} \tag{4}$$

Example I.3. In a reduced-form model, the cause of default is not specified. Instead, observable information is used to specify a conditional default rate or intensity $\lambda(t, \omega)$ and the conditional default probability is

$$F(t,\omega) = 1 - \exp\left(-\int_0^t \lambda(t,\omega)ds\right) \tag{5}$$

by construction.

Example I.4. In the structural I^2 model, investors observe the continuous firm value V. As in the Black-Cox model, default occurs when firm value first hits a barrier. In reality, however, the barrier cannot be observed by investors. Since they have no better alternative, investors apply their view of the likelihood that any particular value may be the default barrier.

Therefore, in I^2 we model the barrier as an unobservable and continuous random variable on $(0, V,(0))$ whose distribution function is denoted by G. The conditional default probability at time t is equal to the probability that the minimum firm value $M(t, \omega)$ seen by time t in state ω is below the default barrier. Since at any future time t, the information about the minimum firm value witnessed by time t is known, we get for the I^2 conditional default probability

$$F(t,\omega) = 1 - G(M(t,\omega)), \qquad t > 0 \tag{6}$$

assuming the barrier is independent of firm values.

The examples above illustrate the hybrid nature of the I^2 model.

- The Merton, Black-Cox, and I^2 models are structural: They model the cause of default. The I^2 conditional default probability is qualitatively different. In the Merton and Black-Cox models, $F(t, \omega)$ is equal to the Bernoulli default indicator process $N(t, \omega)$. In I^2, $F(t, \omega)$ is much richer than its default indicator. The difference arises from the information assumptions underlying the models. In the Merton and Black-Cox models, investors can deduce the distance to default from firm fundamentals. It follows that the conditional default probability is equal to the default process. In the I^2 model, investors cannot deduce the distance to default from firm fundamentals.

- The reduced-form and I^2 models are based on unpredictable default times: Default comes as a surprise and cannot be anticipated. Remarkably, this accounts for the fact that the conditional default probabilities for I^2 and the reduced-form model have continuous paths. Conversely, the jumps in the Merton and Black-Cox conditional default probabilities come from their *predictable* default times. This is an indication of the mathematical relationship between the probabilistic properties of the default time and the analytic properties of the conditional default probability.

The unpredictability of the I^2 default time has significant empirical implications that we explore later. It also has ramifications with regard to pricing credit-sensitive securities. For the Black-Cox and Merton models, the conditional default probability is a zero-one indicator. It is not a useful security pricing tool. For unpredictable models such as I^2, the conditional default probability leads to reduced-form pricing formulas. For the I^2 and reduced-form models, we introduce the non-decreasing and continuous process A, called the pricing trend:[3]

$$A(t,\omega) = -\log(1 - F(t,\omega)) \tag{7}$$

Example I.5. In a reduced-form model, the pricing trend is, by construction, given by the cumulative intensity:

$$A(t,\omega) = \int_0^t \lambda(s,\omega)ds \tag{8}$$

Example I.6. In the I^2 model, the pricing trend is

$$A(t,\omega) = -\log G(M(t,\omega)) \tag{9}$$

The I^2 pricing trend does not come from an intensity h. If it did, the trend could be expressed as

$$A(t,\omega) = \int_0^t h(s,\omega)ds \tag{10}$$

and h would be equal to the derivative of A with respect to time. Under the assumption that the barrier distribution function G is differentiable, the derivative of $A(t, \omega)$ with respect to time t is zero in almost all states w. This means that h is zero and Equation (10) does not hold.

If the pricing trend has continuous paths as in the reduced-form and I^2 models, it can often be used to estimate prices of credit-sensitive securities. Consider, for example, a corporate zero-coupon bond that promises to pay one dollar at some future time T. Suppose investors recover nothing at default.

Let $r(t, \omega)$ denote the risk-free interest rate. If the process p defined by

$$p(t,T) = E_t[e^{-\int_t^T r(s)ds + A(t) - A(T)}], \quad t \leq T \tag{11}$$

is continuous at the default time, the price of the bond at any time $t < \tau$ is given by $p(t, T)$. This condition is satisfied for the I^2 model specification we describe below but not for all reduced-form models.

The pricing formula in Equation (11) gives a formula for the default probability as a special case. Setting $r=0$, we get

$$P[\tau \leq t] = 1 - p(0,t) = 1 - E[e^{-A(t)}] \tag{12}$$

This result is not surprising, given Equation (2) and the relation in Equation (7) between the conditional default probability and the pricing trend.

In reality, bond investors usually do not lose everything at default. Suppose they receive a fraction $R \in [0,1]$ of the bond's market value just before default. Mathematically, this value is $(\tau-, T) = \lim_{t\uparrow\tau} p(t,T)$. This recovery model follows Duffie and Singleton [1999], who introduce it in the context of reduced-form credit models. Giesecke and Goldberg [2003] show that if the process p defined by

$$p(t,T) = E_t[e^{-\int_t^T [r(s)ds + (1-R_s)dA(s)]}], \quad t \leq T \tag{13}$$

is continuous at the default time, then the price of the fractional recovery bond at any time $t < \tau$ is given by $p(t, T)$.

II. COMPENSATORS

The definition in Equation (7) of the reduced-form and I^2 pricing trends extends to any credit model whose conditional default probability is monotone and continuous, and always satisfies $F(t, \omega) < 1$.[4] As long as $F(t, \omega) < 1$, the formula $\log(1-F(t, w))$ given in (7) is mathematically well defined. It cannot be used in (13) to price securities if $F(t, \omega)$ fails to be either continuous or monotone.

