Default and Recovery Implicit in the Term Structure of Sovereign *CDS* Spreads

Jun Pan and Kenneth J. Singleton [1]

This draft: October 30, 2005

[1]Pan is with the MIT Sloan School of Management and NBER, junpan@mit.edu. Singleton is with the Graduate School of Business, Stanford University and NBER, ken@future.stanford.edu.
We benefited from discussions with Darrell Duffie, Francis Longstaff, Jun Liu, Roberto Rigobon, and seminar participants at UCLA, University of Michigan, and 2005 NBER IASE workshop. Scott Joslin provided excellent research assistance. We are grateful for financial support from the Gifford Fong Associates Fund, at the Graduate School of Business, Stanford University and financial support from the MIT Laboratory for Financial Engineering.

Abstract

This paper explores in depth the nature of the risk-neutral credit-event intensities ($\lambda^{\mathbb{Q}}$) that best describe the *term structures* of sovereign *CDS* spreads. We examine three distinct families of stochastic processes: the square-root, lognormal, and three-halves processes. These models employ different specifications of mean reversions and time-varying volatilities to fit both the distributions of spreads, and the variation over time in the shapes of the term structures of spreads. We find that the models imply highly persistent $\lambda^{\mathbb{Q}}$ that are strongly correlated with measures of global credit event risks and the *VIX* index of option-implied volatilities. Moreover, the correlations across countries of the model-implied credit-event intensities are large, and change with credit-market conditions. There are substantial model-implied risk premiums associated with unpredictable future variation in $\lambda^{\mathbb{Q}}$. We show that the term structure of *CDS* spreads allows us to separately identify both the loss rate in the event of default, $L^{\mathbb{Q}}$, and the parameters of the process $\lambda^{\mathbb{Q}}$. Unconstrained estimates of $L^{\mathbb{Q}}$ are much lower than the values typically assumed in the financial industry. Finally, to shed light on the economic consequences of differing levels of $L^{\mathbb{Q}}$ or persistence in $\lambda^{\mathbb{Q}}$, we explore the sensitivity of the prices of options on *CDS* contracts to alternative settings of the parameters governing the default process.

1 Introduction

The burgeoning market for sovereign credit default swaps (*CDS*) contracts offers a nearly unique window for viewing investors' risk-neutral probabilities of major credit events impinging on sovereign issuers, and their risk-neutral losses of principal in the event of a restructuring or repudiation of external debts. In contrast to many "emerging market" sovereign bonds, sovereign *CDS* contracts are designed without complex guarantees or embedded options. Trading activity in the *CDS* contracts of several sovereign issuers has developed to the point that they are more liquid than many of the underlying bonds. Moreover, in contrast to the corporate *CDS* market, where trading has been concentrated largely in the five-year maturity contract, there are actively traded *CDS* contracts at several maturity points between one and ten years. As such, a full *term structure* of *CDS* spreads is available for inferring default and recovery information from market data.

This paper explores in depth the nature of the risk-neutral credit-event intensities ($\lambda^{\mathbb{Q}}$) that best describe the *term structures* of sovereign *CDS* spreads. Since little is known about the nature of the market-implied $\lambda^{\mathbb{Q}}$ processes that underlie the term structures of survival probabilities for sovereign issuers, we examine three distinct families of stochastic processes: the square-root (Cox, Ingersoll, and Ross [1985]), lognormal, and three-halves (Ahn and Gao [1999]). While all three models allow for mean reversion in $\lambda^{\mathbb{Q}}$, they differ in the degree of nonlinearity in their drifts. They also differ in their assumptions about the dependence of the instantaneous volatilities of $\lambda^{\mathbb{Q}}$ on itself: they take the form $(\lambda^{\mathbb{Q}})^{\gamma}$, with γ equal to .5 (square-root), 1 (lognormal), or 1.5 (three-halves).

Equally central to modeling the credit risk of sovereign issuers is the recovery in the event of default. Standard practice in modeling corporate *CDS* spreads is to assume a fixed risk-neutral loss rate $L^{\mathbb{Q}}$, largely because the focus has been on the liquid five-year *CDS* contract.[1] We depart from this literature and exploit the term structure of *CDS* spreads to separately identify both $L^{\mathbb{Q}}$ and the parameters of the process $\lambda^{\mathbb{Q}}$. That we even attempt to separately identify these parameters of the default process may seem surprising in the light of the apparent demonstrations in Duffie and Singleton [1997], Houweling and Vorst [2005], and elsewhere of the infeasibility of achieving this objective. We show that in fact, in market environments where recovery is a fraction of face value, as is the case with *CDS* markets, these parameters can be separately identified through the information contained in the term structure of *CDS* spreads.

We initially follow market practice and fix the risk-neutral loss rate $L^{\mathbb{Q}}$ at 0.75. At $L^{\mathbb{Q}} = 0.75$, the likelihood functions of the data associated with the three models for $\lambda^{\mathbb{Q}}$ imply different trade-offs between mean reversions and time-varying volatilities when fitting both the shapes of the distributions of spreads at a given maturity, and the variation over time in the term structures of spreads. Comparing the standardized pricing errors across models, we find that all three models fit roughly equally well during the relatively calm period from the middle of 2003 until the end of our sample (August, 2004). During the more

[1] See, for example, Berndt, Douglas, Duffie, Ferguson, and Schranzk [2004], Hull and White [2004], and Houweling and Vorst [2005].

turbulent periods of the fall of 2001 and the second half of 2002, the trade-off implicit in the lognormal model appears to represent a good compromise for modeling spreads for most of the maturities examined.

The maximum likelihood (*ML*) estimates of the parameters imply that the risk-neutral– the $\mathbb{Q}$-distribution– of $\lambda^{\mathbb{Q}}$ shows very little mean reversion. In fact, $\lambda^{\mathbb{Q}}$ is explosive under the $\mathbb{Q}$-distribution for several countries and models. In contrast, under the measure associated with the historical data-generating process– the $\mathbb{P}$-distribution– $\lambda^{\mathbb{Q}}$ shows substantial mean reversion, consistent with the stationary property of *CDS* spreads. This large difference between the properties of $\lambda^{\mathbb{Q}}$ under the $\mathbb{Q}$ and $\mathbb{P}$ measures implies, within the context of our models, that an economically important systematic risk is being priced in the *CDS* market.

Further insight into the nature of this risk comes from examination of the correlations among the $\lambda^{\mathbb{Q}}$ across countries, and among the $\lambda^{\mathbb{Q}}$ and observable market risk-related variables. There is a high degree of correlation between the fitted $\lambda^{\mathbb{Q}}$ from our models and the time series of global speculative grade default rates computed by Moody's. Additionally, there is a strikingly large correlation between the *CDS* spreads and the U.S. *VIX* index of equity option-implied volatilities. Both of these findings suggest that $\lambda^{\mathbb{Q}}$ is capturing global credit event risks. Moreover, a substantial portion of the co-movement among the term structures of sovereign spreads across countries appears to be due to the market's appetite for credit exposure at a global level, rather than to any sovereign events *per se*. Consistent with these observations, the degree to which the model-implied $\lambda^{\mathbb{Q}}$ are correlated across countries varies over time and with sovereign credit market conditions.

A key attraction of studying the term structures of *CDS* spreads is that the parameters governing $\lambda^{\mathbb{Q}}$ and $L^{\mathbb{Q}}$ are, in principle, separately identified. Interestingly, when we search over $L^{\mathbb{Q}}$, instead of imposing the market convention of $L^{\mathbb{Q}} = 0.75$, the likelihood function calls for a much smaller value, more in the region of 0.25. Monte Carlo evidence based on our models evaluated at their maximum likelihood estimates suggests that the difference between our unconstrained estimates and those typically used by market makers cannot be attributed to small sample biases in our estimates of $L^{\mathbb{Q}}$.

At the same time, the pricing errors implied by the models evaluated at the maximum likelihood estimates with $L^{\mathbb{Q}}$ set at 0.75 or at 0.25 are very similar. Therefore, to shed light on the practical economic differences between pricing models with high or low loss rates, we examine the model-implied prices of European options on sovereign *CDS* contracts within the lognormal model. Particular attention is given to the sensitivities of option prices to the degree of persistence in $\lambda^{\mathbb{Q}}$ and the level of $L^{\mathbb{Q}}$.

Throughout this analysis we maintain the assumption that a single risk factor underlies the temporal variation in $\lambda^{\mathbb{Q}}$, consistent with most previous studies of *CDS* spreads that have allowed for a stochastic arrival rate of credit events. In the case of our sovereign data, this initial focus is motivated by the high degree of comovement among spreads across the maturity spectrum within each country. For our sample period, this comovement is even greater than the high degree of correlations between yields in highly liquid treasury markets documented, for example, in Litterman and Scheinkman [1991]. In our concluding remarks we assess the potential role for a second risk factor, particularly for understanding *CDS* spreads at the short end of the maturity spectrum.

To our knowledge, the closest precursor to our analysis is the complementary study by Zhang [2003] of *CDS* spreads for Argentina leading up to the default in late 2001. We focus on a different set of countries and broader set of models for the underlying intensity processes (he assumed a completely affine, square-root process), and examine in more depth the implications of model specification for the properties of the default process. Additionally, our sample period begins towards the end of his, is longer in length, and spans a period during which the sovereign *CDS* markets were more developed in breadth and liquidity.

2 The Structure of the Sovereign *CDS* Market

The structure of the standard *CDS* contract for a sovereign issuer shares many of its features with the corporate counterpart. The default protection buyer pays a semi-annual premium, expressed in basis points per notional amount of the contract, in exchange for a contingent payment in the event one of a pre-specified credit events occurs. Settlement of a *CDS* contract is typically by physical delivery of an admissible bond in return for receipt of the original face value of the bonds,[2] with admissibility determined by the characteristics of the reference obligation in the contract.

Typically, only bonds issued in *external* markets and denominated in one of the "standard specified currencies" are deliverable.[3] In particular, bonds issued in domestic currency, issued domestically, or governed by domestic laws are not deliverable. For some sovereign issuers without extensive issuance of hard-currency denominated Eurobonds, loans may be included in the set of deliverable assets. Among the countries included in our analysis, Turkey and Mexico have sizeable amounts of outstanding loans, and their *CDS* contracts occasionally trade with "Bond or Loan" terms. The contracts we focus on are "Bond only."

The key definition included in the term sheet of a sovereign *CDS* contract is the credit event. Typically, a sovereign *CDS* contract lists as events any of the following that affect the reference obligation: (i) obligation acceleration, (ii) failure to pay, (iii) restructuring; or (iv) repudiation/moratorium. Note that "default" is not included in this list, because there is no operable international bankruptcy court that applies to sovereign issuers.

Central to our analysis of the term structure of sovereign *CDS* spreads is the active trading of contracts across a wide range of maturities. In contrast to the U.S. corporate and bank *CDS* markets, where a large majority of the trading volume is concentrated in five-year

[2]Physical delivery is the predominant form of settlement in the sovereign *CDS* market, because both the buyers and sellers of of protection typically want to avoid the dealer polling process involved in determining the value of the reference bond in what is often a very illiquid post-credit-event market place.

[3]The standard specified currencies are the Euro, U.S. dollar, Japanese yen, Canadian dollar, Swiss franc, and the British pound. The option to deliver bonds denominated in these currencies, and of various maturities, into a *CDS* contract introduces a cheapest-to-deliver option for the protection buyer. Our impression, from conversations with traders, is that usually there is a single bond (or small set of bonds) that are cheapest to deliver. So the price of the *CDS* contract tracks this cheapest to deliver bond and the option to deliver other bonds is not especially valuable. In any event, for the purpose of our subsequent analysis, we will ignore this complication in the market.

contracts, the three- and ten-year contracts have each accounted for roughly 20% of the volumes in sovereign markets, and the one-year contract has accounted for an additional 10% of the trading (see Figure 1).[4] While the total volume of new contracts has been much larger in the corporate than in the sovereign market, the volumes for the most actively traded sovereign credits are large and growing. We focus our analysis on Mexico, Russia, and Turkey, three of the more actively traded names.[5]

Figure 1: *CDS* volumes by maturity, as a percentage of total volume, based on *BIS* calculations from CreditTrade data. Source: *BIS* Quarterly Review [2003].

Our sample consists of daily trader quotes of bid and ask spreads for *CDS* contracts with maturities of one, three, five, and ten years. The sample covers the period March 19, 2001 through August 12, 2004. We focus on the data for three countries Mexico, Russia, and Turkey, displayed in Figure 2. At the beginning of our sample period (March, 2001), Mexico had achieved the investment grade rating of Baa3. In February, 2002, Mexico was upgraded one notch to Baa2.[6] Russia was rated B2 in the spring of 2001, was upgraded to Ba3 at the end of November, 2001, and Ba2 in December, 2002. Russia achieved the investment grade rating of Baa3 at the beginning of October, 2003. Turkey maintained the same speculative grade rating, B1, throughout our sample period. However, both in April, 2001 and July, 2002 it was put in the "negative outlook" category. Following the most recent negative outlook, Turkey returned to "stable outlook" in October, 2003. Consistent with the relative credit

[4]Figure 1 is a corrected version of the original appearing in Packer and Suthiphongchai [2003].

[5]Several South American credits– Brazil, Columbia, and Venezuela– are also among the more traded sovereign credits. The behavior of their *CDS* spreads was largely dominated by the political turmoil in Brazil during the summer/fall of 2002. The co-movements among the *CDS* spreads of these countries is an interesting question for future research.

[6]Mexico was subsequently upgraded again one notch to Baa1 in January, 2005, but this is beyond our sample period.

qualities of these countries, the average five-year *CDS* spreads over our sample period are 216, 435, 752 basis points, respectively, for Mexico, Russia, and Turkey.

In addition to the fact that they cover a broad range of credit quality, two important considerations factor into our choice of these three countries: their regional representativeness in the emerging markets and the relative liquidity and thus better data quality of their *CDS* markets compared to those of other countries in the same region. While the first consideration is important for the economic interpretation of our results, the latter consideration plays a more crucial role given our focus on the empirical estimation and evaluation of various default models.

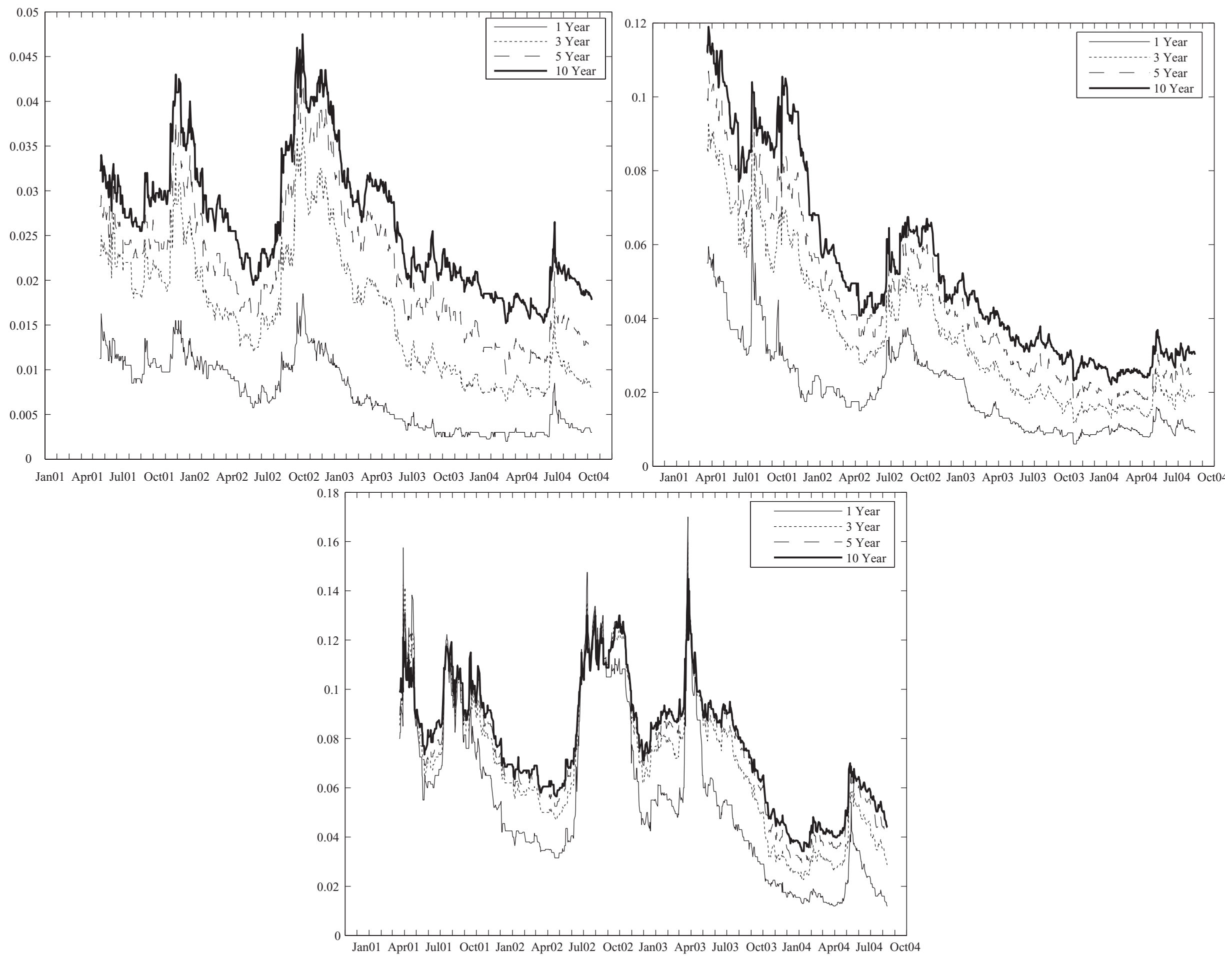

Figure 2: CDS Spreads: Mexico (upper left), Russia (upper right), and Turkey (lower), mid-market quotes.

