

CAT bond spreads: A model with HARA utility and an empirical validation using nonparametric tests

Abstract

Previous empirical studies on catastrophe (CAT) bond premium calculations rely almost exclusively on actuarial models, and usually compare their accuracy strictly in terms of in-sample fit and predictive power. We contribute to this literature by deriving a utility-based specification for pricing CAT bonds under hyperbolic absolute risk aversion (HARA), and by proposing two specification tests that use nonparametric estimation techniques to test simultaneously for all possible mis-specifications. Various pricing models are then estimated and tested with data from the primary market for CAT bonds. Our results suggest that the utility-based model we propose not only is well-suited for explaining the risk-return relationship observed in the CAT bond market but also delivers the best performance among the tested models. We also provide new empirical evidence that the aggregate utility function of CAT investors exhibits decreasing absolute risk aversion.

Keywords: Pricing CAT bonds, HARA, Statistical analyses of CAT bond spreads, Nonparametric specification tests, Econometric pricing models, Out-of-sample analysis.

1 Introduction

Increases in the frequency and severity of natural catastrophes in recent years have propelled the use of alternative risk-transfer instruments for managing catastrophic risks. As a matter of fact, the size of the outstanding catastrophe bond and insurance-linked securities market at the end of 2015 reached an historical record of \$25.903 billion.¹ Catastrophe bonds (CAT bonds) are insurance-linked securities offering an exposure to natural disaster risk in the form of a bond with coupon payments that are contingent upon the absence of a pre-agreed triggering event. In other words, if a natural catastrophe of sufficient magnitude occurs within the life of the CAT bond, which is typically between one and four years, the holder loses some fraction of the par amount and of the remaining coupons. If there is no catastrophic triggering event during the term of the bond, the principal is returned to the investor with a final coupon payment.

Most of the empirical works on CAT bonds seek premium calculation models that best describe the empirical data. Lane (2000) argues that a trade-off between the frequency and the severity of losses must be included in the pricing framework, and proposes a CAT bond pricing model that consists of a Cobb-Douglas production function on the conditional expected loss and the probability of first loss. Major and Kreps (2002) hypothesize a loglinear relationship between the spread and the expected loss, but they apply their model to traditional treaties instead of CAT bonds. Wang (2004) uses distortion operators to transform the probability of loss into an empirical value that can be used for pricing. To explain CAT bonds' spread premiums, Lei et al. (2008) use a multivariate linear regression model that involves the conditional expected loss, the probability of first loss, the probability of exhaustion, and other CAT-bond-specific information such as maturity, issue size, location, trigger type, and rating. Dieckmann (2010) estimates other multivariate linear regressions to establish some stylized facts such as the effect that Hurricane Katrina had on the pricing of CAT bonds. Lane and Mahul (2008) propose a simple linear relation between the spread and the expected loss, and also extend their model to include an additional variable that accounts for cycle effects. Papachristou (2011) criticizes

¹See article "Q4 2015 Catastrophe Bond & ILS Market Report" at <http://www.artemis.bm>.

the use of a simple linear relation between the spread and the expected loss and instead uses a generalized additive model to establish the main factors driving the spreads of CAT bonds. Bodoff and Gan (2012) describe the market clearing issuance price of CAT bonds as a linear function of the expected loss. They argue that the intercept of the linear model is due to the required rate of return on downside risk capital, and that the slope coefficient is due to the uncertainty of the estimate of the expected loss. Galeotti et al. (2013) apply different premium calculation models and compare them with regard to their in-sample accuracy and predictive power. They find that a version of the Wang (2004) transformation model and the linear model are the most accurate. They also find that CAT-bond-specific information does not improve the out-of-sample pricing accuracy of the tested models. Gürtler et al. (2014) present evidence that both the 2007-2008 financial crisis and 2005's Hurricane Katrina significantly affected CAT bond premiums. Recently, Braun (2015) has estimated a series of ordinary least squares regressions with heteroskedasticity- and autocorrelation-consistent standard errors, and presents results confirming that the expected loss is the main driver of the spreads. The territory covered, the sponsor, the reinsurance cycle, and the spreads on comparably rated corporate bonds also exhibit an impact on the pricing relation. In this work, we investigate whether risk aversion is well suited for explaining the risk-return relationship observed in the market for CAT bonds. Since macroeconomic and CAT-bond-specific factors affect the spread to a lesser extent, as documented in previous literature, we leave these out of our study. We instead focus on the relation between the spread and the expected loss, which is by far the most important driver, economically.²

²Galeotti et al. (2013) find that including additional CAT-bond-specific factors (such as the type of peril and the trigger mechanism) and macroeconomic factors (cyclical, seasonal, and business-cycle effects) does not improve out-of-sample results. In fact, Table 1 of Galeotti et al. (2013) shows that, in a stable market environment (i.e., when the effects of the financial crisis of 2008 are excluded), the results of the out-of-sample analysis can be even worse for these extended models. Braun (2015) also confirms the expected loss as the most important spread driver, but finds that the territory covered, the sponsor, the reinsurance cycle, and the spreads on comparably rated corporate bonds also exhibit an impact on the spread. More precisely, Braun (2015) finds that the spread increases by 221.04 bp per percentage point of expected loss. In comparison, the spread widens by 26.57 bp for each percentage point of the BB corporate bond spread and increases by 161.85 bp per point of the rate on line index. Moreover, the spread rises by 175.08 bp for peak territories, declines by 103.58 bp for Swiss Re sponsored deals, and declines by 159.76 bp if the CAT bond exhibits an investment-grade rating. Braun (2015)'s extended model yields a centered adjusted R^2 of 0.89, whereas a simple linear model in the expected loss displays a centered adjusted R^2 of 0.80. Thus, both studies mentioned above indicate that the expected loss is by far the most important CAT bond spread driver. Braun (2015) also states that (unreported) collinearity diagnostics indicate that the regressors are only minimally correlated with each other. Hence, including these additional

Although several actuarial and pure econometric models have been investigated empirically, there are various pricing approaches that have been put forward in the theoretical literature on insurance-linked securities that have not yet been tested empirically in the context of CAT bonds. Popular proposed frameworks are based on arbitrage pricing (Vaugirard, 2003), equilibrium pricing (see, e.g., Cox and Pedersen (2000) and Chang et al. (2010)), stochastic dominance (Perrakis and Bolorforoosh, 2013), ambiguity aversion (Zhu, 2011), utility indifference pricing (Egami and Young, 2008) and even on the reduced-form approach used to price credit derivatives (Jarrow, 2010). However, empirical applications of these theoretical models are currently lacking in the CAT bond literature. In the first section of this paper, we fill this gap by deriving a utility-based specification for pricing CAT bonds under hyperbolic absolute risk aversion (HARA). This model, which we refer to as the HARA model, is well-grounded in equilibrium pricing theory and can be tested readily with CAT bond data.