Here are two credit models where this issue arises:

Example II.1. Suppose that I^2 is modified so the firm value process has jumps. Then the minimum firm value process M has jumps as well, so the conditional default probability is not continuous.

Example II.2. In an economy with several firms, default dependence can be modeled in I^2 through barrier relationships. When the first firm defaults, the value of its default barrier is revealed and provides new information about the level of the barriers of the remaining firms. Investors reassess the default barrier of surviving firms, and the conditional default probabilities jump. This leads to *default contagion*, which is analyzed in Giesecke and Goldberg [2004].

These and many other examples are covered by the more general definition of the pricing trend, which is given by the Stieltjes integral:

$$A(t,\omega) = \int_0^t \frac{dK(s,\omega)}{1 - F(s-,\omega)} \qquad (14)$$

see Giesecke [2001, Definition 4.1]. The notation $F(t-, \omega)$ is shorthand for the left limit ${}_{s\uparrow t}F(s, \omega)$. Here, $K(t, \omega)$ is the *compensator* of $F(t, \omega)$. The compensator is a process that is mathematically more tractable than the conditional default probability. For example, it is monotone even if the conditional default probability is not. The compensator includes all the information about default needed to price credit-sensitive securities.

A compensator can be derived from any bounded stochastic process Z, such as the conditional default probability, which is non-decreasing on average:

$$Z(s,\omega) \le E_s[Z(t)](\omega) \qquad (15)$$

for all $s \le t$.[5]

The compensator is uniquely specified by three properties: 1) The compensator is non-decreasing and starts at zero; 2) the difference between Z and its compensator is a martingale. This is a process that is fair in the sense that the expected gain or loss at every future time is zero; 3) the compensator is a predictable process, even if the underlying process Z is not. A predictable process is similar to a predictable default event. The values are always foreshadowed or "announced" by information observable at earlier times. If the underlying process Z satisfies these three properties, it is its own compensator, and the difference martingale part is zero.

Example II.3. In the reduced-form and I^2 models, the conditional default probability $F(t, \omega)$ is a continuous monotone process that starts at zero. Since a continuous process is predictable, $F(t, \omega)$ is equal to its own compensator. The general definition (14) of the pricing trend

$$A(t,\omega) = \int_0^t \frac{dF(s,\omega)}{1 - F(s,\omega)} = -\log(1 - F(t,\omega)) \qquad (16)$$

reduces to the simpler definition in Equation (7).

Example II.4. In the Merton and Black-Cox models, the default time τ is predictable. Here, the default indicator $N(t, \omega)$, the conditional default probability $F(t, \omega)$, and its compensator coincide. There is no pricing trend since there is a definite chance that $F(t, \omega) = 1$ at a finite time.

Example II.5. Suppose I^2 is modified so that firm value is given by $V(t, \omega) = V(0)\exp[-U(t, \omega)]$ where $U(t, \omega)$ is a Poisson process with intensity α. Suppose further that the default barrier is uniformly distributed. From (6), we have $F(t, \omega) = 1 - \exp[-U(t, \omega)]$ which is monotone but not continuous. The associated compensator can be derived from Meyer's [1966] representation theorem. It is given by

$$K(t,\omega) = \alpha(1 - e^{-1}) \int_0^t e^{-U(s,\omega)} ds \qquad (17)$$

The pricing trend defined by the more general Equation (14) can be used to price securities as in Equation (13).

Exhibit 2 is a flowchart for the construction of an incomplete information credit model based on an unpredictable default. Note that there are two compensators in the diagram. The compensator of the conditional default probability that we discuss here is equal to the conditional default probability $F(t, \omega)$ for I^2 and reduced-form models. In general, it is derived from $F(t, \omega)$ with the Doob-Meyer theorem.

A compensator depends on both an underlying process and an information set. The compensators of the conditional default probabilities described above implicitly use information observable within the context of the model, such as minimum firm value for the Black-Cox

EXHIBIT 2

Construction of Incomplete Information Credit Models Based on Unpredictable Defaults

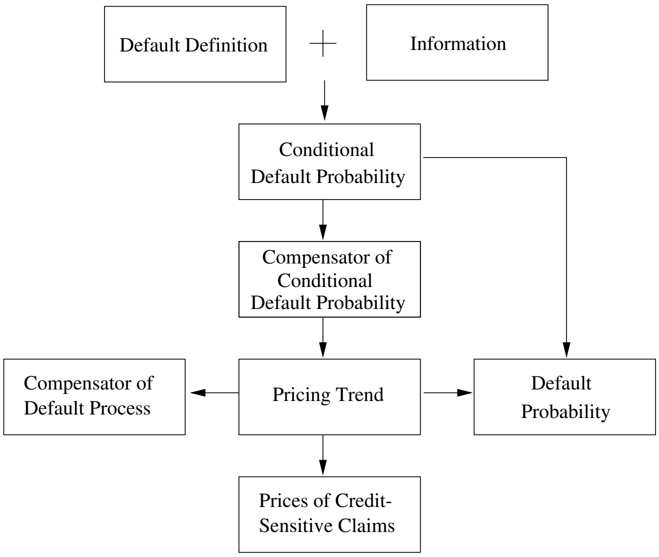

model. In both the Merton and the Black-Cox models, the observable information set includes the default event since the distance to default is observable. In I^2, the distance to default is not observable, so the model information set must be augmented in order to include default.