As shown in Figure 2, the term-structures of *CDS* spreads exhibit interesting dynamics. One immediately noticeable feature present in all three countries is the high level of co-movement among the 1y, 3y, 5y, and 10y *CDS* spreads. Indeed, a principal component (*PC*) analysis of the spreads in each country (see Section 10) shows that the first *PC* explains (by the usual measure in *PC* analyses) 97.8, 98.5, and 98.1 percent of the variation in spreads in Mexico, Russian, and Turkey, respectively. It is these high levels of explained levels of variation that motivate our focus on one-factor models.

Another prominent feature of the *CDS* data is the persistence of upward sloping term structures. This is especially true for the term structures of Mexican and Russian *CDS* spreads: throughout our sample period, the one-year *CDS* spreads were always lower than the respective longer maturity *CDS* spreads and, hence, the term structure was never inverted. For example, the difference between the five-year and one-year Mexican *CDS* spreads was 144 basis points on average, 73 basis points at minimum, and 275 basis points at maximum. Without resorting to institutional features that might separate the one-year from the longer maturity *CDS* contracts, this feature of *CDS* spreads implies an increasing term structure of risk-neutral one-year forward default probabilities.

The slope of the term structure of *CDS* spreads for Turkey was mostly positive. For example, the difference between the five- and one-year *CDS* spreads was on average 195 basis points with a standard deviation of 112 basis points. However, in contrast to the robust pattern of upward sloping spread curves in Mexico and Russia, the term structure of Turkish *CDS* spreads did occasionally invert, especially when credit spreads exploded to high levels due to financial or political crises that were (largely) specific to Turkey. For example, the differences between the five- and one-year *CDS* spreads were −250 basis points on March 29, 2001, −150 basis points on July 10, 2002, and −200 basis points on March 24, 2003. The related events were the devaluation of the Italian lira, political elections in Turkey, and the collapse of talks between Turkey and Cyprus (which had implications for Turkey's bid to join the EU).

Sovereign credit default swaps trade, on average, in larger sizes than in the underlying cash markets: U.S. $5 million, and occasionally much larger, against U.S. $1 - 2 million. The liquidity of the underlying bond market is relevant, because traders hedge their *CDS* positions with cash market instruments and the less liquid is the cash market, the larger the bid/ask spread must be in the *CDS* market to cover the higher hedging costs. Comparing across sovereign *CDS* markets, a given bid/ask spread will sustain a larger trade in the market for Mexico (up to about $40 million) relative to say Turkey (up to about $30 million) (Xu and Wilder [2003]).

For our sample of countries, the bid/ask spreads (in basis points for the five-year contract) ranged between 5 - 30 for Mexico, 10 to 50 for Russia, and 25 - 300 for Turkey (see Figure 3). For high-grade countries with large quantities of bonds outstanding like Mexico, the magnitudes of the bid/ask spreads in the *CDS* markets are comparable to those for their bonds. The bid/ask spreads for the Russian *CDS* contracts were somewhat larger, particularly during the early part of our sample. As Russia's credit quality improved and spreads on Russian *CDS* contracts narrowed, so did the bid/ask spreads on its *CDS* contracts. By the end of our sample period, Mexico and Russia had comparable spreads. Turkey had the largest bid/ask spreads, and its spreads show the most variation. Notably, when Turkey's spreads widened out due to the "local" events chronicled above, so did the bid/ask spreads.

Particularly at the short end of the maturity spectrum, the *CDS* market is often more

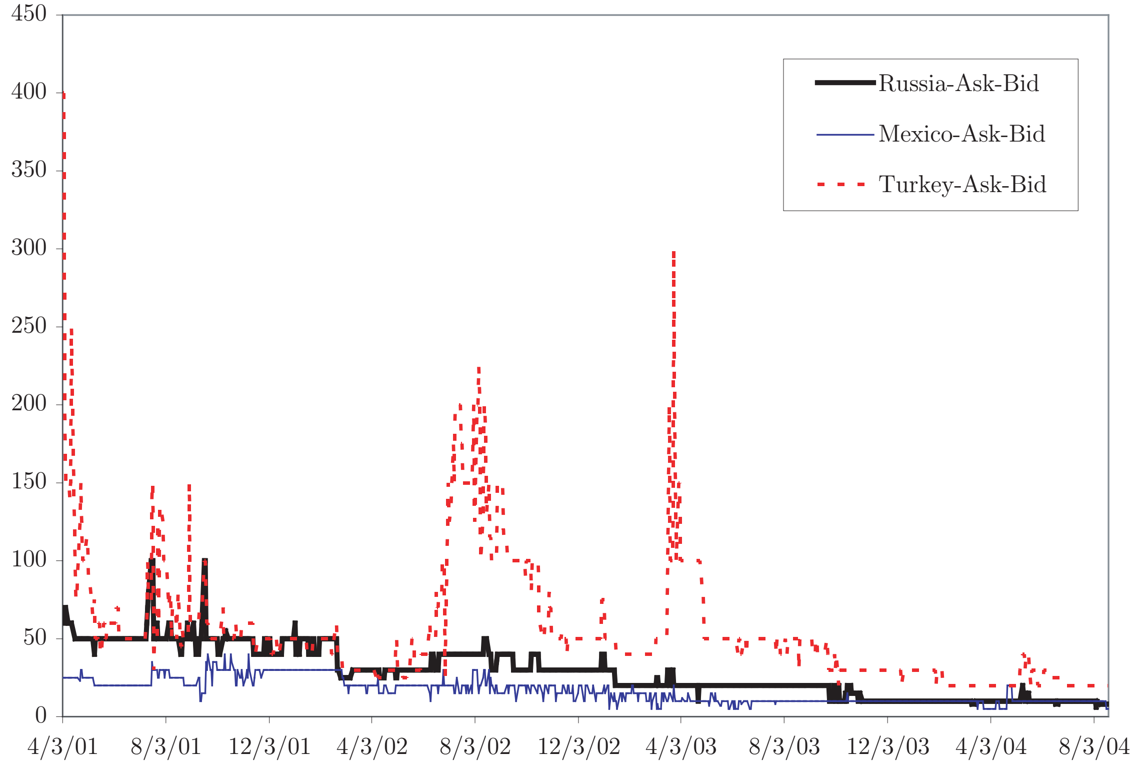

Figure 3: Ask-Bid Spreads (basis points) for five-year *CDS* contracts

liquid than the underlying cash markets and, as such, market participants experience smaller bid/ask spreads. For our sample period, the bid/ask spreads for the one-year contracts (not shown) are comparable in magnitude to those of the five-year contracts. The primary exceptions are during turbulent periods, especially in Turkey, when the levels of *CDS* spreads are large. In such cases, the bid/aks spreads on the one-year contracts are larger then those on the five-year contracts.

3 Models for Pricing *CDS* Contracts

The basic pricing relation for sovereign *CDS* contracts is identical to that for corporate *CDS* contracts. Let M denote the maturity (in years) of the contract, $CDS_t(M)$ denote the (annualized) spread at issue, $R^{\mathbb{Q}}$ denote the (constant) risk-neutral fractional recovery of face value on the underlying (cheapest-to-deliver) bond in the event of a credit event, and $\lambda^{\mathbb{Q}}$ denote the risk neutral arrival rate of a credit event. Then, at issue, a *CDS* contract with semi-annual premium payments is priced as (see, e.g., Duffie and Singleton [2003]):

$$\frac{1}{2}\, CDS_t(M) \sum_{j=1}^{2M} E_t^{\mathbb{Q}}\left[e^{-\int_t^{t+.5j}(r_s+\lambda_s^{\mathbb{Q}})ds}\right] = (1-R^{\mathbb{Q}}) \int_t^{t+M} E_t^{\mathbb{Q}}\left[\lambda_u^{\mathbb{Q}} e^{-\int_t^{u}(r_s+\lambda_s^{\mathbb{Q}})ds}\right] du, \qquad (1)$$

where r_t is the riskless rate relevant for pricing *CDS* contracts. The left-hand-side of (1) is the present value of the buyer's premiums, payable contingent upon a credit event not having occurred. Discounting by $r_t + \lambda_t^{\mathbb{Q}}$ captures the survival-dependent nature of these payments (Lando [1998]). The right-hand-side of this pricing relation is the present value of the contingent payment by the protection seller upon a credit event. We have normalized the face value of the underlying bond to \$1 and assumed a constant expected contingent payment (loss relative to face value) of $L^{\mathbb{Q}} = (1 - R^{\mathbb{Q}})$. In implementing (1), we use a slightly modified version that accounts for the buyer's obligation to pay an accrued premium if a credit event occurs between the premium payment dates.

How should $\lambda^{\mathbb{Q}}$ and $L^{\mathbb{Q}}$ be interpreted, given that default is not a relevant credit event, and *ISDA* terms sheets for plain vanilla sovereign *CDS* contracts reference four types of credit events? To accommodate this richness of the credit process for sovereign issuers, let each of the four relevant credit events have their own associated arrival intensities $\lambda_i^{\mathbb{Q}}$ and loss rates $L_i^{\mathbb{Q}}$. Then, following Duffie, Pedersen, and Singleton [2003] and adopting the usual "doubly stochastic" formulation of arrival of credit events (see, e.g., Lando [1998]), we can interpret the $\lambda_t^{\mathbb{Q}}$ and $L_t^{\mathbb{Q}}$ for pricing sovereign *CDS* contracts as:

$$\lambda_t^{\mathbb{Q}} = \lambda_{acc,t}^{\mathbb{Q}} + \lambda_{fail,t}^{\mathbb{Q}} + \lambda_{rest,t}^{\mathbb{Q}} + \lambda_{repud,t}^{\mathbb{Q}}, \tag{2}$$

$$L_t^{\mathbb{Q}} = \frac{\lambda_{acc,t}^{\mathbb{Q}}}{\lambda_t^{\mathbb{Q}}} L_{acc,t}^{\mathbb{Q}} + \frac{\lambda_{fail,t}^{\mathbb{Q}}}{\lambda_t^{\mathbb{Q}}} L_{fail,t}^{\mathbb{Q}} + \frac{\lambda_{rest,t}^{\mathbb{Q}}}{\lambda_t^{\mathbb{Q}}} L_{rest,t}^{\mathbb{Q}} + \frac{\lambda_{repud,t}^{\mathbb{Q}}}{\lambda_t^{\mathbb{Q}}} L_{repud,t}^{\mathbb{Q}}, \tag{3}$$

where the subscripts represent acceleration, failure to pay, restructuring, and repudiation. In a doubly stochastic setting, conditional on the pathes of the intensities, the probability that any two of the credit events happen at the same time is zero. Thus, $\lambda^{\mathbb{Q}}$ is naturally interpreted as the arrival rate of the first credit event of any type. Upon the occurrence of a credit event of type i, the relevant loss rate is L_i^{Q} and, given that a credit event has occurred, this loss rate is experienced with probability $\lambda_{it}^{\mathbb{Q}}/\lambda_t^{\mathbb{Q}}$. The corresponding $\lambda_i^{\mathbb{Q}}$ and $L_i^{\mathbb{Q}}$ may, of course, differ across countries.

We examine three specifications for the default arrival process: $\lambda^{\mathbb{Q}}$ follows a square-root diffusion, $\ln(\lambda^{\mathbb{Q}})$ follows an Ornstein-Uhlenbeck (Gaussian) process, and $\lambda^{\mathbb{Q}}$ follows a "three-halves" diffusion. In the literature on corporate *CDS* spreads the square-root model was studied by Longstaff, Mithal, and Neis [2005], while Berndt, Douglas, Duffie, Ferguson, and Schranzk [2004] adopted the lognormal model. Zhang [2003] had $\lambda^{\mathbb{Q}}$ following a square-root process in his analysis of Argentinean *CDS* contracts. One of our primary objectives it to provide a systematic comparison of how these models fit credit spreads, in our case sovereign *CDS* spreads.

To set notation, we use the superscript $\mathbb{Q}$ ($\mathbb{P}$) to denote the parameters of the process $\lambda^{\mathbb{Q}}$ under the risk-neutral and historical data-generating distributions, respectively.[7]

[7]There is a bit of ambiguity in our notation here, so the fact that we are discussing the properties of $\lambda^{\mathbb{Q}}$, as a stochastic process, under two different measures, $\mathbb{Q}$ and $\mathbb{P}$, deserves emphasis. At this juncture, $\lambda^{\mathbb{P}}$, the arrival rate of default under the historical measure, is playing no role in our analysis. We comment briefly on the relation between λ^{P} and $\lambda^{\mathbb{Q}}$ in subsequent sections.

Affine Diffusion Model: Among one-factor affine models, the square-root diffusion model is the most natural choice since it guarantees a non-negative intensity process. The specific formulation we adopt is

$$d\lambda_t^{\mathbb{Q}} = \kappa^{\mathbb{P}}(\theta^{\mathbb{P}} - \lambda_t^{\mathbb{Q}})\,dt + \sigma_{\lambda^{\mathbb{Q}}}\sqrt{\lambda_t^{\mathbb{Q}}}\,dB_t^{\mathbb{P}}, \tag{4}$$

$$\eta_t = \frac{\delta_0}{\sqrt{\lambda_t^{\mathbb{Q}}}} + \delta_1\sqrt{\lambda_t^{\mathbb{Q}}} \quad \text{Market Price of Risk.} \tag{5}$$

With this choice of market price of risk, $\lambda^{\mathbb{Q}}$ follows a square-root process under both the actual measure $\mathbb{P}$ and the risk-neutral measure $\mathbb{Q}$. Specifically, under $\mathbb{Q}$,

$$d\lambda_t^{\mathbb{Q}} = \kappa^{\mathbb{Q}}(\theta^{\mathbb{Q}} - \lambda_t^{\mathbb{Q}})\,dt + \sigma_{\lambda^{\mathbb{Q}}}\sqrt{\lambda_t^{\mathbb{Q}}}\,dB_t^{\mathbb{Q}}, \tag{6}$$

where $\kappa^{\mathbb{Q}} = \kappa^{\mathbb{P}} + \delta_1\sigma_{\lambda^{\mathbb{Q}}}$ and $\kappa^{\mathbb{Q}}\theta^{\mathbb{Q}} = \kappa^{\mathbb{P}}\theta^{\mathbb{P}} - \delta_0\sigma_{\lambda^{\mathbb{Q}}}$.

This specification of η_t is non-standard within the literature on pricing defaultable securities. The standard specification has $\delta_0 = 0$ (see Dai and Singleton [2000]), and it is this case that has been considered in previous implementations of the square-root diffusion model for pricing *CDS* contracts (see, e.g., Zhang [2003] and Longstaff, Mithal, and Neis [2005]). Relaxing the constraint $\delta_0 = 0$ allows for both $\kappa\theta$ and κ to differ across $\mathbb{P}$ and $\mathbb{Q}$. Moreover, η_t can change signs over time; in the standard case, the sign of $\eta_t = \delta_1\sqrt{\lambda^{\mathbb{Q}}}$ is fixed by the sign of δ_1.[8] The generalized model rules out arbitrage if $\delta_0 \leq (\kappa^{\mathbb{P}}\theta^{\mathbb{P}})/\sigma^2_{\lambda^{\mathbb{Q}}} - .5$; see Cheridito, Filipovic, and Kimmel [2003] and Collin-Dufresne, Goldstein, and Jones [2004]. In preliminary explorations of fit, we found that the flexibility of having nonzero δ_0 and δ_1 was essential for obtaining a reasonably good fit to the *CDS* data.

An attractive feature of this model, from a computational perspective, is that the terms in (1) of the form $E_t^{\mathbb{Q}}[\exp\{-\int_t^u (r_s + \lambda_s^{\mathbb{Q}})ds\}]$ and $E_t^{\mathbb{Q}}[\lambda_u^{\mathbb{Q}}\exp\{-\int_t^u (r_s + \lambda_s^{\mathbb{Q}})ds\}]$ are known in closed form (see Duffie, Pan, and Singleton [2000]).