As stated earlier, it is generally agreed that the expected loss is the most important spread driver, pointing to the importance of correctly specifying the functional relation between the spread and the expected loss. However, most empirical studies compare the specifications mainly in terms of in-sample fit and predictive power. Although this approach may indicate which of these models is most appropriate, it does not provide evidence of the validity of the selected model. In other words, one model may outperform all the others tested, yet not appropriately describe the actual mechanism that generates the data. We address this issue in the second part of this paper, proposing two specification tests that use nonparametric regressions to test simultaneously for all possible departures from the model of interest. The first test is based upon the idea that a well-specified parametric regression model should not predict significantly less well than a nonparametric one, which will in general closely resemble the true model if the sample size is sufficiently large. The second test is based on the idea that residuals from a well-specified parametric regression should be centered at zero over all values of the independent variable, implying that a nonparametric regression of the residuals on the independent variable

factors is not necessary here, since our focus is the relation between the spread and the expected loss. Moreover, a multi-dimensional extension of the nonparametric approach proposed in this work would be harder to implement in practice.

should be identically zero everywhere. To the best of our knowledge, this is the first time that nonparametric regressions have been used to test CAT bonds' pricing specifications.

In the last section of this work, we investigate the adequacy of our proposed model (the HARA model) by using both the traditional approach and the nonparametric approach we advocate. The empirical data set we use consists of 371 CAT bond tranches issued between June 1997 and March 2015. We also test various premium calculation models that have been proposed in the previous literature. In particular, the linear specification is rejected at the significance level of 1%. In fact, we find that, among the tested models, the HARA specification is the only one not rejected at the significance level of 5%. We also present evidence suggesting that the aggregate utility function of CAT investors exhibits decreasing absolute risk aversion.

Our paper continues as follows. We derive our CAT bond pricing formula with HARA utility in Section 2. We describe the mechanics of the nonparametric specification tests in Section 3. We then perform empirical analyses of CAT bond spreads in Section 4. We summarize and conclude our paper in Section 5.

2 Pricing CAT bonds with HARA utility: Theoretical framework

In this section, we first present a brief overview of the typical CAT bond structure and we then discuss popular approaches for investigating how these transactions can be best priced. More importantly, we derive a utility-based econometric model for pricing CAT bonds under hyperbolic absolute risk aversion, which is in contrast to the actuarially oriented specifications that have been proposed in the extant empirical literature.

2.1 CAT bonds

CAT bonds are insurance-linked securities invented by insurers and reinsurers to shift natural disaster risks to the capital markets. These transactions usually involve a Special-Purpose Vehicle (SPV), located in a tax-efficient jurisdiction, that sells catastrophe protection to a ceding (re)insurer in the form of a reinsurance contract. The SPV then effectively transfers its risk exposure by issuing CAT bond tranches to capital market investors. In order to offer a virtually

pure exposure to the natural disaster risk, the proceeds of the issuance are invested by the SPV in highly rated short-term assets that are held in a collateral account.³

In general, the payout to the ceding (re)insurer is contingent upon a specific underlying loss variable L breaching a pre-agreed attachment point a before maturity,⁴ in which case the collateral is liquidated to reimburse the sponsor up to the par amount P paid by the investor at the issue date. In practice, both excess-of-loss (XOL) and binary (Bin) payouts are possible (see, e.g., Cummins and Trainar (2009)). The corresponding payout functions are respectively given by

$$L_{a,P}^{\text{XOL}} = \begin{cases} 0, & L \leq a, \\ L - a, & a < L \leq a + P, \\ P, & L > a + P, \end{cases} \quad L_{a,P}^{\text{Bin}} = \begin{cases} 0, & L \leq a, \\ P, & L > a, \end{cases} \quad (1)$$

where P is the principal of the bond, a is the attachment point and L is the underlying loss variable. If there is no triggering event during the term of the CAT bond, which is typically between one and four years, the principal P is returned to the investor with a coupon payment of the form $(r + S)P$, where r is a risk-free rate (e.g., LIBOR) and S is the spread risk premium.

2.2 Background on pricing CAT bonds

To investigate the question of how CAT bond transactions can best be priced, as described in the introduction, previous research usually follows either an econometric approach or a theoretical/numerical approach.

Empirically, the expected loss of the CAT bond, defined as $\text{EL} \equiv \mathbb{E}[L_{a,P}]/P$, is found to be the most economically important factor for predicting the spread premium.⁵ Several empirical studies

³Prior to the financial crisis of 2007, the typical CAT bond structure also used a total return swap (TRS) to protect the collateral account. This structure was, however, abandoned in late 2008 as the importance of the TRS counterparty risk was brought to light by the collapse of Lehman Brothers, which was the TRS counterparty of four CAT bonds at the time of its bankruptcy (see, e.g., Tower Watson (2010) for more details).

⁴See, e.g., Mocklow et al. (2002) for an overview of the different trigger mechanisms.

⁵See Galeotti et al. (2013), who find that other CAT-bond-specific information does not improve the out-of-sample pricing accuracy of econometric models. Macroeconomic factors that have been found to be statistically significant explanatory variables are the reinsurance cycle and the spreads on comparably rated corporate bonds. These variables are, nevertheless, economically much less important than the expected loss variable (Braun, 2015, p. 22).

have therefore attempted to identify the most suitable parametric relation between the spread S and the expected loss EL. In general, these econometric models can be written as

$$S = f(\text{EL}; \theta) + \epsilon, \quad \text{EL} = \frac{\mathbb{E}[L_{a,P}]}{P}, \quad (2)$$

where $f(\cdot; \theta)$ is a parametric function with (unknown) parameters θ , EL is the expected loss of the CAT bond, and ϵ is a pricing error term.⁶ By comparing both the in-sample and the out-of-sample pricing accuracy of several such econometric models, earlier research attests that a linear relationship between the spread and the expected loss is appropriate (see, e.g., Bodoff and Gan (2012), Galeotti et al. (2013) and Braun (2015)).