The compensator of the default process with respect to the augmented information set is equal to the pricing trend prior to default and is constant afterwards. An important mathematical result is that the default compensator has continuous paths if and only if default is unpredictable. This result underlies much of the analysis above.

III. I^2 MODEL SPECIFICATION AND CALIBRATION

We make three assumptions. The first two are:

1. Default is triggered when the value of the firm falls below a barrier.
2. Prior to default, the firm value process follows a geometric Brownian motion that can be observed.

The pre-default firm value process is specified by its drift, volatility, and initial value. The drift is the short-term risk-free rate, which is estimated from observed Treasury prices using generalized least squares. The estimates are updated daily.

To estimate the remaining parameters of the firm value process, we view firm equity as a European call option on pre-default firm value. The strike of the option is equal to the short-term debt (≤1 year) as reported in Compustat, and the maturity is one year. The Black-Scholes formula relates equity values $S(t)$ to firm values $V(t)$ and the volatility of firm values σ. Itô's formula implies that equity volatility $\sigma_S(t, \omega)$ at time t in state ω is related to σ through the equation:

$$S(t,\omega)\sigma_S(t,\omega) = \Delta(t,\omega)V(t,\omega)\sigma \tag{18}$$

where $\Delta(t, \omega)$ is the delta of the option at time t in state ω. Given market prices of equity and equity volatility forecast by Barra's U.S. equity trading model, we solve the two equations simultaneously for the corresponding firm values and firm volatility. In Giesecke and Goldberg [2004], we calibrate the model assuming that equity is a barrier option on firm value whose strike price may be stochastic.

A third assumption completes the model specification:

3. The default barrier is a random variable that is independent of firm value and cannot be observed by investors. It follows a scaled beta distribution.

We choose the beta distribution because it provides a flexible class of parametric distributions. A standard beta distribution is supported on the interval (0, 1) and is specified by its mean and variance. We scale the standard beta by a height parameter κ so that the distribution is supported on $(0, \kappa)$.

The height of the beta distribution is expressed in terms of firm leverage. The idea is that when the leverage ratio is at a recent high, short-term uncertainty is also at a recent high. The height κ is updated daily. We set it to the current debt divided by the maximum historical leverage ratio observed over the previous six months. Therefore, current firm value is an upper bound for the default barrier. When the leverage ratio is at a six-month high, κ is equal to current firm value.[6]

The mean m of the barrier distribution is set to the short-term debt as reported in Compustat. The debt is updated quarterly.

The variance of the barrier distribution can be treated as a free parameter that can be used to calibrate the degree of confidence the investor feels about the information that is publicly available. Our base case variance is set to $(\kappa - m)(m/10)$.

The Black-Cox model is the special case of the I^2 model where the variance is zero. The barrier is a con-