Lognormal Diffusion Model: An alternative description of $\lambda_t^{\mathbb{Q}}$ that also assures its non-negativity is the lognormal model:

$$d\log\lambda_t^{\mathbb{Q}} = \kappa^{\mathbb{P}}(\theta^{\mathbb{P}} - \log\lambda_t^{\mathbb{Q}})\,dt + \sigma_{\lambda^{\mathbb{Q}}}\,dB_t^{\mathbb{P}}, \tag{7}$$

$$\eta_t = \delta_0 + \delta_1\log\lambda_t^{\mathbb{Q}} \quad \text{Market Price of Risk.} \tag{8}$$

This market price of risk also allows κ and $\kappa\theta$ to differ across $\mathbb{P}$ and $\mathbb{Q}$, while assuring that λ follows a lognormal process under both measures. Specifically, under the risk-neutral measure $\mathbb{Q}$, defined by the market price of risk η_t,

$$d\log\lambda_t^{\mathbb{Q}} = \kappa^{\mathbb{Q}}(\theta^{\mathbb{Q}} - \log\lambda_t^{\mathbb{Q}})\,dt + \sigma_{\lambda^{\mathbb{Q}}}\,dB_t^{\mathbb{Q}}, \tag{9}$$

[8] See Duffee [2002] and Dai and Singleton [2002] for discussions of the importance of market prices of risk changing signs in term structure models of riskless bonds, and Cheridito, Filipovic, and Kimmel [2003] for evidence that the market price of risk (5) leads to improved fits to data on default-free bond yields over the standard specification with $\delta_0 = 0$.

where $\kappa^{\mathbb{Q}} = \kappa^{\mathbb{P}} + \delta_1 \sigma_{\lambda^{\mathbb{Q}}}$ and $\kappa^{\mathbb{Q}}\theta^{\mathbb{Q}} = \kappa^{\mathbb{P}}\theta^{\mathbb{P}} - \delta_0 \sigma_{\lambda^{\mathbb{Q}}}$.

To price *CDS* contracts with this model, we assume that r_t and $\lambda^{\mathbb{Q}}$ are independent, and then construct a discrete approximation to

$$\int_t^{t_M} E_t^{\mathbb{Q}}\left[\lambda_u^{\mathbb{Q}} e^{-\int_t^u (r_s + \lambda_s^{\mathbb{Q}}) ds}\right] du \;=\; \int_t^{t_M} D(t,u) E_t^{\mathbb{Q}}\left[\lambda_u^{\mathbb{Q}} e^{-\int_t^u \lambda_s^{\mathbb{Q}} ds}\right] du$$

in terms of survival probabilities, where $D(t,u)$ is the price of a default-free zero-coupon bond maturing at date u. The risk-neutral survival probabilities $E_t^{\mathbb{Q}}\left[e^{-\int_t^u \lambda_s^{\mathbb{Q}} ds}\right]$ are computed using a Crank-Nicholson implicit finite difference scheme.

Three-halves Diffusion Model: The third model we examine has the property that $1/\lambda^{\mathbb{Q}}$ followings a square-root diffusion:

$$d\left(\frac{1}{\lambda_t^{\mathbb{Q}}}\right) \;=\; \kappa^{\mathbb{P}}\left(\theta^{\mathbb{P}} - \frac{1}{\lambda_t^{\mathbb{Q}}}\right) dt - \sigma_{\lambda^{\mathbb{Q}}} \sqrt{\frac{1}{\lambda_t^{\mathbb{Q}}}}\, dB_t^{\mathbb{P}} \tag{10}$$

$$\eta_t \;=\; \frac{\delta_0}{\sqrt{\lambda_t^{\mathbb{Q}}}} + \delta_1 \sqrt{\lambda_t^{\mathbb{Q}}} \quad \text{Market Price of Risk.} \tag{11}$$

Given our choice of market price of risk, the implied representation of $\lambda_t^{\mathbb{Q}}$ under $\mathbb{Q}$ is:

$$d\lambda_t^{\mathbb{Q}} = \kappa^{\mathbb{Q}}(\theta^{\mathbb{Q}} - \lambda_t^{\mathbb{Q}})\lambda_t^{\mathbb{Q}}\, dt + \sigma_{\lambda^{\mathbb{Q}}}(\lambda_t^{\mathbb{Q}})^{1.5} dB_t^{\mathbb{Q}}\,, \tag{12}$$

where the diffusion coefficient depends on $\lambda^{\mathbb{Q}}$ raised to the three-halves power (and hence the name of this model). The mappings between parameters under the $\mathbb{Q}$ and $\mathbb{P}$ distributions are: $\kappa^{\mathbb{Q}} = \left(\kappa^{\mathbb{P}} - \delta_0 \sigma_{\lambda^{\mathbb{Q}}}\right)^2 / (\kappa^{\mathbb{P}}\theta^{\mathbb{P}} + \delta_1 \sigma_{\lambda^{\mathbb{Q}}}) - \sigma^2_{\lambda^{\mathbb{Q}}}$ and $\kappa^{\mathbb{Q}}\theta^{\mathbb{Q}} = \left(\kappa^{\mathbb{P}} - \delta_0 \sigma_{\lambda^{\mathbb{Q}}}\right)^3 / (\kappa^{\mathbb{P}}\theta^{\mathbb{P}} + \delta_1 \sigma_{\lambda^{\mathbb{Q}}})^2$. The fact that $1/\lambda^{\mathbb{Q}}$ follows a square-root diffusion is convenient for our estimation since it allows us to derive the likelihood function for $\lambda^{\mathbb{P}}$ and the *CDS* spreads in closed form. As with our formulation of the affine model, η_t may change signs over time if δ_0 and δ_1 have opposite signs. Moreover, the non-linear transformation $1/\lambda^{\mathbb{Q}} \rightarrow \lambda^{\mathbb{Q}}$ induces a a nonlinear drift in $\lambda^{\mathbb{Q}}$.

Comparing across models, we see they differ in the sensitivities of the volatility of $\lambda^{\mathbb{Q}}$ to its own level. The instantaneous volatility of $\lambda^{\mathbb{Q}}$ increases with the square root, level, and three-halves power of $\lambda^{\mathbb{Q}}$ in the affine, lognormal, and three-halves models, respectively. A primary motivation for comparing the fits of these models is our interest in determining which "elasticity of variance" best fits the data. These models also differ in the nature of the nonlinearity in the drifts of $\lambda^{\mathbb{Q}}$. The affine model has a linear conditional mean, the lognormal model is linear in the logarithm of the intensity (but non-linear in the level), and the three-halves model has a quadratic (instantaneous) conditional mean.

4 Loss Given Default

Beyond the specification of the default arrival intensity, a critical input into the pricing of *CDS* contracts is the risk-neutral loss rate due to a credit event, $L^{\mathbb{Q}}$. Convention within both academic analyses and industry practice is to treat this loss rate as a constant parameter of the model. In the context of pricing corporate *CDS* contracts this practice has been questioned in the light of the evidence of a pronounced negative correlation between default rates and recovery over the business cycle (see, e.g., Altman, Brady, Resti, and Sironi [2003] and related publications by the U.S. rating agencies). A business-cycle induced correlation seems less compelling in the case of sovereign risk. Indeed, a theme we consistently heard in conversations with sovereign *CDS* traders is that recovery depends on the size of the country (and the size and distribution of its external debt), but is not obviously cyclical in the same way that corporate recoveries are. In any event, we will follow industry practice and treat $L^{\mathbb{Q}}$ as a constant parameter of our pricing models, appropriately interpreted as the *expected* loss of face value on the underlying reference bond due to a credit event.

Traders are naturally inclined to call upon historical experience in setting loss rates in their pricing models. One source of this information is the agencies that rate sovereign debt issues. For example, Moody's [2003] estimates of the *recoveries* (weighted by issues sizes) on several recent sovereign defaults are: Argentina 28%; Ecuador 45%; Moldova 65%; Pakistan 48%; and Ukraine 69%. As stressed by Moody's, these numbers must be interpreted with some caution, because they are based on the market prices of sovereign bonds shortly after the relevant credit events.

That estimates of recovery may differ, depending on when market prices are sampled and perhaps also across measuring institutions, is confirmed by the recoveries estimated by *Credit Suisse First Boston* (*CSFB*), as reported in the Economist [2004]. The values at default of the bonds involved in Russia's default in May/June 1999 were 23.5% (15.9%) of face value, weighted (unweighted) by issue size. The corresponding numbers for Ecuador's default in October, 1999 were 23.4% (30.0%). Interestingly, at the time of restructuring, which in both of these cases was within a year of the default, the restructured values[9] were substantially higher. For Russia they were 36.6% (38%), and for Ecuador they were 36.2% (49.3%). Singh [2003] provides additional examples of the market prices at the time of default being depressed relative to the subsequent amounts actually recovered, and that this phenomenon was more prevalent for sovereign than for corporate credit events.

For valuing sovereign *CDS* contracts, it is the loss in value on the underlying bonds around the time of the credit event that matters for determining the payment from the insurer to the insured, regardless of whether or not these values accurately reflect the present values of the subsequently restructured debt. This fact, together with the numbers provided by $CSFB$, are consistent with the market practice of setting $L^{\mathbb{Q}}$ at 75%, or at a number somewhere near this level. Of course the relevant loss rate for pricing is $L^{\mathbb{Q}}$, and not the historical loss rate $L^{\mathbb{P}}$. Just as in many discussions of corporate bond and *CDS* pricing, the setting of $L^{\mathbb{Q}}$ based on historical experience requires the assumption that that there is no risk premium on recovery, $L^{\mathbb{Q}} = L^{\mathbb{P}}$. We take up the issue of default risk premiums in more depth in subsequent sections.

[9]This is the market value of the new bonds received as a percentage of of the original face value of the bonds.

5 Maximum Likelihood Estimates at $L^{\mathbb{Q}} = 0.75$

Initially, we follow industry practice and set $L^{\mathbb{Q}} = 0.75$. The parameters were estimated by the method of maximum likelihood, with the conditional distribution of the spreads derived from the known conditional distribution of the state. In the cases of the lognormal and affine models we used the normal (for $\log \lambda^{\mathbb{Q}}$) and non-central chi-square distributions, respectively. In the case of the three-halves model, we used the fact that $1/\lambda^{\mathbb{Q}}$ follows a square-root process, and hence is drawn from a conditional non-central chi-square distribution, to construct the likelihood function.

The five-year *CDS* contract was assumed to be priced perfectly, so that the pricing function could be inverted for $\lambda^{\mathbb{Q}}$. The one-, three-, and ten-year contracts were assumed to be priced with errors. The errors were assumed to be normally distributed with mean zero and standard deviations $\sigma_\epsilon(M)|Bid_t(M) - Ask_t(M)|$, where the $\sigma_\epsilon(M)$ are constants depending on the maturity of the contract, M. Time-varying variances that depend on the bid/ask spread allow for the possibility that the fits of our one-factor models deteriorate during periods of market turmoil when bid/ask spreads widen substantially. Conveniently, $\sigma_\epsilon(M)$ measures volatility in units of bid/ask spreads.

The risk-free interest rate (term structure) was assumed to be constant. We experimented with using a two-factor model (an $A_1(2)$ model in the nomenclature of Dai and Singleton [2000]) within the affine model for $\lambda^{\mathbb{Q}}$, but we obtained virtually identical results to those for constant r. Therefore, we proceeded with the simplifying assumption of a constant r.[10]

The *ML* estimates of the parameters and their associated standard errors are presented in Table 1. Also reported are the maximized values of the log-likelihood function (llk). Two striking patterns regarding mean reversion are evident. First, with the exception of the log-normal model for Turkey, the point estimates for $\kappa^{\mathbb{Q}}$ are negative, implying that $\lambda^{\mathbb{Q}}$ is explosive under $\mathbb{Q}$. At the same time, $\lambda^{\mathbb{Q}}$ is quite strongly mean reverting under $\mathbb{P}$ ($\kappa^{\mathbb{P}} > 0$) for all six model-country pairs. The large differences between $\kappa^{\mathbb{Q}}$ and $\kappa^{\mathbb{P}}$ (in all model-country pairs) is indicative of substantial market risk premiums related to uncertainty about future arrival rates of credit events.[11] In the presence of systematic risk about the future path of $\lambda^{\mathbb{Q}}$, it is natural that the $\mathbb{Q}$-distribution of $\lambda^{\mathbb{Q}}$ reflect a more pessimistic credit outlook than historical experience (in order for risk-neutral pricing to replicate market experience). For the models and sample period examined, the negative credit outlook under $\mathbb{Q}$ manifests itself in an explosive $\lambda^{\mathbb{Q}}$.[12]

[10] A simple arbitrage argument (see, e.g., Duffie and Singleton [2003]) shows that *CDS* spreads are approximately equal to the spreads on comparable maturity, par floating rate bonds from the same issuer as the underlying reference bonds. The prices of these bonds are not highly sensitive to the level of interest rates.

[11] Note that, for the affine model, $\hat{\delta}_0 > 0$ and $\hat{\delta}_1 > 0$ for all three countries and, therefore, there is a non-zero probability that the market price of risk η_t will change signs over time. As noted above, allowing this flexibility substantially improved the fit of the affine model. For the lognormal model, $\hat{\delta}_0 < 0$ and $\hat{\delta}_1 < 0$, so $\eta_t < 0$ throughout our sample period. Of course, the interpretations of η are different across these models, because of the logarithmic transformation.

[12] While model misspecification could also be playing a role in these estimates, we are struck by the systematic pattern of $\kappa^{\mathbb{Q}}$ being negative or only slightly positive across countries and models, even with notably different specifications of the drifts and volatilities in the three models examined.

Table 1: Maximum Likelihood Estimates at $L^{\mathbb{Q}} = 0.75$

	Affine			Lognormal			Three-halves		
	Mexico	Russia	Turkey	Mexico	Russia	Turkey	Mexico	Russia	Turkey
$\kappa^{\mathbb{Q}}$	-0.559	-0.336	-0.221	-0.037	-0.103	0.032	-36.42	-9.49	-4.54
	(.004)	(.003)	(.003)	(.004)	(.004)	(.004)	(.756)	(.153)	(.100)
$\theta^{\mathbb{Q}}\kappa^{\mathbb{Q}}$	10.6bp	11.6bp	46.2bp	0.165	0.426	-0.009	0.067	0.000	0.101
	(0.2bp)	(0.0bp)	(0.9bp)	(.017)	(.012)	(.011)	(.004)	(.002)	(.006)
$\sigma_{\lambda^{\mathbb{Q}}}$	0.202	0.169	0.209	1.26	0.712	0.822	7.524	3.280	2.700
	(.001)	(.001)	(.001)	(.008)	(.007)	(.009)	(.090)	(.044)	(.039)
$\kappa^{\mathbb{P}}$	13.7	2.80	1.61	3.20	1.51	0.844	0.149	0.004	0.002
	(7.5)	(1.54)	(1.01)	(2.20)	(1.01)	(.759)	(1.45)	(.819)	(.655)
$\theta^{\mathbb{P}}$	73bp	219bp	538bp	-5.31	-3.99	-3.38	60.81	20.02	8.72
	(19bp)	(95bp)	(236bp)	(.360)	(.429)	(.796)	(243)	(39.2)	(12.1)
$\sigma_\epsilon(1)$	1.43	1.79	1.61	1.15	1.77	1.18	1.17	1.23	1.24
	(.042)	(.070)	(.059)	(.044)	(.080)	(.034)	(.046)	(.038)	(.040)
$\sigma_\epsilon(3)$	1.39	0.576	0.707	1.05	0.558	0.504	0.891	0.865	0.532
	(.036)	(.020)	(.024)	(.036)	(.019)	(.013)	(.030)	(.031)	(.033)
$\sigma_\epsilon(10)$	0.509	0.689	0.401	0.636	1.11	0.793	0.961	1.76	0.972
	(.013)	(.014)	(.006)	(.016)	(.032)	(.024)	(.026)	(.068)	(.033)
llk	22.990	21.547	19.249	24.725	22.205	20.240	28.794	24.673	22.075

Daily data from March 19, 2001 through August 12, 2004 are used. The sample size is 856 for Mexico, and 866 for Russia and Turkey. llk is the sample average of log-likelihood.

For the affine and three-halves models to be well defined (i.e., for $\lambda^{\mathbb{Q}} > 0$), we require that $\kappa\theta > 0$, under both $\mathbb{Q}$ and $\mathbb{P}$. Since $\kappa^{\mathbb{Q}} < 0$ in these models, $\theta^{\mathbb{Q}}$ no longer has its usual interpretation as the long-run $\mathbb{Q}$-mean of $\theta^{\mathbb{Q}}$. In the lognormal model, $\lambda^{\mathbb{Q}}$ is positive under both $\mathbb{P}$ and $\mathbb{Q}$ by construction.