In contrast with econometric approaches, theoretical/numerical pricing methods are well-developed in the literature and propose a variety of sophisticated CAT bond pricing models. As discussed in the introduction, the arbitrage approach is arguably the most popular one (see, e.g., Lee and Yu (2002), Vaugirard (2003), Nowak and Romaniuk (2013), Ma and Ma (2013) and Lai et al. (2014)), with other notable frameworks being based on equilibrium pricing (see, e.g., Cox and Pedersen (2000) and Chang et al. (2010)), stochastic dominance (Perrakis and Bolorforoosh, 2013), ambiguity aversion (Zhu, 2011), utility indifference pricing (Egami and Young, 2008) and even on the reduced-form credit derivatives pricing approach (see, e.g., Jarrow (2010) and Têtu et al. (2015)). Most of these theoretical/numerical methods are, however, difficult to translate into pure econometric specifications that can be tested empirically with the data that are publicly available, which might explain why previous econometric models do not rely on this literature and are instead actuarially oriented (see, e.g., Galeotti et al. (2013)).

In the next section, we attempt to fill the gap between the econometric and the theoretical approaches for pricing CAT bonds by presenting a new econometric specification that is well-grounded in financial theory in addition to being straightforward to test empirically.

⁶A few other econometric models do not consider the expected loss as the main explanatory variable. Lane (2000) models the absolute premium $S - \text{EL}$ as a Cobb-Douglas production function on the conditional expected loss $\text{CEL} = \mathbb{E}[L_{a,P} | L > a] / P$ and the probability of first loss $\text{PFL} = \text{Prob}(L_{a,P} > 0)$. Another approach, proposed by Wang (2000, 2004), uses distortion operators to transform the probability of loss into an empirical one.

2.3 Derivation of the HARA specification

We start from equilibrium pricing theory and make use of a few modeling assumptions and mathematical approximations to derive a utility-based econometric model for pricing CAT bonds under hyperbolic absolute risk aversion. We consider a single-period economy in which the uncertainty is represented by the set Ω of possible states of the world ω at the terminal date.

Let us price the following CAT bond transaction. At the initial date $t = 0$, the par amount P is paid by the investor. At the terminal date $t = 1$, the investor receives a payment of the form $(1 + r + S)P$, where r is the risk-free rate and S is the spread risk premium. At the same time, the investor suffers a loss $L_{a,P} \geq 0$ (see, e.g., Equation (1)). To obtain the fair spread premium S , we simply equate the net present value of the transaction to zero and solve for S :

$$\frac{(1 + r + S)P - V_1(L_{a,P})}{1 + r} - P = 0 \quad \Rightarrow \quad S = \frac{V_1(L_{a,P})}{P}, \quad (3)$$

where $V_1(L_{a,P})$ denotes the price at time $t = 1$ of the (random) payoff $L_{a,P}$.

To simplify valuation, we suppose that the competitive equilibrium admits a representative agent who owns the endowments and whose problem is to maximize the aggregate utility of the agents (i.e., the insurer and the investor).⁷ It can be shown that the first-order condition of this problem produces a pricing kernel (i.e., state price per unit of probability) that depends only on the aggregate utility and wealth of the agents (see, e.g., Chapter 7 of Back (2010)). If we denote the aggregate utility functions at time $t = 0$ and time $t = 1$ by $U_0(\cdot)$ and $U_1(\cdot)$ respectively, the pricing kernel is then given by

$$\xi(W_\omega) \equiv \frac{d\phi(\omega)}{f(\omega)d\omega} = \frac{U'_1(W_\omega)}{U'_0(W_0)}, \quad \forall \omega \in \Omega, \quad (4)$$

where $d\phi(\omega)$ is the Arrow-Debreu price of state ω , $f(\omega)$ is the probability density of state ω , W_0 is the aggregate wealth at time $t = 0$ and W_ω is the aggregate wealth at time $t = 1$ in the state ω .

⁷The construction of a representative agent is frequently done in finance to simplify valuation (see, e.g., Rubinstein (1974), Breeden and Litzenberger (1978) and Constantinides (1982)) and can be shown to produce the same equilibrium as in a competitive Arrow and Debreu (1954) economy.

For simplicity, let us consider a CAT bond with binary payout $L_{a,P} = \mathbf{1}_{\{L_\omega > a\}}P$,⁸ where $\mathbf{1}_{\{\cdot\}}$ is the indicator function, L_ω is the value of the underlying loss variable in the state ω and a is the attachment point of the CAT bond. Using the pricing kernel (4) for valuation, the spread premium of this CAT bond is thus given by

$$S = \int_{\omega \in \Omega_a} \frac{d\phi(\omega)}{C} = \int_{\omega \in \Omega_a} \frac{\xi(W_\omega)}{C} f(\omega) d\omega, \quad \Omega_a \equiv \{\omega \in \Omega \mid L_\omega > a\}, \quad (5)$$

where the constant $C = \int_{\omega \in \Omega} d\phi(\omega)$ ensures that $S = 1$ if $\Omega_a = \Omega$ (as dictated by Equation (3)).

We assume that the aggregate wealth at time $t = 1$ is of the form $W_\omega = W - L_\omega$, where L_ω is the value of the loss variable in state ω , and W is a constant (or a certainty equivalent) that represents the total wealth of the insurer and the investor in the absence of natural disasters. The previous equation can then be written as

$$S = \int_{\omega \in \Omega_a} \frac{\xi(W - L_\omega)}{C} f(\omega) d\omega \approx \frac{\xi(W - a)}{C} \pi_a, \quad \pi_a \equiv \int_{\omega \in \Omega_a} f(\omega) d\omega, \quad (6)$$

where the right-hand side follows from a first-order approximation of the integral,⁹ and where π_a is the CAT bond's probability of being triggered, i.e., the probability that $L_\omega > a$. The role of the attachment point in the valuation formula (6) is twofold. First, a higher value of the attachment point a implies a lower trigger probability π_a , which reduces the spread premium. Second, a higher value of a means that the asset hedges against states of the world in which aggregate wealth is lower, which increases the spread by increasing the value $\xi(W - a)$ of the pricing kernel. The next step is to capture the dynamic between the CAT bond's spread S and its trigger probability π_a . To do so, we first use Equation (4) to obtain the first derivative of the kernel:

$$\xi'(W_\omega) = \frac{U_1''(W_\omega)}{U_0'(W_0)} = -A_1(W_\omega) \frac{U_1'(W_\omega)}{U_0'(W_0)} = -A_1(W_\omega) \xi(W_\omega),$$

⁸This assumption seems reasonable given that the CAT bonds in our empirical data set (see Section 4.1) have an average conditional expected loss of about 77%.