EXHIBIT 3
One-Year Default Probabilities for United Airlines 2001-2003

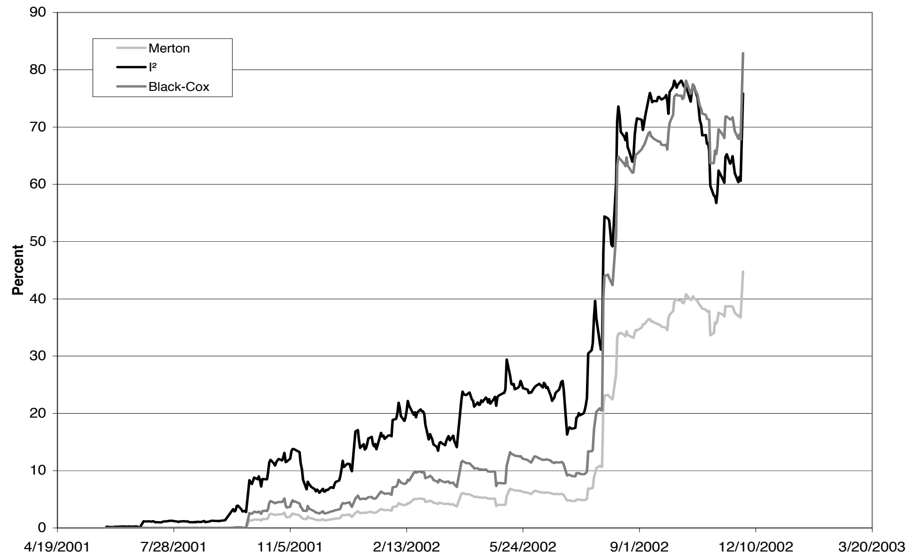

EXHIBIT 4
One-Year Default Probabilities for United Airlines 2001-2002

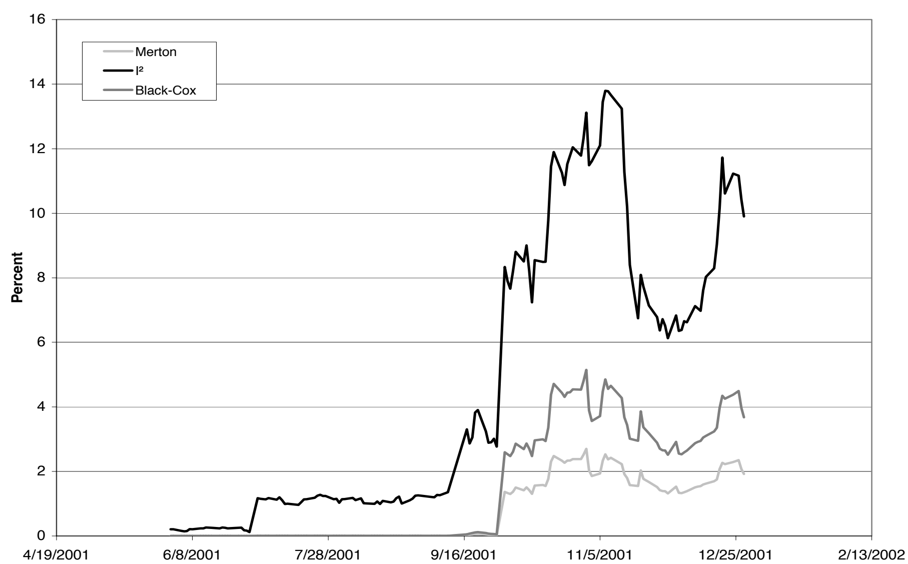

stant and equal to the short-term debt from Compustat. The Merton model is a degenerate Black-Cox model where default is triggered when firm value falls below the barrier at debt maturity.

IV. EMPIRICAL RESULTS

We use several case studies to compare the model forecasts of the calibrated Merton, Black-Cox, and I^2 models with respect to default probabilities and credit spreads.

Early Reaction

Exhibits 3 and 4 show time series of one-year United Airlines default probabilities generated by three different models. The lowest probabilities are generated by the Merton model. The other two curves are generated by first-passage time models, which imply higher default probabilities.

The silhouettes of the curves forecast by the first-passage models are quite similar. The differences are due to the different information assumptions underlying the

EXHIBIT 5
Time Series of Leverage Ratios for General Electric

two models. While the Black-Cox model posits that the default barrier is certain, the I^2 model assumes that the default barrier is random and unobserved.

Thanks to the higher uncertainty built into the I^2 model, it generates slightly higher probabilities than the Black-Cox model throughout most of the study period. This uncertainty allows for the possibility that the default barrier is closer than publicly advertised. The situation reverses in November 2002 when the leverage ratio gets very close to 1. Here, the default barrier is so close to the value of the firm that the uncertainty allows mainly for the possibility that the default barrier might not be as near as we think.[7]

Our next case study is General Electric, which is of higher credit quality than United. Exhibit 5 shows a time series of GE leverage ratios. The values show a slow and steady increase from 0.18 to almost 0.4 between April 2001 and the end of 2002.

In Exhibit 6 we see that the I^2 model shows a slow and steady increase in one-year default probabilities for GE, ranging from a few basis points to almost 100 basis points over the same time horizon. The Black-Cox model looks flat on this time scale. The Black-Cox and Merton models also react, as seen in Exhibit 7. Note the similarity in the shapes of the graphs in Exhibits 6 and 7.

These case studies illustrate the fact that I^2 reacts earlier than the Black-Cox and Merton models. This is because the I^2 model is sensitive to a history of leverage ratios rather than just the current value. This sensitivity comes from the specific form of the I^2 pricing trend in Equation (9), which depends on the minimum firm value and the default barrier distribution.

Positive Short-Term Credit Spreads

September 17, 2001, was the first trading day after the World Trade Center was attacked and destroyed. On this day, the value of United Airlines stock fell more than 40%, resulting in a large jump in the leverage ratio. Exhibit 8 shows term structures of credit spreads for United Airlines immediately before and after the crisis.[8] The new high in leverage ratio raised credit spreads across the board and created a positive short spread; see Exhibit 9. The Black-Cox and Merton models cannot forecast positive short spreads because of the assumption of predictability.

In mid-April 2002, renowned money manager Bill Gross of Pimco publicly took General Electric to task for some of its borrowing practices as well as its failure to communicate with its investors. In response, the value of GE stock fell almost 15% over a period of several days.

Exhibit 10 shows term structures of credit spreads for GE on dates before and after the fall. It is interesting not only that credit spreads rose during this five-day period, but also that the shape of the curve changed dramatically by steepening at the short end. Of course, the spread values are low, on the order of a few basis points. The likelihood of default is still very low, but between April 10 and 15, 2002, there was an increase in uncertainty about the true