Turning to the magnitudes of the pricing errors for the *CDS* contracts with maturities of one, three, and ten years, the estimates of $\sigma_\epsilon(M)$ in Table 1 measure, in units of the bid/ask spread, the volatility of the pricing errors. We see that, based on this metric, no model clearly dominates the others for all countries and all contract maturities. This impression is confirmed by Figure 4 which compares the pricing errors (actual minus model-implied *CDS* spreads) for Mexico on the one-, three-, and ten-year contracts. For all three maturities, the errors for the lognormal model tend to lie between those for the affine and three-halves models. However, the ordering of the errors changes with maturity. The affine model produces the largest errors at the one-year maturity, while the three-halves model produces the largest errors for the ten-year contract. For both the three- and ten-year contracts, the pricing errors are quite close to zero, except during the summer of 2002 when *CDS* spreads widened substantially (see Figure 2). There is more evidence of persistent under or over pricing of the one-year contract, a finding that we will revisit in Section 10. Overall, however, the patterns in Figure 4 indicate that the lognormal model represents a

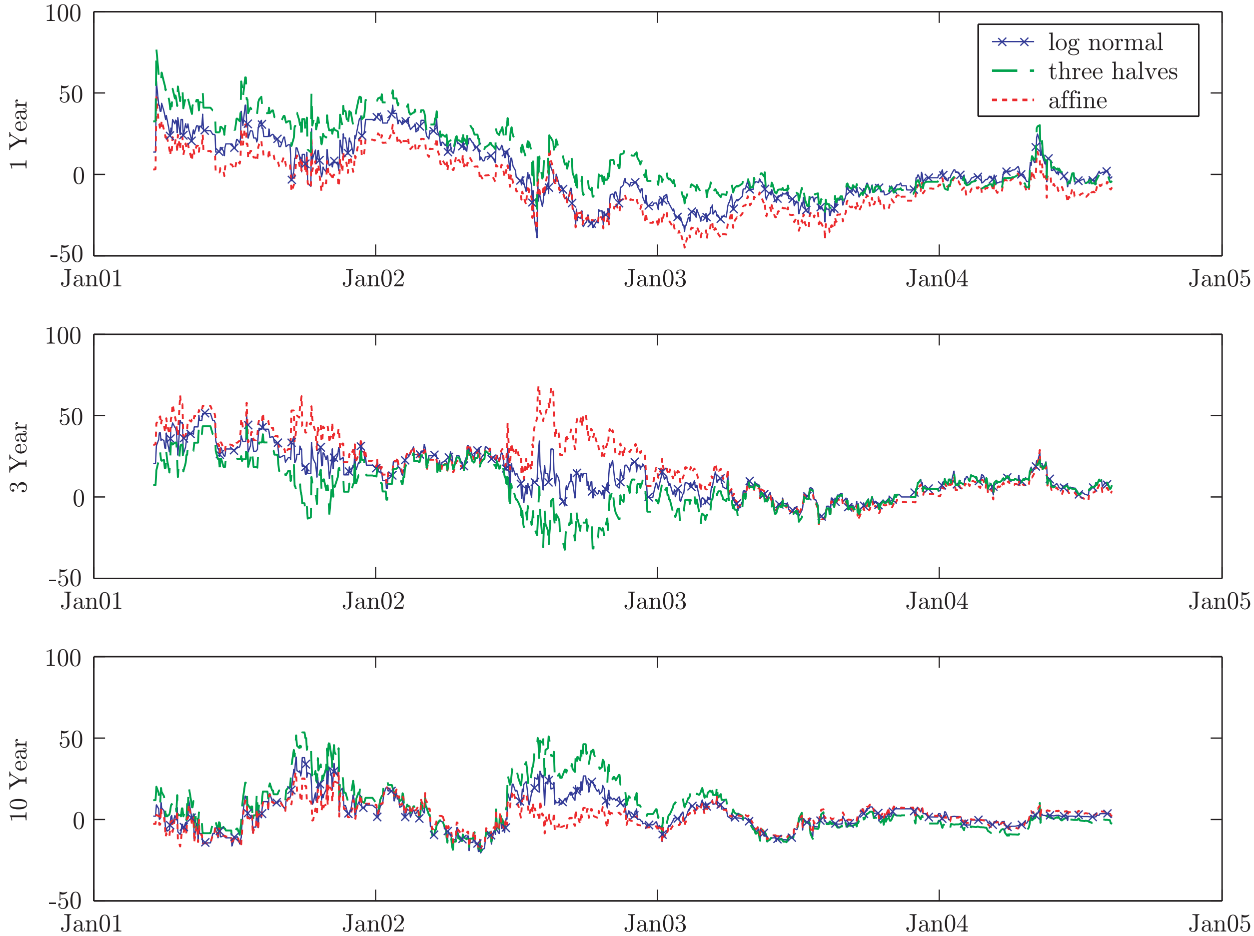

Figure 4: Pricing errors for *CDS* contracts with maturities of one, three, and ten years for Mexico, based on *ML* estimates with $L^{\mathbb{Q}} = .75$

good compromise between the three models. Similar patterns emerge from examination of the pricing errors for the contracts for Russia and Turkey and, therefore, we focus primarily on the lognormal model going forward.

6 Can We Separately Identify $\lambda^{\mathbb{Q}}$ and $L^{\mathbb{Q}}$?

Up to this point we have fixed $L^{\mathbb{Q}}$, thereby facilitating the econometric identification of the parameters governing the intensity process $\lambda^{\mathbb{Q}}$. A common impression among academics and practitioners alike is that fixing $L^{\mathbb{Q}}$ is necessary to achieve identification. This is certainly true in an economic environment in which contracts are priced under the *fractional recovery of market* value convention (RMV) introduced by Duffie and Singleton [1999]. In such a pricing framework, the product $\lambda^{\mathbb{Q}} \times L^{\mathbb{Q}}$ determines prices in the sense that the time-t

spread on a *CDS* contract takes the form

$$S_t^{RMV} = g(\lambda_t^{\mathbb{Q}} L^{\mathbb{Q}}), \tag{13}$$

for some function g. That $\lambda^{\mathbb{Q}}$ and $L^{\mathbb{Q}}$ enter symmetrically implies that they cannot be separately identified using *CDS* data alone.

In the pricing framework of *fractional recovery of face* value (RFV) (see Duffie [1998] and Duffie and Singleton [1999]), which is the most natural pricing convention for *CDS* contracts, $\lambda^{\mathbb{Q}}$ and $L^{\mathbb{Q}}$ play distinct roles. Specifically, the CDS pricing relation in (1) takes the form

$$CDS_t = L^{\mathbb{Q}} f(\lambda_t^{\mathbb{Q}}). \tag{14}$$

Comparing equation (13) against (14), we can see that the joint identification problem in the RMV framework is no longer present for *CDS* prices. For example, the explicit linear dependence of CDS_t on $L^{\mathbb{Q}}$ implies that the ratio of two *CDS* spreads on contracts of different maturities does not depend on $L^{\mathbb{Q}}$, but does contain information about $\lambda^{\mathbb{Q}}$.

Now what is conceptually true need not be true in actual implementations of these pricing models, as is illustrated by the very similar prices for par coupon bonds under the pricing conventions RMV and RFV displayed in Duffie and Singleton [1999]. To gauge the degree of numerical identification in practice, we perform the following analysis. Suppose that $\lambda^{\mathbb{Q}}$ follows a square-root process, $L^{\mathbb{Q}}$ is constant, and hence that $y_t = L^{\mathbb{Q}}\lambda^{\mathbb{Q}}$ also follows a square-root process,

$$dy_t = \kappa_y(\theta_y - y_t)\, dt + \sigma_y \sqrt{y_t}\, dB_t, \tag{15}$$

where $\kappa_y = \kappa_\lambda$, $\theta_y = L^{\mathbb{Q}}\theta_\lambda$, and $\sigma_y = \sqrt{L^{\mathbb{Q}}}\sigma_\lambda$. Using this model we ask what happens to spreads as $L^{\mathbb{Q}}$ is varied holding y fixed. For this exercise, "fixed y" means that the level of $y = L^{\mathbb{Q}}\lambda^{\mathbb{Q}}$ as well as its parameter values θ_y, κ_y, and σ_y are fixed. This, in turn, implies that any variation in $L^{\mathbb{Q}}$ is accompanied by an adjustment of $\lambda^{\mathbb{Q}} = y/L^{\mathbb{Q}}$ and its parameter values:

$$d\lambda_t^{\mathbb{Q}} = \kappa_y \left(\frac{\theta_y}{L^{\mathbb{Q}}} - \lambda_t^{\mathbb{Q}} \right) dt + \frac{\sigma_y}{\sqrt{L^{\mathbb{Q}}}} \sqrt{\lambda_t^{\mathbb{Q}}}\, dB_t. \tag{16}$$

For our illustration, the mean-reversion rate κ_y is set to 0.25, the volatility parameter σ_y is set to 0.2, and the long-run mean θ_y is set to 400*bps*.

Figure 5 illustrates the $L^{\mathbb{Q}}$-sensitivity of two securities: S^{RMV} is the credit spread of a par-coupon bond priced under the RMV framework, CDS is the *CDS* spread priced under the RFV convention in equation (1). (As usual, our *CDS* pricing takes accruals into account.) As expected, as long as the level and parameter values of y are held fixed, the value of S^{RMV} is invariant to changes in the level of $L^{\mathbb{Q}}$. In other words, $\partial S^{RMV}/\partial L^{\mathbb{Q}}|_{\text{fixed y}} = 0$.

In contrast, the *CDS* spreads are sensitive to variation in $L^{\mathbb{Q}}$. For example, in the top left-side panel, the level of y is fixed at 100 basis points (*bps*) above its long-run mean θ_y. While the one-year *CDS* spread exhibits very little sensitivity to changes in $L^{\mathbb{Q}}$, the five- and ten-year spreads show significant amounts of variation with changes in $L^{\mathbb{Q}}$. Moreover, the sensitivity of spreads to change in $L^{\mathbb{Q}}$ is higher at low levels of L^{Q}. Changing the level of y to its long-run mean (middle left-side panel) and to 100*bps* below its long-run mean

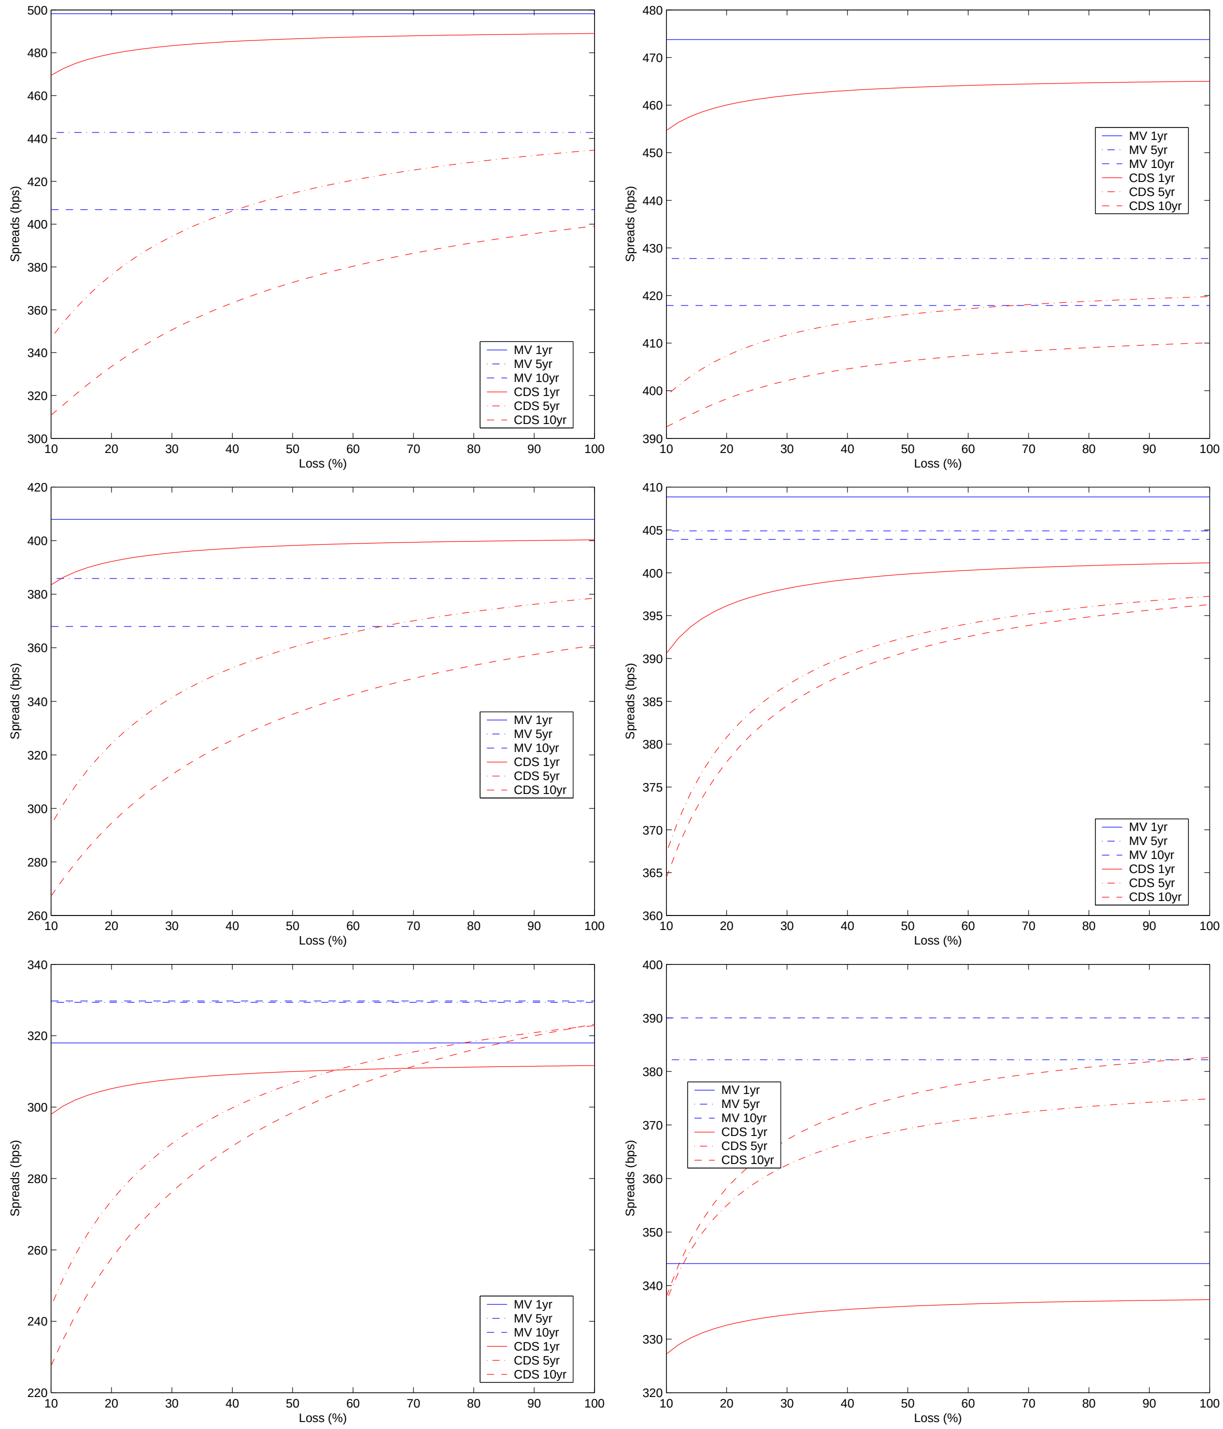

Figure 5: The sensitivity of CDS spreads to loss rate $L^{\mathbb{Q}}$, for fixed value of $y = Loss \times \lambda^{\mathbb{Q}}$. From the top to bottom panels, the levels of y are $500bps$, $400bps$, and $300bps$. For the left-side panels, $\kappa^{\mathbb{Q}} = 0.25$, while for right-side panels $\kappa^{\mathbb{Q}} = 1.0$. In all cases, $\sigma_y = 0.2$.

(bottom left-side panel), the slopes of the term structures of *CDS* spreads change, but the basic patterns of $L^{\mathbb{Q}}$-sensitivity of *CDS* spreads remain the same. Finally, comparing the left-side panels against the right-side panels, which differ only in the assumed rate of mean-reversion ($\kappa_y = 1$ on the right-side panels), we also see that the sensitivity is higher when the mean reversion rate is lower.