⁹We found that higher-order approximations resulted in much more complicated derivations. A first-order approximation is preferred here since our goal is to provide a closed-form pricing formula that can be evaluated without the use of numerical methods. It could be interesting to include higher orders in future work.

where

$$A_1(W_\omega) \equiv -\frac{U_1''(W_\omega)}{U_1'(W_\omega)}$$

is the aggregate absolute risk aversion. We can then use this identity together with the valuation formula (6) to obtain a differential equation:

$$\begin{aligned} C \frac{dS}{d\pi_a} &= \frac{d}{d\pi_a} \xi(W-a) \pi_a, \\ &= \xi(W-a) - \xi'(W-a) \pi_a \frac{da}{d\pi_a}, \\ &= \xi(W-a) + A_1(W-a) \xi(W-a) \pi_a \frac{da}{d\pi_a}, \\ &= \xi(W-a) \left(1 + A_1(W-a) \pi_a \frac{da}{d\pi_a} \right), \\ \Rightarrow \frac{dS}{d\pi_a} &= \frac{S}{\pi_a} \left(1 + A_1(W-a) \pi_a \frac{da}{d\pi_a} \right). \end{aligned}$$

Rearranging the terms yields

$$\frac{dS}{S} = \frac{d\pi_a}{\pi_a} \left(1 + A_1(W-a) \pi_a \frac{da}{d\pi_a} \right). \quad (7)$$

To solve this differential equation, one needs to make a parametric assumption regarding the distribution of the underlying loss variable L_ω . Popular models for calculating premiums of reinsurance contracts include the exponential and Pareto distributions (see, e.g., Burnecki et al. (2011)). Here, the exponential distribution is a convenient first approximation since it provides a straightforward solution to the differential equation above. In fact, under the assumption that the loss variable L_ω follows an exponential distribution with rate parameter λ , we first note that

$$\pi_a \equiv \text{Prob}(L_\omega > a) = e^{-a\lambda} \quad \Rightarrow \quad a = -\frac{1}{\lambda} \ln \pi_a \quad \Rightarrow \quad \frac{da}{d\pi_a} = -\frac{1}{\pi_a \lambda}. \quad (8)$$

Using these formulas in Equation (7) and simplifying yields

$$\frac{dS}{S} = \frac{d\pi_a}{\pi_a} \left(1 - \frac{A_1(W + \lambda^{-1} \ln \pi_a)}{\lambda} \right). \quad (9)$$

Solving this differential equation leads to an expression for the equilibrium spread S of the CAT bond as a function of its trigger probability π_a . We can guess that the solution is going to depend on the form of the absolute risk aversion function $A_1(\cdot)$. In the case of risk neutrality (i.e., $A_1 = 0$) it can be shown that the solution has the linear form: $S = \pi_a$. As we show below, a nonlinear relation between S and π_a indicates the presence of risk aversion.

We consider the hyperbolic absolute risk aversion (HARA) form (see, e.g., Mossin (1968)), which is a rich family of utility functions that is especially convenient for mathematical modeling (see, e.g., Merton (1971)). In the following, we mainly make use of the fact that all utility functions in the HARA family have an absolute risk aversion function of the form

$$A_1(W_\omega) = \frac{1}{\delta + \gamma W_\omega}. \quad (10)$$

The HARA family is rich in the sense that, by adjusting the coefficient γ , one can obtain a utility function with an absolute risk aversion that is constant ($\gamma = 0$), increasing ($\gamma < 0$) or decreasing ($\gamma > 0$). Moreover, by adjusting the coefficient δ , one can obtain a utility function with a relative risk aversion that is constant ($\delta = 0$), increasing ($\delta > 0$) or decreasing ($\delta < 0$).¹⁰

To obtain the pricing dynamic under HARA utility, we rewrite Equation (9) using Equation (10):

$$\frac{dS}{S} = \frac{d\pi_a}{\pi_a} \left(1 - \frac{1}{\lambda(\delta + \gamma(W + \lambda^{-1} \ln \pi_a))} \right) = \frac{d\pi_a}{\pi_a} \left(1 - \frac{1}{\lambda(\delta + \gamma W) + \gamma \ln \pi_a} \right). \quad (11)$$

We then introduce the constant $\alpha \equiv \lambda(\delta + \gamma W)$ to simplify the notation. The equation above becomes

$$\frac{dS}{S} = \frac{d\pi_a}{\pi_a} - \frac{1}{\alpha} \frac{d\pi_a}{\pi_a (1 + \gamma/\alpha \ln \pi_a)}. \quad (12)$$

¹⁰Relative risk aversion is defined as $R_1(W_\omega) = W_\omega A_1(W_\omega)$.