EXHIBIT 6

One-Year Default Probabilities for General Electric

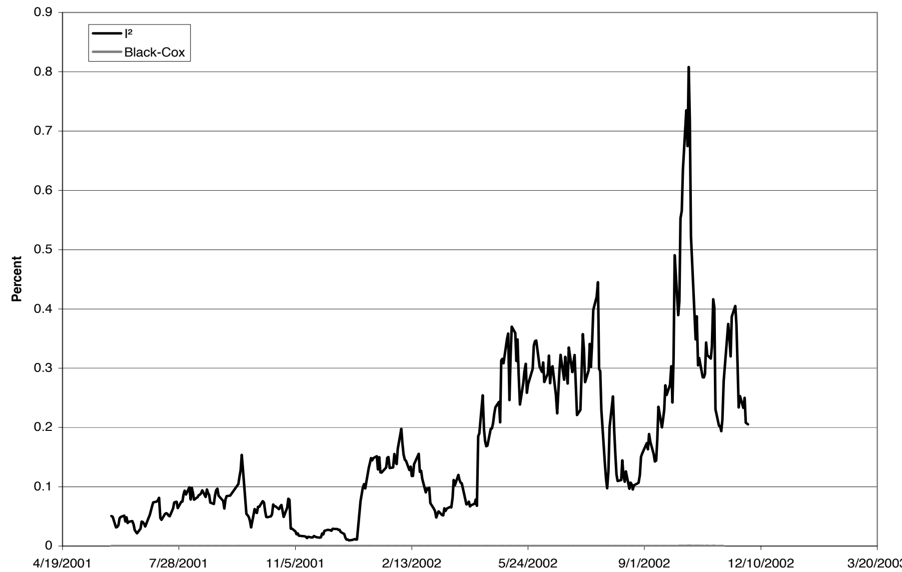

EXHIBIT 7

One-Year Default Probabilities for General Electric—Black-Cox and Merton Models

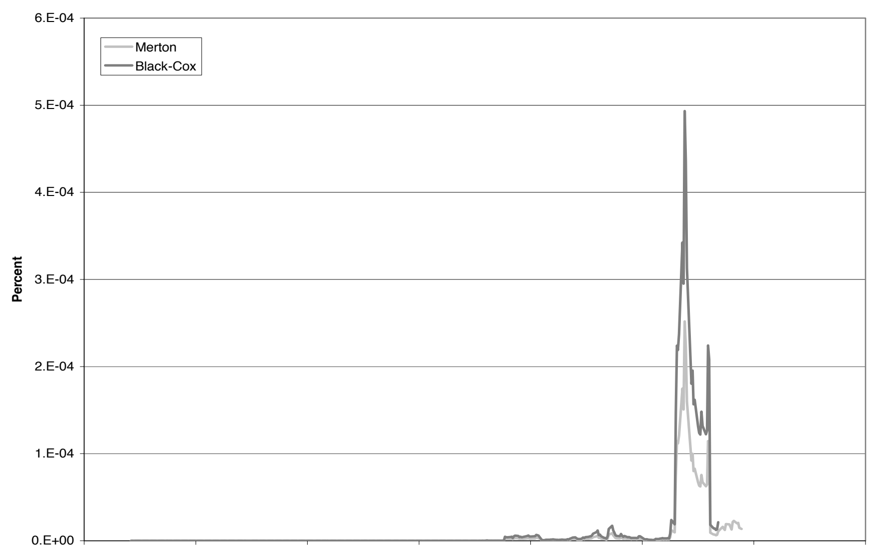

level of GE liabilities, and this was reflected immediately in the I^2 model.

Default Barrier Calibrated to Quality of Information

In Exhibit 11, we display one-year United Airlines default probabilities forecast by the Black-Cox model along with the base case I^2 model and several I^2 variants. The variants $I^2(k)$ are made by multiplying the base case variance of the default barrier distribution by a factor k. The exhibit indicates that if the leverage seems reasonable, greater uncertainty in the model translates into higher forecast default probabilities.

IBM represents the diametric opposite of United Airlines. Over the past several years, IBM's leverage ratio

Exhibit 8
Term Structures of I^2 Default Probabilities for United Airlines 9/10-9/17/01