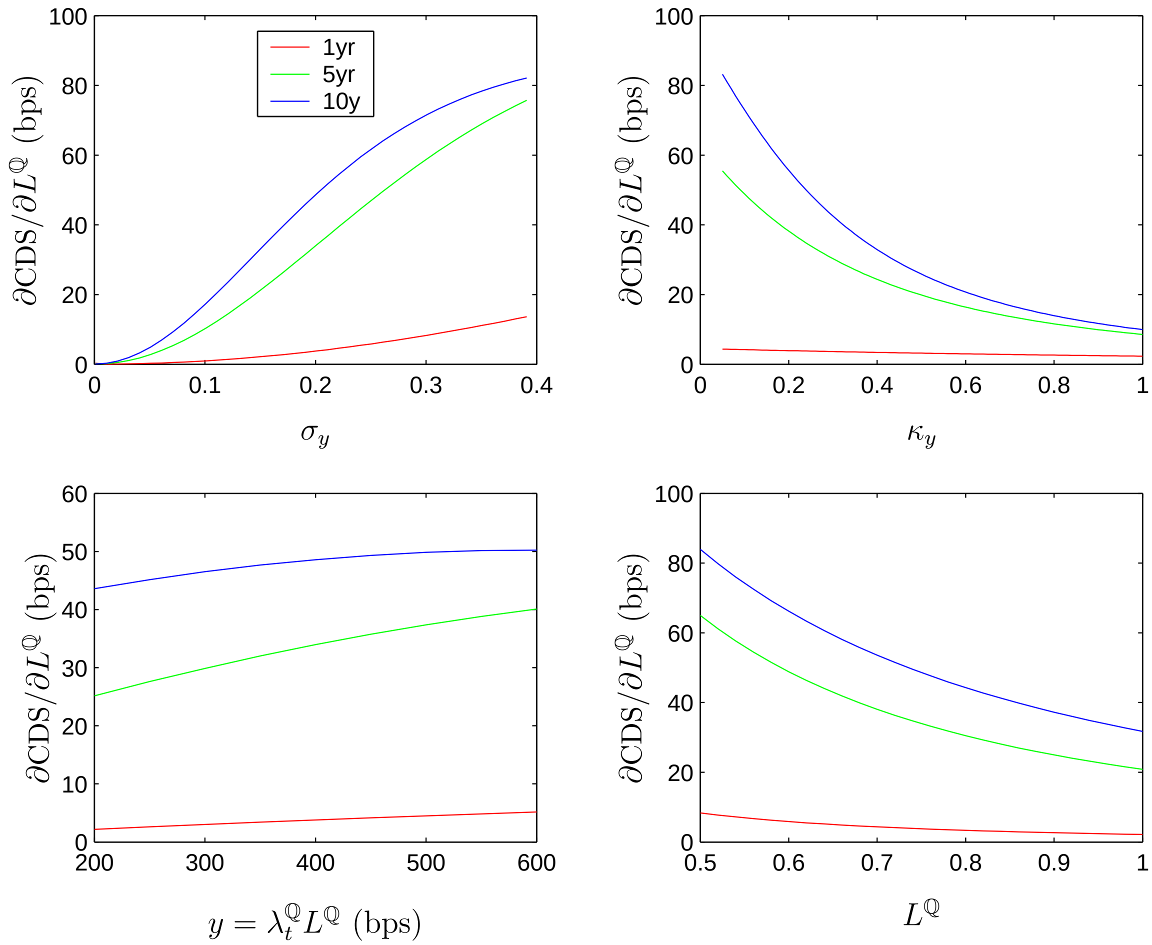

Figure 6: The partial derivative of CDS spread with respect to loss rate $L^{\mathbb{Q}}$ with fixed y: the level and parameter values of $\lambda^{\mathbb{Q}}$ are adjusted so that the process $y = Loss \times \lambda^{\mathbb{Q}}$ is kept fixed (both level and parameter values). In all figures, the base case parameters are: $\theta_y = 400$ bps, $\kappa_y = 0.25$, $\sigma_y = 0.2$, and $L^{\mathbb{Q}} = 0.75$.

An alternative means of gauging the degree of identification is by direct calculation of the partial derivative $\partial CDS/\partial L^{\mathbb{Q}}|_y$. As shown in Figure 6, $L^{\mathbb{Q}}$-sensitivity of *CDS* spreads varies substantially across the admissible regions of the parameter space. Fixing $L^{\mathbb{Q}} = 75\%$, the top two panels of Figure 6 show that $\partial CDS/\partial L^{\mathbb{Q}}|_y$ are quite sensitive to changes in volatility (σ_y) and mean-reversion (κ_y). In particular, identification is strong when either volatility is relatively high or when the mean-reversion rate is low. In contrast, $\partial CDS/\partial L^{\mathbb{Q}}|_y$

is relatively insensitive to changes in y, as illustrated in the lower left-side panel. In general, $L^{\mathbb{Q}}$-sensitivity increases markedly with the maturity of a *CDS* contract, consistent with our prior that access to the term structure of *CDS* spreads enhances the numerical identification of $L^{\mathbb{Q}}$ separately from the parameters governing $\lambda^{\mathbb{Q}}$.

7 Maximum Likelihood Estimates including $L^{\mathbb{Q}}$

Encouraged by these observations about the feasibility of identifying both the recovery and default parameters, we proceed to compute *ML* estimates with $L^{\mathbb{Q}}$ treated as an unconstrained parameter.[13] Relaxing the constraint on $L^{\mathbb{Q}}$ also tends to lower the pricing errors (in the sense of producing smaller values of $\sigma_\epsilon(M)$), particularly for the case of Turkey. However, there is not a uniform reduction in $\sigma_\epsilon(M)$ for all countries and models. A more formal assessment of fit comes from comparing the values of *llk* for the constrained and unconstrained models. For example, in the case of Mexico under the lognormal model, *llk* is 24.725 and 24.906 with and without the constraint on $L^{\mathbb{Q}}$ imposed. The implied likelihood-ratio test statistic is 309, well above the 1% critical value of 10.8 for a $\chi^2(1)$. This rejection of the constraint $L^{\mathbb{Q}} = 0.75$ at conventional significance levels obtains for all three countries and for both the lognormal and affine models.

The estimates for $L^{\mathbb{Q}}$ are the most intriguing aspect of these unconstrained estimates. As discussed in Section 4, while estimates of recovery may differ, the market practice is to set $L^{\mathbb{Q}}$ at or around the 75% level. By contrast, our estimated values of $L^{\mathbb{Q}}$ are in the range of 16% to 44% for the lognormal and affine models. Accompanying the lower estimated $L^{\mathbb{Q}}$ (relative to the constrained value of 0.75) are substantial increases in the estimated value of $\theta^{\mathbb{Q}}\kappa^{\mathbb{Q}}$. Given that the corresponding estimates $\kappa^{\mathbb{Q}}$ are similar across the constrained and unconstrained models, these observations imply that the lower estimated value of $L^{\mathbb{Q}}$ match up with much larger estimates of the base arrival rate of credit events that applies when $\lambda^{\mathbb{Q}}$ is near zero. This finding is intuitive: in order to price *CDS* contracts as if agents are risk neutral, a lower loss rate in the event of default (a positive development for investors) must be accompanied by a higher risk-neutral base arrival rate of credit events.

For all of the country/model pairs, $\lambda^{\mathbb{Q}}$ continues to follow a $\mathbb{Q}$-explosive process, with the exception of the lognormal model for Turkish data. This explosive behavior is also an integral part of fitting the term structure of sovereign *CDS* spreads. When we constrained $\kappa^{\mathbb{Q}} \geq 0.001$, the likelihood search algorithm was driven to the boundary of the parameter space, and this was accompanied by a significant reduction in *llk*. Moreover, for several country/maturity pairs, and in both the affine and lognormal models, imposing this constraint led to increases in the standard deviations of the pricing errors, $\sigma_\epsilon(M)$.

Nevertheless, it is interesting to note the consequences for $L^{\mathbb{Q}}$ of imposing the constraint $\kappa^{\mathbb{Q}} > 0$. In the lognormal models, for example, the estimated values of $L^{\mathbb{Q}}$ increase from 21.4% to 73.2% for Mexico and from 36.1% to 100% for Russia. For the affine models of all three countries the estimated values of L^{Q} reach the boundary of 100% when $\lambda^{\mathbb{Q}}$ is

[13] All other aspects of the estimation remain as in the constrained case, with $L^{\mathbb{Q}} = 0.75$.

Table 2: Maximum likelihood estimates of the *log-normal* and *affine* models with all parameters including $L^{\mathbb{Q}}$.

	Lognormal					Affine					
	Unconstrained			Stationary		Unconstrained			Stationary		
	Mexico	Russia	Turkey	Mexico	Russia	Mexico	Russia	Turkey	Mexico	Russia	Turkey
$\kappa^{\mathbb{Q}}$	-0.098	-0.135	0.010	≈ 0	≈ 0	-0.614	-0.353	-0.358	≈ 0	≈ 0	≈ 0
	(.005)	(.005)	(.004)	(NA)	(NA)	(.006)	(.004)	(.005)	(NA)	(NA)	(NA)
$\theta^{\mathbb{Q}}\kappa^{\mathbb{Q}}$	0.690	0.571	0.348	-0.0046	0.040	48bp	18bp	351bp	33bp	56bp	121bp
	(.022)	(.017)	(.010)	(.076)	(.026)	(1bp)	(1bp)	(5bp)	(0.2bp)	(1bp)	(2bp)
$\sigma_{\lambda^{\mathbb{Q}}}$	0.900	0.556	0.689	1.33	0.831	0.181	0.165	0.322	0.101	0.123	0.206
	(.021)	(.008)	(.006)	(.035)	(.016)	(.002)	(.001)	(.002)	(.001)	(0.001)	(0.002)
$\kappa^{\mathbb{P}}$	1.66	1.15	0.339	2.89	1.35	1.84	2.06	0.80	1.23	1.37	1.00
	(1.18)	(.640)	(.444)	(2.01)	(0.88)	(1.22)	(0.85)	(0.42)	(0.68)	(0.40)	(0.50)
$\theta^{\mathbb{P}}$	-4.12	-3.38	-3.10	-5.39	-4.41	316bp	328bp	1064bp	106bp	154bp	364bp
	(.422)	(.397)	(2.05)	(.450)	(.533)	(112bp)	(130bp)	(663bp)	(45bp)	(58bp)	(190bp)
$L^{\mathbb{Q}}$	0.214	0.361	0.235	0.732	1	0.161	0.439	0.237	1	1	1
	(.010)	(.013)	(.005)	(.102)	(.085)	(.002)	(.006)	(.002)	(0.01)	(0.01)	(0.01)
$\sigma_{\epsilon}(1)$	1.17	1.65	1.02	1.21	1.95	1.22	1.73	1.09	6.73	4.46	1.88
	(.045)	(.077)	(.026)	(.046)	(.088)	(.049)	(.075)	(.025)	(.206)	(.224)	(.062)
$\sigma_{\epsilon}(3)$	1.06	0.581	0.499	0.933	0.653	1.37	0.560	0.509	2.09	1.48	0.544
	(.039)	(.021)	(.014)	(.033)	(.022)	(.042)	(.021)	(.014)	(.071)	(.072)	(.014)
$\sigma_{\epsilon}(10)$	0.588	1.08	0.524	0.736	1.73	0.535	0.746	0.552	0.572	1.69	1.25
	(.015)	(.029)	(.011)	(.021)	(0.073)	(.013)	(.017)	(.012)	(.014)	(.054)	(.030)
llk	24.906	22.323	20.725	24.673	21.539	23.729	21.662	19.619	21.591	18.992	18.262

In the "Stationary" cases, $\kappa^{\mathbb{Q}}$ is constrained to be positive. Daily data from March 19, 2001 through August 12, 2004 are used. The sample size is 856 for Mexico, and 866 for Russia and Turkey. llk is the sample average of log-likelihood.

constrained to be $\mathbb{Q}$-stationary. This finding, again, reflects an intuitive trade-off: if $\lambda^{\mathbb{Q}}$ is $\mathbb{Q}$-stationary, then survival probabilities are higher (compared to the explosive case). For risk-neutral pricing to match market prices, $L^{\mathbb{Q}}$ must be increased, substantially according to the results in Table 2.

From the perspective of statistical inference, the likelihood ratio statistics suggest that the global optimum is obtained with a low loss rate and high mean arrival rate of credit events. However, as noted above, relaxing the constraint that $L^{\mathbb{Q}} = 0.75$ does not lead to a marked reduction in the estimated $\sigma_\epsilon(M)$ for most countries and maturities. Furthermore, due to the recent maturation of the sovereign *CDS* markets, we are working with an undesirably short sample,[14] one that does not include many of the most prominent sovereign credit events during the past decade. Therefore, based on the evidence presented so far, it seems premature to strongly favor the unconstrained estimates over those with $L^{\mathbb{Q}} = 0.75$ (or some similarly large number). To shed further light on the properties of models with (low $L^{\mathbb{Q}}$, high $\theta^{\mathbb{Q}}\kappa^{\mathbb{Q}}$) versus (high $L^{\mathbb{Q}}$, low $\theta^{\mathbb{Q}}\kappa^{\mathbb{Q}}$), we investigate the model-implied survival probabilities and options prices below.

Finally, a natural question at this juncture is whether, with sample sizes that are available in the *CDS* markets, one can reasonably expect to estimate $L^{\mathbb{Q}}$ precisely. To address this question we conduct a small-scale Monte-Carlo exercise. Specifically, we simulate affine model-implied one-, three-, five-, and ten-year *CDS* spreads, and add normally distributed pricing errors to the one-, three- and ten-year *CDS* spreads. The resulting (noisy) simulated *CDS* data is then used to construct *ML* estimates of the underlying parameters. This was repeated one-hundred times, and the means and standard deviations of the *ML* estimates are displayed in Table 3. The true parameter values are set close to the $L^{\mathbb{Q}} = 0.75$ case for Russia with two exceptions. First, to gauge the effect of $\kappa^{\mathbb{Q}} < 0$, we consider two cases, one with explosive $\mathbb{Q}$-intensity ($\kappa^{\mathbb{Q}} < 0$), and the other with stationary $\mathbb{Q}$-intensity ($\kappa^{\mathbb{Q}} > 0$). Second, to simplify the estimation process, we use a common coefficient $\sigma_\epsilon(M)$ for the volatility of one-, three-, and ten-year *CDS* pricing errors.

The standard deviations of the simulated estimates are of the same orders of magnitude as the standard errors reported from our *ML* results, and the means of the simulated estimates are close in magnitude to the true parameter values. Moreover, for econometric identification, whether or not the default intensity is $\mathbb{Q}$-explosive appears to be inconsequential. The degree of persistence in $\kappa^{\mathbb{Q}}$ matters, of course, as was documented in Figure 6, but so long as $\lambda^{\mathbb{Q}}$ is reasonably persistent the likelihood function appears to exhibit sufficient curvature for reliable estimation of $L^{\mathbb{Q}}$. There is a downward small-sample bias, but it very small relative to the difference between 0.75 and the low estimated levels of $L^{\mathbb{Q}}$ for the unconstrained models displayed in Table 2.

[14] For the reasons discussed below, our concern is not so much with the length of the sample per se, but rather about the possibility that our sample period does not adequately reflect the relative frequencies of periods with high and low probabilities of credit events.

Table 3: Simulation results for the *affine* model

	$\theta^{\mathbb{P}}$	$\kappa^{\mathbb{P}}$	$\sigma_{\lambda^{\mathbb{Q}}}$	$\kappa^{\mathbb{Q}}$	σ_{ϵ}	$L^{\mathbb{Q}}$	$\theta^{\mathbb{Q}}\kappa^{\mathbb{Q}}$
explosive case							
true param	219bp	2.7880	0.1691	-0.3361	0.5000	0.7500	12bp
mean(estm)	224bp	3.1417	0.1704	-0.3458	0.5043	0.7265	12bp
std(estm)	41bp	0.8002	0.0007	0.0017	0.0069	0.0278	1bp
stationary case							
true param	219bp	2.7880	0.1691	0.1000	0.5000	0.7500	611bp
mean(estm)	232bp	3.2271	0.1711	0.0848	0.5046	0.7148	633bp
std(estm)	55bp	0.9935	0.0044	0.0073	0.0074	0.0135	7bp

Simulations are performed under the "true" parameter values with the same sample size as that of our *CDS* data. The mean and standard deviation of the estimates are calculated with 100 simulation runs.

8 On Priced Risks in Sovereign *CDS* Markets

The large differences between the properties of $\lambda^{\mathbb{Q}}$ under the risk-neutral and the actual measure suggest that there is a systematic risk related to changes in future arrival rates of sovereign credit events that is priced in the *CDS* market. To shed further light on the nature of this risk, we begin with a more in depth exploration of the magnitudes of market risk premia, as reflected in the model-implied survival probabilities. Then we examine the correlations among the model-implied $\lambda^{\mathbb{Q}}$ and their changes over time with credit market conditions. Finally, we relate the co-movements among the credit-event intensities $\lambda^{\mathbb{Q}}$ to the historical patterns in global default rates and market volatility.