We then integrate both sides using the following identities (see, e.g., Lipschutz et al. (2012)):

$$\int \frac{d\pi_a}{\pi_a} = \ln \pi_a + \text{constant of integration},$$

$$\int \frac{d\pi_a}{\pi_a(1 + \gamma/\alpha \ln \pi_a)} = \frac{\alpha}{\gamma} \ln \left(1 + \frac{\gamma}{\alpha} \ln \pi_a \right) + \text{constant of integration},$$

and we obtain

$$\ln S = \ln \pi_a - \frac{1}{\gamma} \ln \left(1 + \frac{\gamma}{\alpha} \ln \pi_a \right) + \ln S_1,$$

where $\ln S_1$ denotes the constant of integration, which is equal to the natural logarithm of the spread premium at $\pi_a = 1$. Using $S_1 = 1$ as boundary condition,¹¹ and taking the exponential of both sides of this equation yields the following expression for the spread of the CAT bond:

$$S = \pi_a \left(1 + \frac{\gamma}{\alpha} \ln \pi_a \right)^{-1/\gamma}. \quad (13)$$

The parameter γ represents the slope of the aggregate risk tolerance function $T_1(W_\omega) \equiv A_1(W_\omega)^{-1}$ (see Equation (10)), and the parameter $\alpha \equiv (\delta + \gamma W)/\lambda^{-1}$ is the ratio of the aggregate risk tolerance in the absence of a loss (i.e., if $L_\omega = 0$) to the expected value of the underlying loss variable (i.e., $\mathbb{E}[L_\omega] = \lambda^{-1}$). Note that when the trigger probability is $\pi_a = 0$, solving the limit in Equation (13) yields $S = 0$ as it should be.¹²

If we denote by EL the expected loss of the CAT bond (instead of π_a), Equation (13) then suggests the following regression model (hereafter referred to as the HARA model):

$$S = \text{EL} \left(1 + \frac{\gamma}{\alpha} \ln \text{EL} \right)^{-\frac{1}{\gamma}} + \epsilon, \quad (14)$$

where ϵ is an error term. This regression model assumes that the parameters γ and α are the same for all CAT bond transactions in the data sample. This is unlikely to be true since these parameters depend on the expected economic loss of the underlying peril, as well as on the

¹¹In Equation (5), we obviously have $V_1(L_{a,p}) = P$ when the trigger probability is $\pi_a = 1$, which implies $S = 1$.

¹²It can be shown that Equation (13) is in fact well-defined even for small values of the trigger probability π_a despite the presence of the logarithmic term $\ln \pi_a$.

aggregate risk aversion and wealth of the ceding insurer and CAT investor, all of which may vary across the sample.

With the limiting case of constant absolute risk aversion (i.e., $\gamma \rightarrow 0$), it can be shown that the model (14) yields the following specification (hereafter referred to as the CARA model):

$$S = EL^{1-\frac{1}{\alpha}} + \epsilon. \quad (15)$$

It is interesting to note that with the limiting case of risk neutrality both the HARA and CARA models reduce to $S = EL + \epsilon$, i.e., the predicted spread equals the expected loss. In contrast, the presence of risk aversion implies that the predicted spread is above the expected loss, meaning that a positive absolute risk premium (i.e., $S - EL > 0$) is required by CAT investors in equilibrium. In our framework, this is because natural disasters lower the aggregate wealth,¹³ and CAT bonds are thus priced as assets that hedge against states of the world with lower wealth.

In sum, our proposed HARA model complements the actuarially oriented models that have been proposed in the previous literature (see, e.g., Bodoff and Gan (2012), Galeotti et al. (2013) and Braun (2015)), and helps to fill the gap between the econometric and the theoretical approaches for pricing CAT bonds. In the next section, we present nonparametric specification tests that will be used to determine the adequacy of CAT bond pricing specifications. We will then use these tests to compare the HARA model with some of the most popular specifications that have been proposed in the previous literature.

3 Nonparametric specification tests

The adequacy of CAT bond pricing models is generally tested by comparing various specifications in terms of in-sample fit and predictive power, such as in Galeotti et al. (2013) and Braun (2015). This approach can indicate which of these models is most appropriate, but it does not provide evidence of the validity of the selected model. In other words, one model may outperform all the others tested, yet not appropriately describe the mechanism that generated the data. Here we

¹³See the discussion just before Equation (6).

propose to use an alternative approach that allows us to test simultaneously all possible departures from the model of interest, by using nonparametric estimation techniques. This general idea, which is sometimes referred to as nonparametric specification testing, can be implemented in various ways, as described, for example, in Azzalini et al. (1989), Zheng (1996), Davison and Hinkley (1997) and Li and Racine (2007).

We use two different nonparametric specification tests. In the first test, we estimate the relationship between the two variables using both the parametric model of interest and a nonparametric model. The nonparametric model can estimate the relationship between two variables without any assumption of a parametric form for the model. With this flexibility, the nonparametric estimate will in general closely resemble the true model if the sample size is sufficiently large. The parametric estimate, however, may not fit the data as well if the model is mis-specified. We can thus test for the adequacy of the parametric model by comparing the fits of the two estimators, as defined by the mean squared error (MSE). The second test uses the fact that, for an adequately specified parametric model, the residual error should be centered at zero over all values of the independent variable. Thus, a nonparametric regression of the residuals on the independent variable should be identically zero everywhere, which can be tested using the method just described for the first test. Note that both tests involve some slightly complicated test statistics, for which the null distribution is not easily derived. We thus estimate the significance levels of the observed statistic using the bootstrap, as described in Efron and Tibshirani (1994) and Davison and Hinkley (1997). Details of both tests are given in the following sections, but first let us motivate the use of nonparametric regression for specification tests by way of a simple illustrative example.

3.1 Motivation for nonparametric specification tests

To justify the nonparametric specification tests used in this work, we consider a simple numerical experiment in which the data-generating process has the following nonlinear specification:

$$y = \ln(x + 1) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, 0.03), \quad (16)$$

where x is the explanatory variable and $\mathcal{N}(0, 0.03)$ denotes the normal distribution with zero mean and variance 0.03. From this specification, we first simulate a data sample $\{(x_i, y_i)\}$ of size $N = 300$ in which the variable x is uniformly distributed in the interval $[0, 3]$. We then fit the following regression equations to the simulated data:

$$y_i = a + bx_i + \epsilon_i \quad \text{and} \quad y_i = a + b \ln(x_i + 1) + \epsilon_i. \quad (17)$$

The first function describes an (incorrect) linear model, whereas the second one has the appropriate nonlinear specification. The simulated sample and the two fitted curves are presented in the top graph of Figure 1, in which a very subtle nonlinearity in the data can be seen. Standard goodness-of-fit tests on the residuals of the linear model indicate no departures from normality.¹⁴ Moreover, the adjusted R^2 of the two regressions are quite close (0.816 for the linear model and 0.825 for the nonlinear one). This example shows how challenging it can be to determine whether a particular specification should be discarded in favor of another one.¹⁵

Next we present nonparametric tests that have adequate flexibility to detect all kinds of systematic errors associated with a mis-specified model. In the case of the example illustrated here, we will show that using these tests results in the rejection of the linear specification at the significance level of 0.02%.