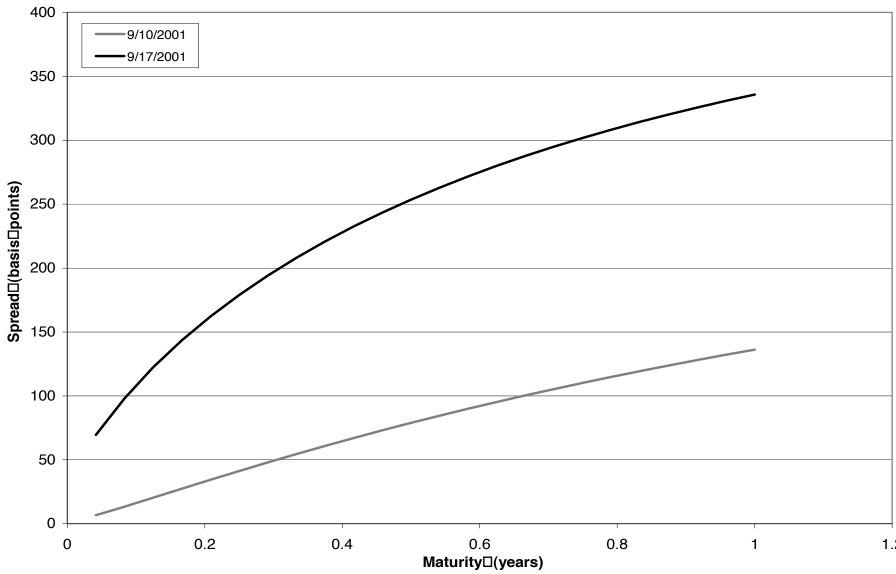

Exhibit 9
Time Series of United Airlines Short Spreads Generated by I^2 Model

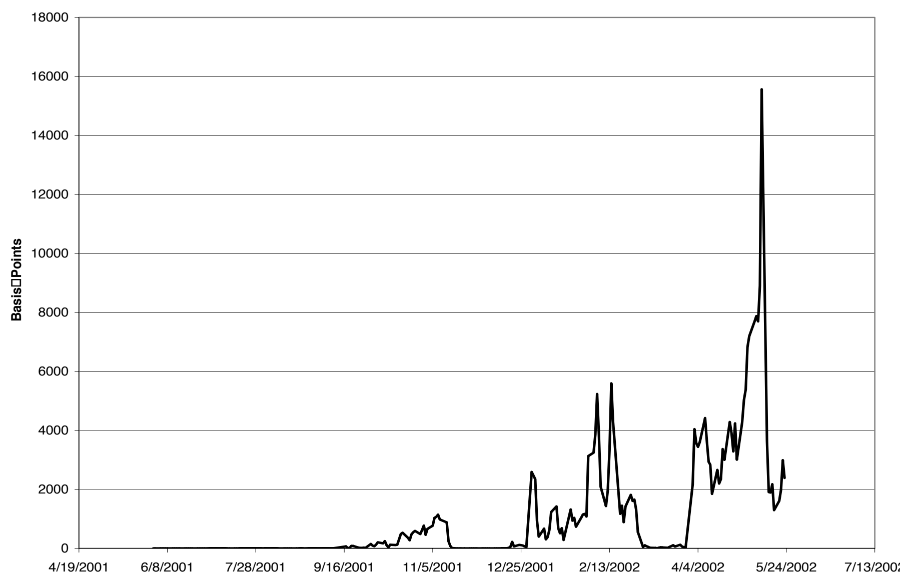

has been under 0.1, roughly 5% to 10% the leverage ratio at United. Given these numbers, any reasonable structural model will forecast that there is virtually no chance IBM will imminently default on its debt. The I^2 model is no exception. An investor who believes the leverage ratio is higher than indicated by the published numbers, though, could increase the variance of the default barrier distribution.

Exhibit 12 shows one-year IBM default probability forecasts by the base case I^2 model and several variants. In the period examined, the highest one-year default probability forecast by the base case I^2 model is quite low, less than 30 basis points. As one's confidence in published numbers diminishes, however, we can increase the variance around the default barrier location and see a corresponding increase in forecast default probability.

EXHIBIT 10
Term Structures of Credit Spreads for General Electric Generated by I^2 Model

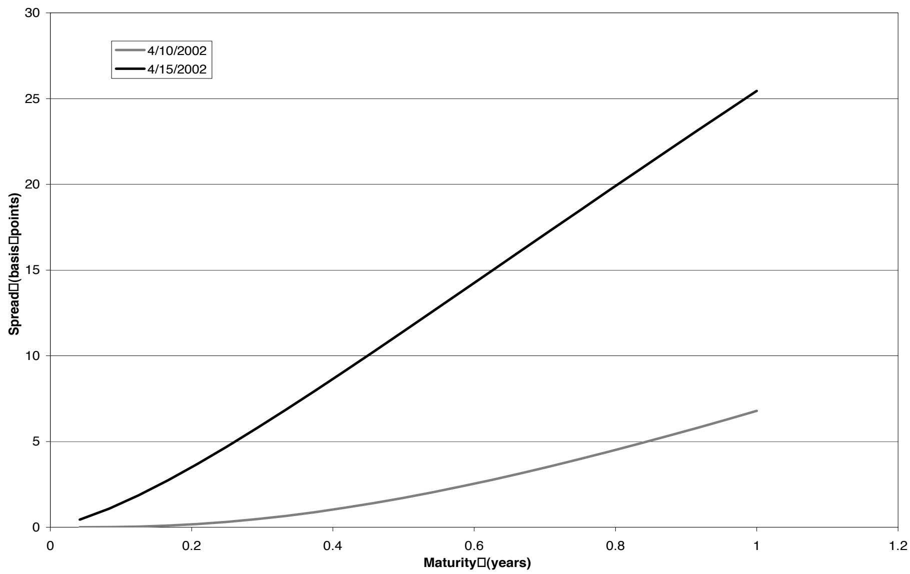

EXHIBIT 11
One-Year Default Probabilities for United Airlines Generated by Various Models

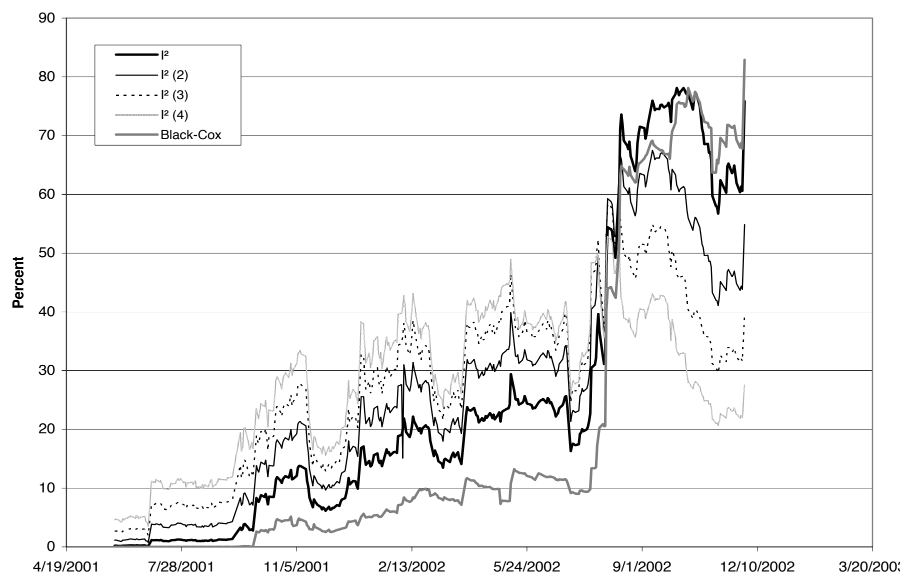

Exhibit 13 shows the base case I^2 and Black-Cox models together. While both models forecast low probabilities of default over a one-year horizon, the Black-Cox numbers are much lower—but they are not zero, as shown in Exhibit 14.

Together, Exhibits 13 and 14 show a similar silhouette in the time series of default probabilities forecast by the two models. Both models indicate that IBM debt is of the highest quality. Only the I^2 model, however, can take account of the degree of uncertainty an investor feels about the location of the default barrier.