Within the reduced form pricing framework adopted in this study, the risk-neutral survival probability is given by

$$\mathcal{S}_t(M) \equiv E_t^{\mathbb{Q}}\left[e^{-\int_t^{t+M} \lambda^{\mathbb{Q}}(u)\,du}\right], \tag{17}$$

for horizon M. It is these survival probabilities that effectively determine the discounting of the cash flows on the *CDS* contract.[15] Figure 7 displays the five- and ten-year $\mathbb{Q}$-survival probabilities for all three countries, based on the lognormal model with $L^{\mathbb{Q}} = 0.75$. The ordering of these probabilities is consistent with the countries' ratings. Russia shows substantial improvement in credit quality (higher survival probabilities) over the sample. Nevertheless, on a risk-neutral basis, Russia's survival probabilities (implicit in the *CDS* spreads) are notably lower than Mexico's. This is true even at the end of our sample period when Russia was

[15] More precisely, the insurance premiums are discounted by $E_t^{\mathbb{Q}}[e^{-\int_t^s (r(u)+\lambda^{\mathbb{Q}}(u))\,du}]$. If the short-rate and default processes are independent, then this expression factors into $D(t, t+M)\mathcal{S}_t(M)$. In our pricing framework, r is constant, so temporal variation in the credit spreads used in discounting is determined entirely by the term structure of risk-neutral survival probabilities $\mathcal{S}_t(M)$.

rated Baa3 while Mexico was only one notch higher at Baa2. Even identical ratings would not imply identical $\mathbb{Q}$ survival probabilities, of course, because (i) ratings are relatively static measures of credit risk; and (ii) two countries with the same $\mathbb{P}$ credit risk may have different risk-neutral survival probabilities. It seems plausible that the risk premiums associated with $\lambda^{\mathbb{Q}}$ were different for Mexico and Russia, given their different histories.

The $\mathbb{Q}$-survival probabilities for Turkey are substantially below those for Mexico and Russia. In addition, these probabilities are low at the beginning of our sample and dip down in the summer of 2002, consistent with the negative outlook from Moody's. There is also a sizeable spike downwards in March, 2003 when the discussions with Cyprus deteriorated.

To compare the implied $\mathbb{Q}$-survival probabilities across models, we focus on the case of Mexico. Figure 8 displays the results for the lognormal (LN) and affine (SR) models. For the constrained models with $L^{\mathbb{Q}} = 0.75$, the model-implied survival probabilities are virtually on top of each other. We do not display $\mathcal{S}_t(s)$ for the three-halves model, because (for the case of $L^{\mathbb{Q}} = 0.75$) it also is virtually on top of the two cases that are displayed. In part, this similarity is induced by the inversion of the model for the implied levels of $\lambda^{\mathbb{Q}}$. In all cases, $\lambda^{\mathbb{Q}}$ is chosen to fit the five-year *CDS* rate perfectly, and this rate is determined by the survival probabilities. However, these patterns are also influenced, through the present value formula (17), by the autocorrelation and volatility properties of $\lambda^{\mathbb{Q}}$. Evidently, though the estimated parameters are quite different across models, their combined effect is to produce the same survival probabilities.

When $L^{\mathbb{Q}}$ is treated as a free parameter, we get some separation between the survival probabilities in the lognormal and affine models. This is induced largely by the differences in the estimated $L^{\mathbb{Q}}$. A lower loss rate is associated with a lower $\mathbb{Q}$-survival probability, as expected.

Another interesting calculation is the pseudo-survival probability

$$\mathcal{PS}_t(M) \equiv E_t^{\mathbb{P}}\left[e^{-\int_t^{t+M} \lambda^{\mathbb{Q}}(u)\,du}\right]. \tag{18}$$

If market participants are neutral towards the risk of variation over time in $\lambda^{\mathbb{Q}}$, then these expectations should replicate the risk-neutral survival probabilities. In other words, a comparison of $\mathcal{S}_t(s)$ and $\mathcal{PS}_t(s)$ gives us a graphical assessment of the quantitative importance of the risk premiums associated with unpredictable variation over time in $\lambda^{\mathbb{Q}}$.[16] Figure 9 displays $\mathcal{PS}_t(M)$ for $M = 5$ years and the case of $L^{\mathbb{Q}} = 0.75$. The time series of $\mathcal{PS}_t(5)$ are much smoother than those of $\mathcal{S}_t(5)$, because of the relatively high rates of mean reversion of $\lambda^{\mathbb{Q}}$ under $\mathbb{P}$ versus under $\mathbb{Q}$. Additionally, comparing across $\mathbb{P}$ distributions, $\mathcal{PS}_t(5)$ for Mexico is much smoother than the counterpart for Turkey because of the relatively high rate of mean reversion of $\lambda^{\mathbb{Q}}$ under $\mathbb{P}$ for Mexico versus Turkey.

The differences between $\mathcal{S}_t(M)$ and $\mathcal{PS}_t(M)$ are large for all three countries, particularly during the second halves of 2001 and 2002. Furthermore, again for all three countries, there is a substantial narrowing of these differences towards the end of our sample period.

[16] We stress that, in neither case, are we computing the historical survival probabilities $E_t^{\mathbb{P}}[e^{-\int_t^s \lambda^{\mathbb{P}}(u)\,du}]$. For this purpose we would need data on the physical intensity $\lambda^{\mathbb{P}}$ and, as emphasized by Jarrow, Lando, and Yu [2005] and Yu [2002], this information cannot be extracted from bond or *CDS* spread data alone.

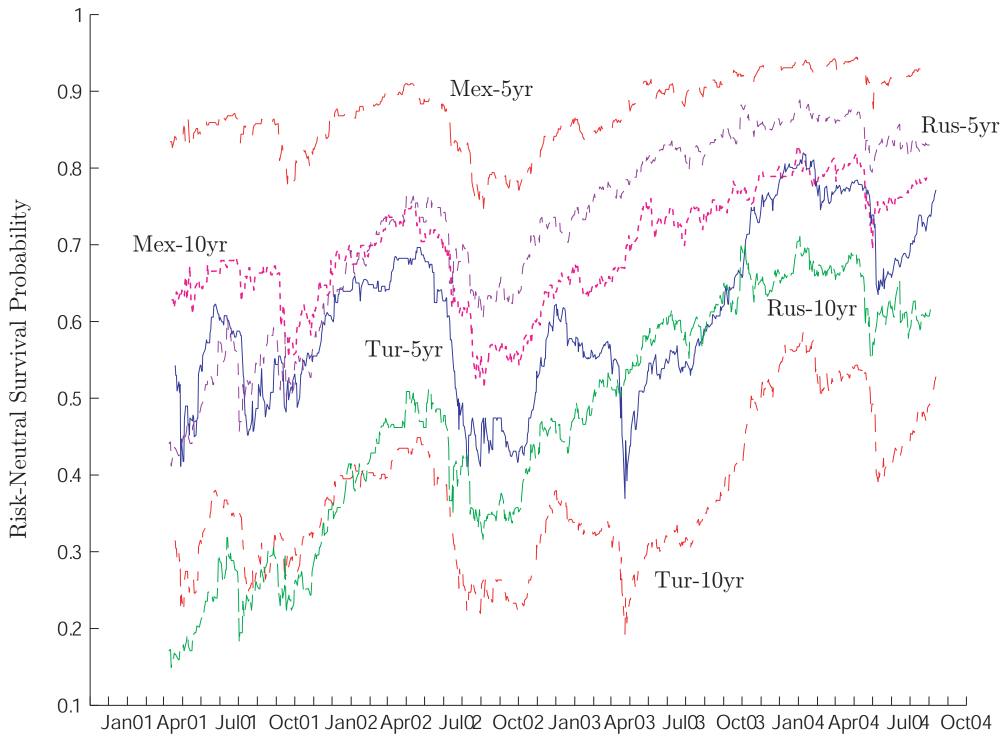

Figure 7: $\mathbb{Q}$-Survival: Lognormal Models, $L^{\mathbb{Q}} = 0.75$

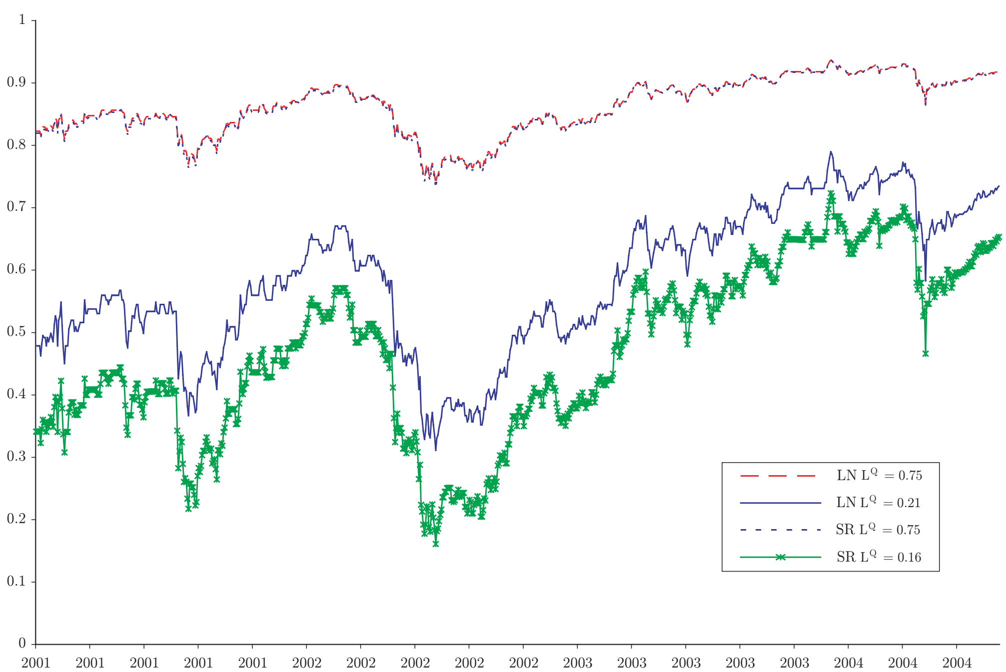

Figure 8: $\mathbb{Q}$-survival probabilities for the various model for Mexico. LN denotes lognormal and SR denotes the affine square-root diffusion. Loss rates are indicated as $L^{\mathbb{Q}}$.

This suggests that, concurrent with the narrowing of spreads late in our sample (see Figure 2), market risk premiums related to uncertainty about the risk-neutral arrival rate of credit events declined substantially. Around the beginning of 2004, the effects of these risk premiums on (5 year) survival probabilities were very close to zero.

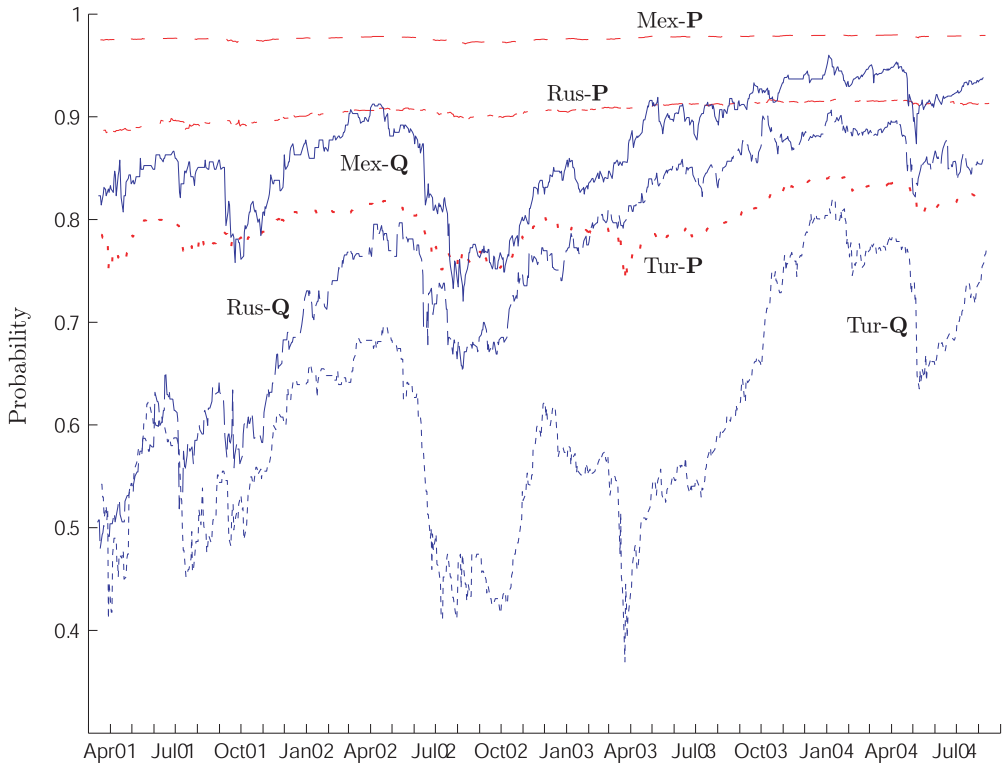

Figure 9: $E_t^{\mathbb{P}}\left[e^{-\int_t^{t+5yr}\lambda^{\mathbb{Q}}(u)\,du}\right]$, Lognormal $\lambda^{\mathbb{Q}}$, $L^{\mathbb{Q}} = 0.75$

Interestingly, there was then a big "jump" in the effects of risk premiums during May, 2004, concurrently with the jump up in *CDS* spreads. This is an interesting episode, because the economic forces in play had a simultaneous and substantial effect on the spreads of all three countries (Figure 2). During the second quarter of 2004 there was a substantial increase in non-farm payrolls in the U.S. This, combined with comments by representatives of the Federal Reserve, led market participants to expect a tightening of monetary policy. A reason that these concerns had a substantial and widespread effect on spreads is that both financial institutions and hedge funds had substantial positions in "carry trades" in place. They were borrowing short-term in dollars and investing in long-term bonds, often high-yield and emerging market bonds issued in various currencies. The unexpected strength in the U.S. economy led to an unwinding of some of these trades and, consequently, an across the board adjustment in spreads on corporate and sovereign credits.[17] This episode illustrates

[17]These concerns were widely noted in the media at the time. "In a single day, May 7, yields on Brazilian bonds jumped 1.52 percentage points as the unexpectedly strong jobs report in the U.S. increased the likelihood of higher short-term rates. (Henry [2004])." See also the discussion in Cogan [2005].

the importance of changes in investors' appetite for exposure to credit, as a global risk class, for co-movements in yields. The induced changes in yields on the sovereign credits examined here had nothing directly to do with the inherent credit qualities of the issuers.

Comparison of Figures 3 and 9 suggests that episodes when the risk premiums associated with variation in $\lambda^{\mathbb{Q}}$ were large (as measured by the difference between $\mathcal{S}_t(M)$ and $\mathcal{PS}_t(M)$) were also episodes when the bid/ask spreads on the *CDS* contracts were large, particularly for the case of Turkey. These patterns suggest that deteriorations in credit quality may have been accompanied by deteriorations in liquidity in the Turkish sovereign bond or *CDS* markets, or both.[18] The changes in bid/ask spreads for Russia and Turkey were much smaller and, for these countries, the ratio $(ask - bid)/bid$ remained between 5 and 10% for most of our sample (Mexico experienced somewhat higher proportionate spreads during the first quarter of 2002). Therefore, changes in liquidity in the *CDS* markets do not appear to have been a major source of the variation in survival probabilities exhibited in Figures 7 and 9.

Over the entire sample period, the pair-wise correlations among the (logarithms of the) model-implied intensities $\lambda^{\mathbb{Q}}$ were $corr(Mex, Tur) = 85.3\%$, $corr(Mex, Rus) = 81.5\%$, and $corr(Rus, Tur) = 69.3\%$. These magnitudes are roughly consistent with those documented for sovereign bond spreads (e.g., Rigobon [2001], Yafeh, Sussman, and Mauro [2001], Cifarelli and Paladino [2002]). Also of interest are the correlations during subperiods when a sovereign issuer had different credit ratings (though with the usual caveats about reduced sub-sample sizes). For example, splitting our sample based on Russia's upgrade to investment grade in late 2003, we find that $corr(Rus^{SG}, Tur^{SG}) = 33.7\%$ and $corr(Mex^{IG}, Rus^{SG}) = 64.4\%$ when Russia had a speculative grade (SG) rating, and $corr(Rus^{IG}, Tur^{SG}) = 76.2\%$ and $corr(Mex^{IG}, Rus^{IG}) = 72.2\%$ when Russia had an investment grade (IG) rating. While there is substantial evidence that correlations among bond yields change with issuers' ratings (e.g., Kaminsky and Schmukler [2001] and Rigobon [2002]), the patterns in our data are perhaps more naturally interpreted as a consequence of changes in credit market conditions. During the early sub-sample, when Russia was rated SG, credit spreads were relatively large and Turkey experienced several local "credit shocks" (see Figure 2). It is not surprising then that $corr(Rus^{SG}, Tur^{SG})$ is relatively low. On the other hand, during the later sample, credit conditions had improved globally and spreads were narrower and less volatile for all three countries. The finding that $corr(Mex^{IG}, Rus^{IG})$ is only marginally larger than $corr(Mex^{IG}, Rus^{SG})$ is consistent with this interpretation.