3.2 Testing the in-sample accuracy

If a parametric model is well-specified, then it should not predict significantly less well than a nonparametric one. In fact, a well-specified parametric model can generate even better in-sample accuracy than nonparametric models. This section presents a procedure that, based on this intuition, tests a functional regression specification by determining whether its mean squared error is significantly higher than the mean squared error of a nonparametric regression.

¹⁴We obtain a p -value of 0.4474 for the Kolmogorov-Smirnov test, a p -value of 0.6627 for the Shapiro-Wilk test, and a p -value of 0.5559 for the Jarque-Bera test.

¹⁵Note that this numerical experiment was repeated a thousand times and that the results were invariably similar.

Here, we only consider parametric models that can be written in the following form:

$$y = f(x; \theta) + \epsilon, \quad \epsilon \sim \mathcal{D}(0, \sigma^2(x)), \quad (18)$$

where $f(\cdot; \theta)$ is a function with (unknown) parameters θ , and ϵ is uncorrelated noise from a distribution $\mathcal{D}(0, \sigma^2(x))$ that has zero mean and variance $\sigma^2(x)$. Note that this is a very general class of models that nests several of the most popular econometric CAT bond pricing models that have been proposed in the literature.

To test for the presence of systematic errors in the parametric specification, we compare it to a nonparametric fit. Nonparametric regression models can in general be written in the form

$$y = m_x(x) + \epsilon, \quad \epsilon \sim \mathcal{D}(0, \sigma^2(x)), \quad (19)$$

where $m_x(\cdot)$ is the regression function associated with point x . The specific nonparametric model that will be used in this work is a local polynomial kernel regression, a description of which can be found in the appendix.

The general idea is to then use the fitted nonparametric regression to check for all kinds of mis-specifications in the fit of the parametric model. One such possible approach is to test whether the parametric model predicts at least as well as the nonparametric one, which should be the case in the absence of mis-specifications. In practice, this test can be implemented by looking at the difference between the mean squared errors of the parametric and nonparametric regressions, which are respectively given by

$$\text{MSE}_p = \frac{1}{N} \sum_{i=1}^N \left(y_i - f(x_i; \hat{\theta}) \right)^2, \quad \text{MSE}_{np} = \frac{1}{N} \sum_{i=1}^N \left(y_i - \hat{m}_{x_i}(x_i) \right)^2,$$

where N is the total number of observations, $\hat{\theta}$ is the estimator of the parametric model and $\hat{m}_{x_i}(\cdot)$ is the estimator of the nonparametric model at point x_i .

More generally, one could be interested in determining whether the parametric model is well-

specified within a precise domain \mathcal{X} of the independent variable x . In this case, the relevant MSEs are

$$\text{MSE}_{\text{p}}^{\mathcal{X}} = \frac{1}{N_{\mathcal{X}}} \sum_{i=1}^N \left(y_i - f(x_i; \hat{\theta}) \right)^2 \mathbf{1}_{\{x_i \in \mathcal{X}\}}, \quad \text{MSE}_{\text{np}}^{\mathcal{X}} = \frac{1}{N_{\mathcal{X}}} \sum_{i=1}^N \left(y_i - \hat{m}_{x_i}(x_i) \right)^2 \mathbf{1}_{\{x_i \in \mathcal{X}\}}, \quad (20)$$

where $N_{\mathcal{X}} = \#\{x_i | x_i \in \mathcal{X}\}$ is the number of observations within the tested region \mathcal{X} and $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function. The parametric specification can then be tested formally in the region of interest by using the spread $\text{MSE}_{\text{p}}^{\mathcal{X}} - \text{MSE}_{\text{np}}^{\mathcal{X}}$ as the test statistic.

The bootstrap can be used to obtain the distribution of this test statistic under the null hypothesis that the parametric model specification is correct. However, since it is generally recommended that one stabilizes the sampling distribution of a test statistic when using the bootstrap (see, e.g., Tibshirani (1988)), we use the following statistic:

$$\hat{t}_{\mathcal{X}} = \frac{\text{MSE}_{\text{p}}^{\mathcal{X}} - \text{MSE}_{\text{np}}^{\mathcal{X}}}{\text{MSE}_{\text{np}}^{\mathcal{X}}}. \quad (21)$$

The steps for computing the test statistic (21) and its p -value are summarized below:

1. Fit the parametric model to the data to obtain the estimator $\hat{\theta}$ and the mean squared error $\text{MSE}_{\text{p}}^{\mathcal{X}}$. Fit the nonparametric model to the data to obtain its mean squared error $\text{MSE}_{\text{np}}^{\mathcal{X}}$. Compute the test statistic as $\hat{t}_{\mathcal{X}} = (\text{MSE}_{\text{p}}^{\mathcal{X}} - \text{MSE}_{\text{np}}^{\mathcal{X}}) / \text{MSE}_{\text{np}}^{\mathcal{X}}$.
2. Estimate the variance function $\sigma^2(\cdot)$ of the error term in the parametric model as well as its probability distribution \mathcal{D} . Denote these estimators by $\hat{\sigma}^2(\cdot)$ and $\hat{\mathcal{D}}$ respectively.¹⁶
3. Simulate a sample $\{y'_i = f(x_i; \hat{\theta}) + \eta_i\}$ with $\eta_i \sim \hat{\mathcal{D}}(0, \hat{\sigma}^2(x_i))$, using the $\{x_i\}$ of the original data. Compute the statistic $\hat{t}'_{\mathcal{X}}$ of the simulated data using the procedure in step 1.
4. Repeat step 3 a large number of times and compute the p -value of $\hat{t}_{\mathcal{X}}$ as $\frac{\#\{\hat{t}'_{\mathcal{X}} \geq \hat{t}_{\mathcal{X}}\}}{\#\{\hat{t}'_{\mathcal{X}}\}}$.

The null hypothesis that the parametric model has generated the original data can then be rejected if the p -value is sufficiently small (e.g., $\leq 5\%$).

¹⁶Various methods of implementing this step are presented in Section 4.4.