V. CONCLUSION

Incomplete information is the basis of a powerful approach to analyzing credit risk. By providing a common

Exhibit 12

One-Year Default Probabilities for IBM Generated by Various Models

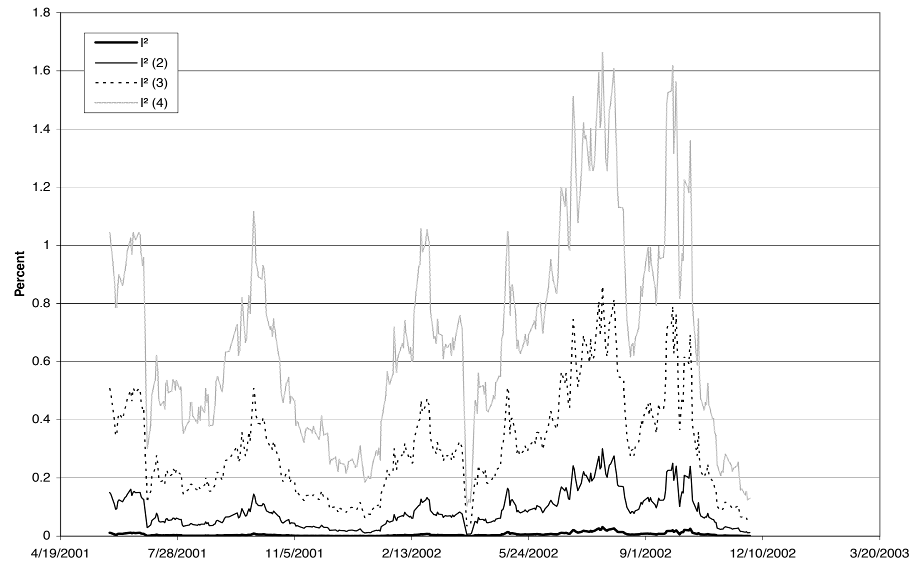

Exhibit 13

One-Year Default Probabilities for IBM Generated by Base Case I² Model and Black-Cox Model

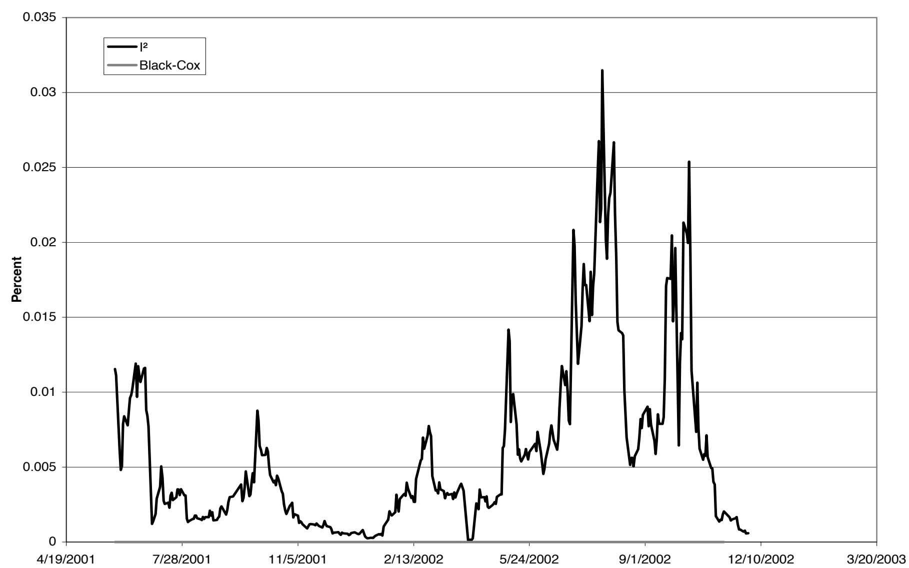

perspective on structural and reduced-form models, the incomplete information framework facilitates the inclusion of short-term uncertainty into a cause-and-effect model. It illuminates structural and reduced-form hybrid models that incorporate the best features of both approaches while avoiding many of their shortcomings.

As in the Black and Cox [1976] and Nielsen, Saá-Requejo, and Santa-Clara [1993] models, the I^2 credit model is based on a first-passage definition of default. By assuming incomplete information about the firm value that triggers default, I^2 acquires reduced-form pricing formulas that can be used to calibrate to market data.

We compare the Merton, Black-Cox, and I^2 default probability and credit spread forecasts in several case studies. Unlike the other two models, I^2 forecasts positive short spreads consistent with empirical observation. It supports a flexible family of shapes for the short end of the term structure of credit spreads. Using a history of firm values, rather than just the current value, enables I^2 to react