To shed further light on the nature of the systematic, priced risk underlying the co-movements among *CDS* spreads, we examined the correlations between the model-implied $\lambda^{\mathbb{Q}}$ and the historical defaulted bond debt on a global basis, as computed by Moody's. Figures 10(a) and 10(b) display the global default dollars against the default intensity $\lambda^{\mathbb{Q}}$, both standardized to zero mean and unit variance, for Mexico and Turkey. There is clearly a

[18] Concurrent movements in liquidity and credit quality is often observed in credit markets. As shown by Duffie and Singleton [1999], the pricing formulas we use can be adapted to accommodate liquidity risk by adjusting the discount rate from $r_t + \lambda_t^{\mathbb{Q}}$ to $r_t + \lambda_t^{\mathbb{Q}} + \ell_t$, where ℓ_t is a measure of liquidity costs. Longstaff, Mithal, and Neis [2005] use this extended framework in their analysis of corporate bond and *CDS* contracts. They assume that $\ell_t = 0$ in their pricing of corporate *CDS* contracts or, equivalently, that *CDS* spreads are driven nearly entirely by variation in $\lambda^{\mathbb{Q}}$.

tendency for the default intensity to be wide during periods of above average global defaults, and narrow when defaults are low. Additionally, consistent with the preceding discussion, the correlations between the $\lambda^{\mathbb{Q}}$'s and the global default series is higher during the recent period when default rates and spreads were both low relative to their means. Default intensities appear to be wider in the month subsequent to the release of the global default dollars, an impression that is confirmed by regression analysis.

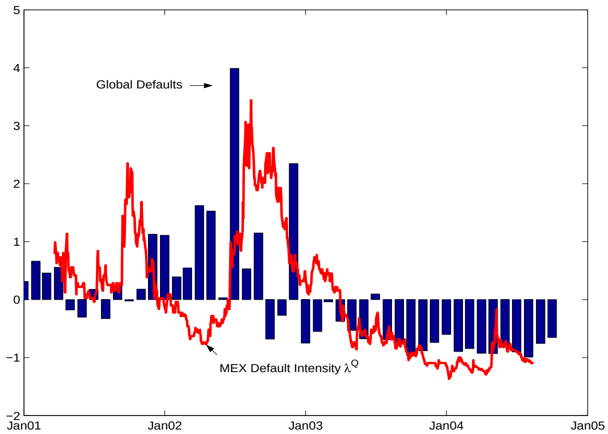

(a) Mexican default intensity versus global defaults

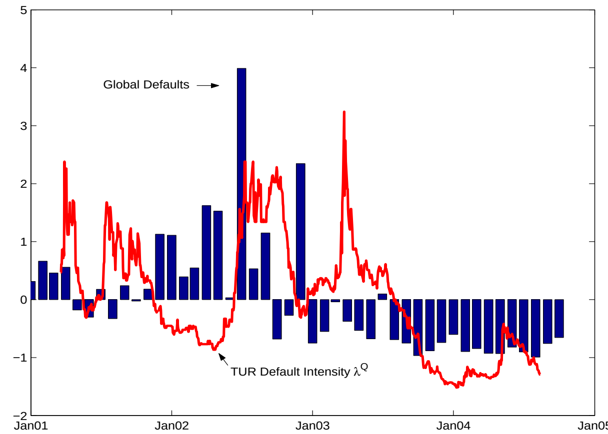

(b) Turkish default intensity versus global defaults

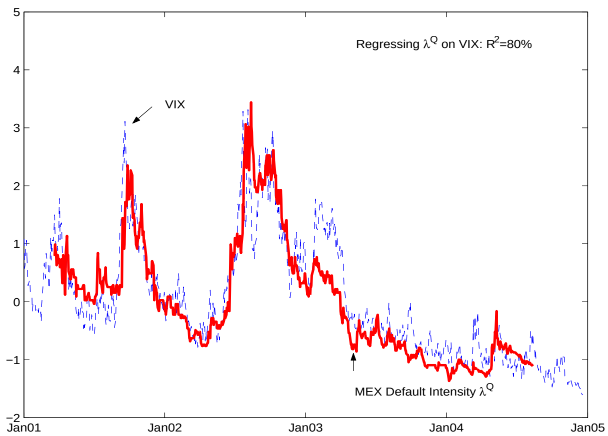

(c) Mexican default intensity versus *VIX*

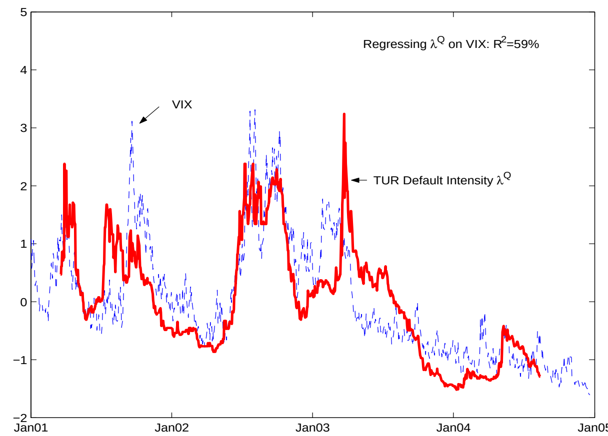

(d) Turkish default intensity versus *VIX*

Figure 10: Figures (a) and (b) display default intensity $\lambda^{\mathbb{Q}}$ against Moody's global dollars of defaulted bond debt. Figures (c) and (d) display the *VIX* volatility index versus default intensity $\lambda^{\mathbb{Q}}$. All of the series are standardized to mean zero and unit variance.

Perhaps even more striking is the degree of correlation between the default intensity $\lambda^{\mathbb{Q}}$ and the *VIX* volatility index, an index of volatility in the U.S. equity markets (see Figures 10(c) and 10(d)). The correlations among the spreads and *VIX* are 89.5% and 77.1% for Mexico and Turkey, respectively.[19] A strong correlation between *VIX* volatility and

[19]The correlation between the Russian default intensity and the *VIX* index is smaller, about 54%. This is

U.S. corporate credit spreads has been extensively documented (see, e.g., Collin-Dufresne, Goldstein, and Martin [2001] and Schaefer and Strebulaev [2004]). That *VIX*, a domestic equity volatility index, is also highly correlated with spreads on sovereign entities as widely dispersed as Mexico and Turkey suggests that *VIX* is a proxy for global "event risk" in credit markets.

There are also several notable periods when default intensity did not closely track either the global defaults or the *VIX* index, particularly in the case of Turkey. (This is also true of Russia, which is not shown here.) As noted above, the period in early 2003, for example, when Turkish *CDS* spreads and default intensity widened out substantially, was when talks between Turkey and Cyprus deteriorated. From Figure 2 it is seen that this period was also one when there was inversion in the *CDS* term structure for Turkey. We suspect that it is the influence of these country-specific influences, and the associated sharp inversions and subsequent reversions to upward sloping curves, all within a matter of months, that explains why Turkey shows the least tendency towards a $\mathbb{Q}$-explosive $\lambda^{\mathbb{Q}}$.

9 Pricing Options on *CDS* Contracts

Model-implied option prices are potentially informative in assessments of fit of the various models examined for two reasons.[20] First, we assess the potential role of options in assisting with the separate identification of $L^{\mathbb{Q}}$ and $\lambda^{\mathbb{Q}}$. As we will see, the payoff structure on a *CDS* option is such that there is a linear dependence on $L^{\mathbb{Q}}$, but a non-linear dependence on $\lambda^{\mathbb{Q}}$ that is above and beyond the non-linearity inherited from the *CDS* pricing. Second, the optionality inherent in these contracts ties their pricing closely to the conditional volatility of the default intensity. The latter, in turn, is affected by both the volatility coefficient σ_λ and the risk-neutral mean-reversion coefficient κ^Q. Of particular interest is whether the model-implied option prices shed light on the relative plausibility of $\mathbb{Q}$-explosive versus $\mathbb{Q}$-stationary intensity processes.

A European-style payers option gives the buyer the right but not the obligation to buy protection on the underlying. For example, at the time of its expiration τ, a payers option on an M-year CDS will be exercised with a profit if the CDS price at that time is more expensive than the pre-arranged strike price K. Otherwise, the option will expire worthless. Letting $\mathcal{V}(\tau, M, CDS_\tau(M), L^{\mathbb{Q}})$ denote the value at date τ of a *CDS* with maturity M, premium $CDS_\tau(M)$, and loss rate $L^{\mathbb{Q}}$, the option payoff at time τ is

$$\left[\mathcal{V}(\tau, M, CDS_\tau(M), L^{\mathbb{Q}}) - \mathcal{V}(\tau, M, K, L^{\mathbb{Q}})\right]^+, \tag{19}$$

where, at the time of expiration, $\mathcal{V}$ is obtained by plugging the premium ($CDS_\tau(M)$ or K) into the pricing formula (1).

a reflection of the the fact that Russian spreads were largely trending downwards during our sample period, unlike the *VIX* index.

[20] Credit options now trade actively on indices in the investment grade and high yield market, as well as on some liquid single names. While they are not currently actively traded on many sovereign names, markets are developing.

Making the substitutions for $\mathcal{V}$ from (1), and recognizing that the second term involving $L^{\mathbb{Q}}$ is common to both terms, the payoff (19) simplifies to

$$[CDS_\tau(M) - K]^+ \; g(\lambda_\tau^{\mathbb{Q}}), \tag{20}$$

where the "duration" $g(\lambda_\tau^{\mathbb{Q}})$ of the contract is given by

$$g(\lambda_\tau^{\mathbb{Q}}) \equiv \frac{1}{2}\sum_{j=1}^{2M} E_\tau^{\mathbb{Q}}\left(e^{-\int_\tau^{\tau+.5j}(r_s+\lambda_s^{\mathbb{Q}})ds}\right). \tag{21}$$

Finally, upon noting that $CDS_\tau = L^{\mathbb{Q}} f(\lambda_\tau^{\mathbb{Q}})$, with

$$f(\lambda_\tau^{\mathbb{Q}}) = \frac{1}{g(\lambda_\tau^{Q})}\int_\tau^{\tau+M} E_\tau^{\mathbb{Q}}\left[\lambda_u^{\mathbb{Q}} e^{-\int_\tau^u (r_s+\lambda_s^{\mathbb{Q}})ds}\right] du, \tag{22}$$

the time-0 value of a payers option (with strike price K and time to expiration τ) on an M-year CDS is[21]

$$\text{Option}_0^{\text{payers}} = E_0^{\mathbb{Q}}\left[e^{-\int_0^\tau (r_s+\lambda_s^{\mathbb{Q}})ds}\left[L^{\mathbb{Q}} f\left(\lambda_\tau^{\mathbb{Q}}\right) - K\right]^+ g(\lambda_\tau^{\mathbb{Q}})\right]. \tag{23}$$

The time-τ payoff is discounted using the default-adjusted discount rate $r_t + \lambda_t^{\mathbb{Q}}$ to account for the knockout feature embedded in the single name contracts.[22] Similarly, a receivers option can be priced using the time-τ payoff of $\left[K - L^{\mathbb{Q}} f\left(\lambda_\tau^{\mathbb{Q}}\right)\right]^+$.

To explore the option pricing implications of our *ML* parameter estimates based on the *CDS* data we focus on the following three versions of the lognormal model: Case 1: $L^{\mathbb{Q}}$ constrained at 75% ; Case 2: $L^{\mathbb{Q}}$ is unconstrained; and Case 3: $L^{\mathbb{Q}}$ is unconstrained and $\mathbb{Q}$-stationarity ($\kappa^{\mathbb{Q}} > 0$) of the default intensity is enforced. For all three cases, Table 4 reports the model-implied three-month options on one-, five-, and ten-year *CDS*'s across a range of moneyness. The at-the-money (*ATM*) option on an M-year *CDS* is set to strike at the forward CDS price (the current price of a *CDS* to be started in three months with a maturity of M years). As expected, the *ATM* payers and sellers options are close in value because of our choice of the *ATM* strike price; moving across moneyness, the in-the-money (*ITM*) options are more expensive than the out-of-the-money (*OTM*) options; and increasing the maturity of the underlying *CDS* results in higher option values because the duration of the underlying contract increases with maturity.

The most interesting comparisons are between the three cases with different $(L^{\mathbb{Q}}, \kappa^{\mathbb{Q}})$. The model-implied option prices are significantly lower in Case 2 (with unconstrained $L^{\mathbb{Q}}$) than those in Case 1 (which constrains $L^{\mathbb{Q}}$ at 75%). This is largely a reflection of the much lower loss rates in the unconstrained models (e.g., for Mexico, $L^{\mathbb{Q}} = 21.4\%$ in Case 2). This sensitivity of option prices to $L^{\mathbb{Q}}$ illustrates the advantage of having reliable price information

[21] As mentioned in Section 3, we modify the fixed side g to account for accruals.

[22] Single name options cancel if there is a relevant credit event before the time to expiration of the option. Therefore, the payoff is realized only if the underlying issuer "survives" to the option expiration date.

Table 4: Model-Implied CDS Option Prices in Basis Points

Country	Case	Payers Options 20% ITM	Payers Options 10% ITM	Payers Options ATM	Payers Options 10% OTM	Payers Options 20% OTM	Receivers Options 20% OTM	Receivers Options 10% OTM	Receivers Options ATM	Receivers Options 10% ITM	Receivers Options 20% ITM
Three-month Options on One-Year CDS											
Mexico	1	28	24	21	18	15	12	17	22	27	33
	2	23	19	15	12	10	7	11	16	21	27
	3	29	25	22	19	16	12	17	22	27	33
Russia	1	65	50	37	28	20	14	24	38	54	72
	2	59	43	30	21	14	9	18	30	46	64
	3	69	54	42	33	25	18	29	43	59	77
Turkey	1	144	112	86	65	49	35	57	86	121	161
	2,3	121	89	64	45	31	22	40	65	95	131
Three-month Options on Five-Year CDS											
Mexico	1	298	238	188	147	115	84	129	185	250	323
	2	226	167	120	85	58	38	70	114	168	232
	3	292	231	180	139	107	75	119	174	239	311
Russia	1	481	363	268	193	137	100	173	268	383	518
	2	390	277	189	124	79	52	105	182	283	403
	3	481	357	257	182	127	89	161	258	379	520
Turkey	1	673	477	326	216	140	91	189	330	512	729
	2,3	388	257	159	93	52	29	78	162	277	417
Three-month Options on Ten-Year CDS											
Mexico	1	519	384	276	194	133	91	167	271	401	552
	2	325	220	141	86	51	33	77	146	240	352
	3	494	357	251	172	115	78	152	257	389	542
Russia	1	802	575	397	262	168	127	243	407	615	863
	2	566	383	243	145	82	58	132	248	407	600
	3	828	571	376	237	144	91	203	377	608	884
Turkey	1	949	631	394	234	134	74	197	401	682	1023
	2,3	432	275	161	88	46	21	71	165	298	464

Option prices are quoted in basis points and are priced under the lognormal model using the MLE estimates reported in Tables 1 and 2: Case 1 is with $L^{\mathbb{Q}}$ constrained at 75%; Cases 2 and 3 are with $L^{\mathbb{Q}}$ unconstrained, and, in addition, Case 3 enforces $\mathbb{Q}$-stationary of the default intensity $\lambda^{\mathbb{Q}}$.

for options for achieving separate identification of $L^{\mathbb{Q}}$ and $\lambda^{\mathbb{Q}}$. As illustrated in Duffie and Singleton [1999], the availability of prices of securities with nonlinear payoffs is necessary for identification when pricing is based on fractional recovery of market value. While, as we have shown above, identification is achieved even with linear instruments in a model with fractional recovery of face value (which is the relevant case for *CDS* contracts), the results in Table 4 show that identification will be strengthened by incorporation of options data.[23]

Given similar levels of the mean loss rate $L^{\mathbb{Q}}$, there are (at least) two potential effects on option prices of increasing the degree of persistence in $\lambda^{\mathbb{Q}}$. First, there is the direct affect on the discount factor for the option payoff. A more persistent $\lambda^{\mathbb{Q}}$, and in particular one that is explosive, will amplify the discounting of future cash flows. In addition, ceteris paribus, a reduction in $\kappa^{\mathbb{Q}}$ increases the volatility of $\lambda^{\mathbb{Q}}$ and, hence, also the volatility of the option payoff. We expect that the former, discount-rate effect, to be most important for long dated options, and the latter payoff-volatility effect to be more important the longer the maturity of the underlying.