Applying this approach with $\mathcal{X} = (-\infty, +\infty)$ to the example introduced in the previous section yields the results presented in the lower graphs of Figure 1, where 10,000 draws are used to estimate a p -value of 0.0002 for the linear model of Equation (17) and a p -value of 0.3483 for the nonlinear one. Thus, despite the fact that these two models have quite similar adjusted R^2 (0.816 for the linear model and 0.825 for the nonlinear one), and that the residuals do not indicate departure from normality in either case, the nonparametric approach has sufficient power to reject the linear specification in favor of the correct one.

3.3 Testing for residual patterns

We will now present another nonparametric specification test that will also be used to strengthen our results. The intuition underlying this test is that the residuals from a well-specified parametric regression must be patternless. In other words, there should be no significant relationship between the residuals of the parametric regression and the explanatory variables. If such a relation is detected, the implication is that there exists a form of dependence that is not captured by the regression and that the corresponding model is thus mis-specified.

To apply this test, we first estimate the parametric model (18) to obtain its mean squared error MSE_p and residuals $\{\hat{\epsilon}_i = y_i - f(x_i; \hat{\theta})\}$. The objective is then to test whether $\mathbb{E}[\epsilon | x] = 0, \forall x$. To achieve this, we first carry out a nonparametric regression on the parametric residuals:

$$\hat{\epsilon}_i = g_{x_i}(x_i) + \nu_i, \quad i = 1, \dots, N, \quad (22)$$

where $g_{x_i}(\cdot)$ is the regression function associated with x_i and ν_i is an error term. This regression implies that $\mathbb{E}[\epsilon | x] = g_x(x)$, which is tested against the null hypothesis that $\mathbb{E}[\epsilon | x] = 0, \forall x$. Again, this specification test can be implemented by looking at the difference between the MSEs of the two regressions. The corresponding test statistic is then

$$\hat{r}_{\mathcal{X}} = \frac{\text{MSE}_p^{\mathcal{X}} - \text{MSE}_{\text{res}}^{\mathcal{X}}}{\text{MSE}_{\text{res}}^{\mathcal{X}}}, \quad (23)$$

where \mathcal{X} is the region in which the specification is tested, and

$$\text{MSE}_{\text{res}}^{\mathcal{X}} = \frac{1}{N_{\mathcal{X}}} \sum_{i=1}^N (\hat{\epsilon}_i - \hat{g}_{x_i}(x_i))^2 \mathbf{1}_{\{x_i \in \mathcal{X}\}}, \quad \text{MSE}_{\text{p}}^{\mathcal{X}} = \frac{1}{N_{\mathcal{X}}} \sum_{i=1}^N (\hat{\epsilon}_i - 0)^2 \mathbf{1}_{\{x_i \in \mathcal{X}\}}. \quad (24)$$

To determine the significance of this statistic, we can use the bootstrap as in the previous section.

The steps for carrying out the specification test are enumerated below:

1. Fit the parametric model to the data to obtain the estimator $\hat{\theta}$, the mean squared error $\text{MSE}_{\text{p}}^{\mathcal{X}}$ and the residuals $\{\hat{\epsilon}_i = y_i - f(x_i; \hat{\theta})\}$. Fit a nonparametric model to the residuals to obtain the mean squared error $\text{MSE}_{\text{res}}^{\mathcal{X}}$ and the statistic $\hat{r}_{\mathcal{X}} = (\text{MSE}_{\text{p}}^{\mathcal{X}} - \text{MSE}_{\text{res}}^{\mathcal{X}}) / \text{MSE}_{\text{res}}^{\mathcal{X}}$.
2. Estimate the variance function $\sigma^2(\cdot)$ of the error term in the parametric model as well as its probability distribution \mathcal{D} . Denote these estimators by $\hat{\sigma}^2(\cdot)$ and $\hat{\mathcal{D}}$ respectively.¹⁷
3. Simulate a sample $\{y'_i = f(x_i; \hat{\theta}) + \eta_i\}$ with $\eta_i \sim \hat{\mathcal{D}}(0, \hat{\sigma}^2(x_i))$, using the $\{x_i\}$ of the original data. Compute the statistic $\hat{r}'_{\mathcal{X}}$ of the simulated data using the procedure in step 1.
4. Repeat step 3 a large number of times and compute the p -value of $\hat{r}_{\mathcal{X}}$ as $\frac{\#\{\hat{r}'_{\mathcal{X}} \geq \hat{r}_{\mathcal{X}}\}}{\#\{\hat{r}'_{\mathcal{X}}\}}$.

The null hypothesis that the parametric model has generated the original data can then be rejected if the p -value is sufficiently small (e.g., $\leq 5\%$).

Applying this test with $\mathcal{X} = (-\infty, +\infty)$ and 10,000 draws to the example from Figure 1 results in a test statistic of $\hat{r} = 0.0729$ with a p -value of 0.0002 for the (incorrect) linear model. The linear model is thus rejected at the significance level of 0.02%. In contrast, the (correct) nonlinear model has a test statistic of $\hat{r} = 0.0214$ and an estimated p -value of 0.3354.

To sum up, in this section, we have presented two nonparametric tests that have enough flexibility to be used to test for all kinds of systematic errors in functional regression specifications. We have also illustrated their usefulness with a simple numerical example in which a linear regression explains a very large proportion of the variability in the data despite being mis-specified by design. This example is especially relevant for the case of CAT bond pricing models, since earlier research

¹⁷Various methods of implementing this step are presented in Section 4.4.

indicates that a linear relationship between the spread and the expected loss is appropriate (see, e.g., Bodoff and Gan (2012), Galeotti et al. (2013) and Braun (2015)). In the following sections, we apply the nonparametric specification tests to verify whether this is true. We also test other econometric models to determine which ones the data show to be well-specified.

4 Empirical analysis of CAT bond spreads

We now investigate the adequacy of various econometric CAT bond pricing specifications using traditional and nonparametric approaches. We start by presenting the empirical data set, then fit the parametric models and compare them using the traditional approach, namely by inspecting the residuals and by computing standard goodness-of-fit measures. To formally test each functional specification, we then apply the nonparametric tests presented in the previous section under several modeling assumptions so as to ensure the robustness of the findings. Among the tested models, we find the HARA model to be the only one not rejected at the significance level of 5%. In particular, the linear model, which is a popular specification in the empirical literature on CAT bonds, is rejected at the significance level of 1%. The adequacy of the HARA model is further confirmed in an out-of-sample study using a Monte Carlo cross-validation approach.¹⁸

4.1 Description of the data

The empirical data set used in this work is collected from the annual reports published by Lane Financial LLC,¹⁹ and provides a total of 375 CAT bond tranches, issued between June 1997 and March 2015, for which we have the per annum spread and the expected loss (see Figure 2).²⁰ These annual reviews also contain deal statistics and market characteristics that help to describe the data sample. Regarding the territories covered, about half of the CAT bonds in the sample cover the United States, approximately one quarter are multi-territory, and the rest

¹⁸These results can be found in the appendix.