EXHIBIT 14
One-Year Default Probabilities for IBM Generated by Black-Cox Model

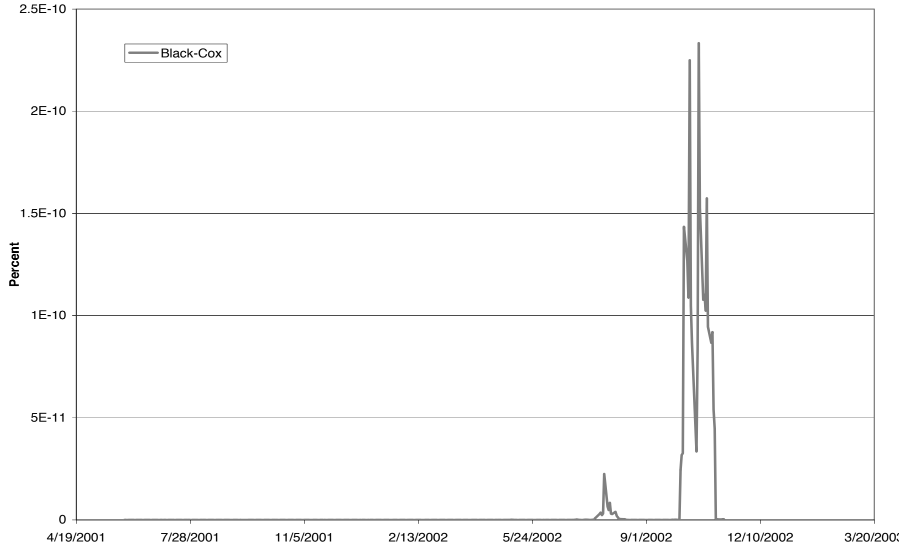

quickly to information as it becomes available.

In work elsewhere, we extend I^2 to price securities by analyzing the credit risk premium and implied recovery. We describe an I^2 calibration based on a pool of equity, bond, and credit default swap data. We also demonstrate that contagion can be modeled when there is incomplete information in multiple-firm markets. We have only begun to explore the possibilities.

ENDNOTES

The authors thank Greg Anderson, Tim Backshall, Roveen Bhansali, Ursula Gritsch, Guy Miller, and Vijay Poduri for their contributions to this article. The data for the empirical studies was generously supplied by MSCI Barra, Inc. The data for the empirical studies were generously supplied by MSCI Barra, Inc. Kay Giesecke acknowledges financial support by Deutsche Forschungsgemeinschaft.

[1]We use the term *predictable* because it is a mathematical standard. Unfortunately, it often leads to confusion. In our context, the term does not imply that the model is without uncertainty. To the contrary: Both the Merton and the Black-Cox models assume that future firm value is uncertain.

[2]Mathematical precision requires a fixed probability space $(\Omega, \mathcal{F}, P)$. Here, Ω is a set representing the possible states of the world, $\mathcal{F}$ is a sigma-algebra that determines the resolution to which states can be distinguished. The symbol P denotes a probability measure that can be thought of as either a physical measure or a pricing measure.

[3]*Integrated hazard rate* or hazard function is a standard synonym for *pricing trend*. We use the latter term for brevity and also for its relationship with our application. Also, the standard term might be misleading; sometimes there is no hazard rate at all.

[4]The last condition means there is always some chance that a firm will survive beyond any fixed time. The Merton and Black-Cox models of Examples I.1 and I.2 do not satisfy this condition.

[5]This is a special case of the Doob-Meyer Theorem, which asserts that this is always possible for any uniformly integrable submartingale. See Dellacherie and Meyer [1982] for details.

[6]Our empirical studies show that six months is an optimal horizon. A longer horizon is likely to include change in regime. A shorter horizon results in forecasts that are too volatile.

[7]In practice, we set an upper bound on the forecast default probability on the grounds that it is impossible to distinguish very risky firms using a model.

[8]The credit spreads are based on zero-recovery zero-coupon bonds with prices given by equation (11).

REFERENCES

Artzner, Philippe, and Freddy Delbaen. "Default Risk Insurance and Incomplete Markets." *Mathematical Finance,* 5 (1995) pp. 187-195.

Black, Fischer, and John C. Cox. "Valuing Corporate Securities: Some Effects of Bond Indenture Provisions." *Journal of Finance,* 31 (1976), pp. 351-367.

Black, Fischer, and Myron Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy,* 81 (1973), pp. 81-98.

"Creditgrades." RiskMetrics Group, Technical Document, 2002.

Dellacherie, Claude, and Paul-Andre Meyer. *Probabilities and Potential*. Amsterdam: North-Holland, 1982.

Duffie, Darrell, and David Lando. "Term Structures of Credit Spreads with Incomplete Accounting Information." *Econometrica,* 69(3) (2001), pp. 633-664.

Duffie, Darrell, and Kenneth J. Singleton. "Modeling Term Structures of Defaultable Bonds." *Review of Financial Studies,* 12 (1999), pp. 687-720.

Giesecke, Kay. "Default and Information." Working paper, Cornell University, 2001.

Giesecke, Kay, and Lisa Goldberg. "Sequential Defaults and Incomplete Information." Forthcoming, *Journal of Risk*, 2004.

——. "Calibrating Credit with Incomplete Information." Working paper, Cornell University, 2004.

——. "The Market Price of Credit Risk." Working paper, Cornell University, 2003.

Jarrow, Robert A., and Stuart M. Turnbull. "Pricing Derivatives on Financial Securities Subject to Credit Risk." *Journal of Finance,* 50(1) (1995), pp. 53-86.

Longstaff, Francis A., and Eduardo S. Schwartz. "A Simple Approach to Valuing Risky Fixed and Floating Rate Debt." *Journal of Finance,* 50(3) (1995), pp. 789-819.

Merton, Robert C. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance,* 29 (1974), pp. 449-470.

Meyer, Paul-Andre. *Probability and Potentials*. London: Blaisdell, 1966.

Nielsen, Lars, Jesus Saá-Requejo, and Pedro Santa-Clara. "Default Risk and Interest Rate Risk: The Term Structure of Default Spreads." Working paper, INSEAD, 1993.

To order reprints of this article, please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.