The latter, payoff effect, is illustrated in Table 4. Taking the case of Mexico, for instance, the difference between the option prices for Cases 1 and 3 (explosive and stationary $\lambda^{\mathbb{Q}}$, respectively) are clearly largest when the ten-year *CDS* contract is the underlying. Moreover, even for three-month options, the differences are quite large, ranging between 20 and 30 basis points, or up to as much as 10% of the option value. Even larger differences are obtained for the *OTM* receiver options for Russia. However, in this case, there is more of a difference between the estimated values of $L^{\mathbb{Q}}$ (0.75 versus 1.0) across Cases 1 and 3.

Finally, it is also interesting to examine the model-implied option prices in terms of Black implied volatility. Following Black [1976], one could model the three-month forward-start *CDS* of maturity M as a lognormal process of its own and back out Black-implied volatility of such forward start *CDS* from the option prices reported in Table 4. As shown in Table 5, the lognormal model with constant volatility employed in this paper generates a flat implied volatility surface across moneyness. There is a downward sloping pattern in the term-structure of implied volatility, however. Not surprisingly, the pricing differences across the various cases reported in Table 4 are also reflected in the differing levels of Black implied volatility. For comparison, the realized volatilities of one-, five-, and ten-year *CDS*'s for Mexico, annualized, are 110%, 58%, and 49%, respectively, using the daily log-returns of the one-, five-, and ten-year Mexico *CDS*. As shown in Table 5, the implied volatilities for the forward-start *CDS*'s with the same maturities are of the same order of magnitudes. Of course, this is not surprising given that the model parameters are estimated to match the moments, including the second moment, implied in the term structure of *CDS* prices. For Russia, the corresponding annualized volatilities are 81%, 54%, and 50%, respectively; and for Turkey, they are 101%, 61%, and 57%.

[23] More precisely, we found that reliable identification is achieved when $\lambda^{\mathbb{Q}}$ is sufficient persistent and volatile. With options data, the range of parameters over which separation will be reliable will surely be increased.

Table 5: Black Implied Volatility of Model-Implied CDS Option Prices

Country	Case	Payers Options 20% ITM	Payers Options 10% ITM	Payers Options ATM	Payers Options 10% OTM	Payers Options 20% OTM	Receivers Options 20% OTM	Receivers Options 10% OTM	Receivers Options ATM	Receivers Options 10% ITM	Receivers Options 20% ITM
Three-month Options on One-Year CDS											
Mexico	1	126	126	126	126	126	131	131	131	131	131
	2	93	93	93	93	93	96	96	97	97	97
	3	131	131	131	131	130	133	133	133	133	133
Russia	1	74	74	74	74	74	74	74	75	75	75
	2	62	61	61	61	61	61	61	61	61	61
	3	83	83	83	83	83	83	84	84	85	85
Turkey	1	79	79	79	79	79	79	79	79	80	80
	2,3	67	66	66	65	65	67	67	66	66	66
Three-month Options on Five-Year CDS											
Mexico	1	93	92	91	90	89	90	90	89	88	88
	2	71	69	67	66	65	65	64	63	63	62
	3	90	88	87	86	85	85	85	84	83	83
Russia	1	73	72	71	70	70	72	72	71	71	70
	2	61	59	58	57	56	57	56	56	55	55
	3	68	67	66	66	66	67	67	66	66	67
Turkey	1	57	57	56	56	57	57	57	57	57	58
	2,3	45	44	44	45	45	44	45	45	46	47
Three-month Options on Ten-Year CDS											
Mexico	1	68	67	66	65	65	66	65	65	64	64
	2	49	48	48	48	48	50	50	50	51	51
	3	60	60	60	60	60	61	61	62	62	62
Russia	1	60	59	59	58	57	61	61	60	60	59
	2	49	49	48	47	47	50	49	49	49	49
	3	52	52	51	51	51	52	52	52	52	52
Turkey	1	45	45	45	46	46	45	46	46	47	48
	2,3	38	39	39	40	41	38	39	40	41	43

The Black implied volatilities are quoted on a percentage basis.

10 Is One Factor Enough?

Up to this point we have chosen to focus on a single-factor model for $\lambda^{\mathbb{Q}}$, largely because, for a given sovereign, the first *PC* of the *CDS* spreads explains a very large percentage of the variation for all maturities. In concluding our analysis, we expand on the empirical plausibility of the one-factor assumption and discuss some reasons why an additional factor might be useful for understanding spreads, particularly at the short end of the maturity spectrum.

Table 6 displays the factor loadings and the percentage variation explained from projections of the *CDS* spreads onto the first two *PC*s of the data.[24] The cumulative R^2, for the first two principal components (*PC1* and *PC2*) combined, is obtained by summing the R^2s across each of the two *PC*s. As noted at the outset of our analysis, *PC1* explains a large percentage of the variation in spreads for all countries and all maturities. Indeed, for maturities of three years and longer, *PC1* accounts for at least 99% of the variation in all of the spreads. Moreover, parallel to the findings for the term-structures of the *US* treasury or swap markets (Litterman and Scheinkman [1991]), the first *PC* emerges as a "level" factor, as reflected in the roughly constant factor loadings across maturities (for a given sovereign). The first part of Table 7 displays the R^2's from the regressions of the time series of model-implied $\lambda^{\mathbb{Q}}$ onto *PC1*. We see that our one-factor model is largely picking up this first *PC*, as expected.

Table 6: Principal Component Analysis of CDS Spreads: Mexico, Russia and Turkey

	Mexico				Russia				Turkey			
	PC1		PC2		PC1		PC2		PC1		PC2	
	$\hat{\beta}$	R^2	$\hat{\beta}$	R^2	$\hat{\beta}$	R^2	$\hat{\beta}$	R^2	$\hat{\beta}$	R^2	$\hat{\beta}$	R^2
OLS regression of CDS spreads on the first two principal components												
1y	0.27	84.8%	0.82	14.3%	0.27	88.6%	0.84	11.2%	0.58	96.9%	0.79	3.0%
3y	0.54	98.8%	0.31	0.6%	0.49	99.6%	0.18	0.2%	0.51	99.6%	-0.17	0.2%
5y	0.58	99.3%	-0.32	0.5%	0.55	99.9%	-0.01	0.0%	0.46	98.7%	-0.39	1.2%
10y	0.55	98.7%	-0.37	0.8%	0.62	99.0%	-0.51	0.9%	0.43	97.8%	-0.45	1.8%
OLS regression of AR(1) residuals of *CDS* spreads on the first two principal components												
1y	0.26	49.2%	0.82	39.4%	0.37	53.8%	0.85	43.4%	0.63	91.1%	0.72	8.4%
3y	0.53	86.7%	0.30	2.3%	0.48	85.8%	0.12	0.9%	0.52	93.8%	-0.18	0.8%
5y	0.57	94.2%	-0.20	0.9%	0.54	92.8%	-0.19	1.8%	0.44	92.9%	-0.35	4.3%
10y	0.57	89.6%	-0.45	4.8%	0.58	86.4%	-0.47	8.8%	0.37	79.2%	-0.56	12.8%

At the same time, Table 6 reveals that *PC2* plays a prominent role in explaining the

[24] This *PC* analysis was conducted using the covariance matrix of the levels of spreads. The results for the *AR*(1) residuals refer to a *PC* analysis of the residuals from regressing the spread for each country and maturity onto a constant and its own lagged value.

movements in the short end of the *CDS* curves, particularly the one-year spreads for Mexico and Russia. Figure 4 revealed some systematic mispricing of the one-year *CDS* contract for Mexico in all three models. So a natural question at this juncture is whether the pricing errors for our models are correlated with *PC2*. As shown in the bottom portion of Table 7 they are in fact strongly correlated, especially for Mexico and Russia. The very strong correlation is displayed graphically in Figure 11.

	MEX		RUS		TUR	
OLS regression of $\lambda^{\mathbb{Q}}$ on PC 1						
	β	R^2	β	R^2	β	R^2
	0.23	98.2%	0.35	99.8%	0.50	98.9%
OLS regression of CDS pricing errors on PC 2						
	β	R^2	β	R^2	β	R^2
1yr	-1.00	91.4%	-0.83	71.0%	-1.01	47.2%
3yr	-0.62	66.5%	-0.18	28.2%	-0.19	33.1%
10yr	0.12	4.2%	0.50	24.3%	0.11	5.6%

Table 7: $\lambda^{\mathbb{Q}}$, Pricing Errors, and PC Factors

Whereas for treasury markets the second *PC* and second factor in a dynamic term structure model typically behave like a "slope" factor, that is not the case for these sovereign *CDS* markets. It is primarily the one-year contract (for Mexico and Russia) that is mispriced by a one-factor model. Moreover, inspection of the loadings associated with *PC2* (see Table 7) reveals a kink between the one- and three-year maturities. That is, all of this evidence points to the one-year contract having special behavior relative to the contracts with longer maturities. Based on conversations with traders, it seems that the most likely explanation for this "anomalous" behavior of the one-year contract is due to a liquidity or supply/demand premium. We are told that large institutional money management firms often use the short-dated *CDS* contract as a primary trading vehicle for expressing views on sovereign bonds. The sizable trades involved in these transactions introduce an idiosyncratic "liquidity" factor into the behavior of the one-year contract. The role of such trading pressures on *CDS* spreads is something we hope to explore in future research.

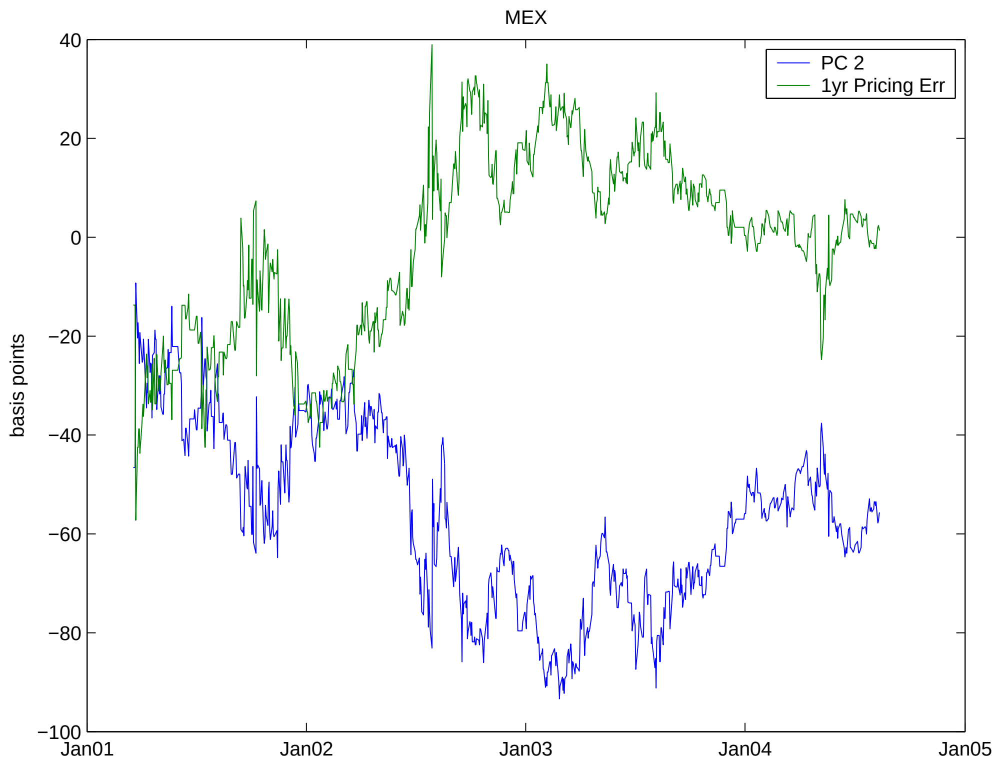

Figure 11: 1 Year Pricing Errors and *PC2*: Mexico

References

Ahn, D. and B. Gao (1999). A Parametric Nonlinear Model of Term Structure Dynamics. *Review of Financial Studies 12*, 721–762.

Altman, E., B. Brady, A. Resti, and A. Sironi (2003). The Link between Default and Recovery Rates: Theory, Empirical Evidence and Implications. Working Paper, Stern School of Business New York University.

Berndt, A., R. Douglas, D. Duffie, M. Ferguson, and D. Schranzk (2004). Measuring Default Risk Premia from Default Swap Rates and EDFs. Working Paper, Stanford University.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics 3*, 167–179.

Cheridito, R., D. Filipovic, and R. Kimmel (2003). Market Price of Risk in Affine Models: Theory and Evidence. Working Paper, Princeton University.

Cifarelli, G. and G. Paladino (2002). An Empirical Analysis of the Co-Movement Among Spreads on Emerging-Market Debt. Working Paper, University of Florence.

Cogan, P. (2005). If They all Rush For the Exit at the Same Time... *Financial Times May*

27, 1–2.

Collin-Dufresne, P., R. Goldstein, and C. Jones (2004). Can Interest Rate Volatility Be Extracted From the Cross Section of Bond Yields? An Investigation of Unspanned Stochastic Volatility. Working Paper, NBER 10756.

Collin-Dufresne, P., R. S. Goldstein, and J. Martin (2001). The Determinants of Credit Spread Changes. *Journal of Finance 56*, 2177–2208.

Cox, J., J. Ingersoll, and S. Ross (1985). A Theory of the Term Structure of Interest Rates. *Econometrica 53*, 385–407.

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

Dai, Q. and K. Singleton (2002). Expectations Puzzles, Time-Varying Risk Premia, and Affine Models of the Term Structure. *Journal of Financial Economics 63*, 415–441.

Duffee, G. R. (2002). Term Premia and Interest Rate Forecasts in Affine Models. *Journal of Finance 57*, 405–443.

Duffie, D. (1998). Defaultable Term Structure Models with Fractional Recovery of Par. Working Paper, Graduate School of Business, Stanford University.

Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica 68*, 1343–1376.

Duffie, D., L. Pedersen, and K. Singleton (2003). Modeling Credit Spreads on Sovereign Debt: A Case Study of Russian Bonds. *Journal of Finance 55*, 119–159.

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

Duffie, D. and K. Singleton (1999). Modeling Term Structures of Defaultable Bonds. *Review of Financial Studies 12*, 687–720.

Duffie, D. and K. Singleton (2003). *Credit Risk*. Princeton: Princeton University Press.

Economist (2004). Argentinas Default, Shooting the Messenger. September 25 Issue, page 91.

Henry, D. (2004). Ahead of a Fed Move, a Cash Cow Runs Dry. *Business Week May 24*, 1–2.

Houweling, P. and T. Vorst (2005). Pricing Default Swaps: Empirical Evidence. forthcoming, *Journal of International Money and Finance*.

Hull, J. and A. White (2004). The Relationship Between Credit Default Swap Spreads, Bond Yields, and Credit Rating Announcements. Working Paper, University of Toronto.

Jarrow, R. A., D. Lando, and F. Yu (2005). Default Risk and Diversification: Theory and Empirical Implications. *Mathematical Finance 15*, 1 – 26.

Kaminsky, G. and S. Schmukler (2001). Emerging Markets Instability: Do Sovereign Ratings Affect Country Risk and Stock Returns? Working Paper, George Washington University.

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

Litterman, R. and J. Scheinkman (1991). Common Factors Affecting Bond Returns. *Journal of Fixed Income 1*, 54–61.

Longstaff, F., S. Mithal, and E. Neis (2005). Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit-Default Swap Market. *Journal of Finance 60*, 2213–2253.

Moody's (2003). Sovereign Bond Defaults, Rating Transitions, and Recoveries (1985-2002). Working Paper, Moody's Investors Service.

Packer, F. and C. Suthiphongchai (2003). Sovereign Credit Default Swaps. *BIS Quarterly Review December*, 81–88.

Rigobon, R. (2001). Contagion: How to Measure It? Working Paper, NBER Working Paper No. 8118.

Rigobon, R. (2002). The Curse of Non-investment Grade Countries. *Journal of Development Economics 69*, 423–449.

Schaefer, S. and I. Strebulaev (2004). Structural Models of Credit Risk are Useful: Evidence from Hedge Ratios on Corporate Bonds. Working Paper, Graduate School of Business, Stanford University.

Singh, M. (2003). Recovery Rates from Distressed Debt– Empirical Evidence from Chapter 11 Filings, International Litigation, and Recent Sovereign Debt Restructurings. Working Paper, IMF 03/161.

Xu, D. and C. Wilder (2003). Emerging Market Credit Derivatives. Working Paper, Deutsche Bank.

Yafeh, Y., N. Sussman, and P. Mauro (2001). Emerging Market Spreads: Then Versus Now. Working Paper, Hebrew University of Jerusalem.

Yu, F. (2002). Modeling Expected Return on Defaultable Bonds. *Journal of Fixed Income 12*, 69–81.

Zhang, F. (2003). What Did the Credit Market Expect of Argentina Default? Evidence from Default Swap Data. Working Paper, Federal Reserve Board.