¹⁹See, e.g., Lane and Beckwith (2015) for a description of the CAT bond issues from Q2 2014 to Q1 2015.

²⁰Other CAT-bond-specific information such as the sponsor, SPV, conditional expected loss, probability of first loss, principal amount, issue date and maturity are also available. Note, however, that the territory covered, the peril and the trigger mechanism are only sporadically available from the annual reports published by Lane Financial LLC.

cover Europe, Japan or other areas. In terms of the underlying peril, about half are multi-peril CAT bonds, and the rest are mainly wind-specific or earthquake-specific.

The data sample is illustrated in Figure 2 as a scatter plot of the spread versus the expected loss. Following Braun (2015), we use standard techniques to test for influential observations (see, e.g., Myers (1990)). In particular, externally studentized residuals are used to identify potential outliers, and Cook’s distance is used to measure the impact of the observations on the estimated regression coefficients. The four transactions labeled in Figure 2 are influential observations detected using this approach when fitting the models presented in Section 4.2.²¹ These observations are further investigated using jackknife-after-bootstrap (Efron, 1992) on the regression’s MSE.²² We find that, for all of the models presented in Section 4.2, these four observations have the potential to affect the regressions and/or the outcomes of the specification tests, and thus conclude that it is better to remove them from the sample. We are thus left with 371 CAT bond transactions for the empirical analysis.

4.2 Competing econometric models

For comparison with our HARA model, we consider some of the most popular econometric CAT bond pricing models that have been proposed and studied in the literature.

In the first and simplest alternative, the spread is a linear function of the expected loss:

$$S = b_0 + b_1EL + \epsilon. \tag{25}$$

This specification, which has been studied several times before (see, e.g., Lane and Mahul (2008), Bodoff and Gan (2012), Galeotti et al. (2013) and Braun (2015)), assumes that the spread consists of a fixed basic premium b_0 plus a fixed multiple b_1 of the expected loss. According to Bodoff and Gan (2012), the constant b_0 may derive from a required rate of return on capital for CAT risk,

²¹For instance, when fitting the linear model (25), which is recommended by previous empirical research (e.g., Galeotti et al. (2013)), the Cook’s distances of these observations are (from left to right in Figure 2) 0.028, 0.867, 0.226 and 0.653, and the corresponding externally studentized residuals are 4.434, 11.962, 2.217, and -3.543.

²²This analysis is available from the authors upon request.

and the expected loss multiplier b_1 could be related to uncertainty in the expected loss estimate.

The second alternative we consider is the quadratic extension of the linear model and it is introduced here to account for the presence of nonlinearity in the empirical data (see Figure 2):

$$S = b_0 + b_1EL + b_2EL^2 + \epsilon. \quad (26)$$

The third specification below has also been investigated in Braun (2015) and was originally developed by the insurance-linked securities (ILS) fund Fermat Capital:

$$S = EL + b_0\sqrt{EL(1 - EL)} + \epsilon, \quad (27)$$

where the parameter b_0 can be interpreted as the ILS Sharpe ratio.

The last alternative assumes the spread to be a power function of the expected loss and was originally introduced by Major and Kreps (2002) in the context of traditional treaties:²³

$$S = b_0EL^{b_1} + \epsilon. \quad (28)$$

Note that the Major-Kreps model (28) has the same form as the CARA model (15), up to a scale factor b_0 . In fact, this scale factor appears in the derivation of the CARA model in the form of an integration constant,²⁴ but is constrained to equal one so as to satisfy the boundary condition $S(EL = 1) = 1$. The derivation of our model presented in Section 2.3 thus provides theoretical background for the Major-Kreps model in addition to a generalization of this model (i.e., the HARA model (14)).

²³This pricing specification was, however, first applied to CAT bonds by Galeotti et al. (2013) and then by Braun (2015).

²⁴See the steps just prior to Equation (13).

4.3 Estimation of the regression models

We will now start to investigate the empirical properties of various econometric CAT bond pricing models, and compare them with the CARA model (15) and the HARA model (14) proposed in this work. We first fit the regression models using our data set, which consists of 371 CAT bond tranches that were issued between June 1997 and March 2015 (see Figure 2). The linear model (25), the quadratic model (26) and the Fermat Capital model (27) are estimated by means of the ordinary least squares approach using R Core Team (2015)'s software. The Major-Kreps model (28), the CARA model (15) and the HARA model (14) are estimated using Elzhov et al. (2013)'s R library, which implements the Levenberg-Marquardt nonlinear least squares approach (see, e.g., Moré (1978)).²⁵ We also apply nonparametric bootstrapping as described in Efron and Tibshirani (1994) and Davison and Hinkley (1997) to obtain variance estimates of the estimated coefficients. The estimated coefficients and their bootstrapped standard errors (estimated from 10,000 draws) are presented in Table 1 for each regression model. The mean squared error MSE_p and the adjusted R^2 of each regression are also displayed. We see that the highest in-sample accuracy is achieved with the HARA model, which has the smallest MSE_p and the largest adjusted R^2 . However, as in the numerical example of Figure 1, the gain over alternative models does not appear to be substantial. This motivates us to push the investigation further using nonparametric tests, but let us gain insights into the adequacy of the econometric models by looking at the fitted regression curves displayed in Figure 3.

In Figure 3, we first notice that the linear model seems to overprice CAT bonds with expected losses below 100 bp, and that the Fermat Capital model and the CARA model seem to underprice these same CAT bonds. Hence, these specifications do not appear capable of generating the appropriate form of nonlinear relation between the spread and the expected loss. In contrast,

²⁵In Galeotti et al. (2013) and Braun (2015), the Major-Kreps model is estimated using a logarithmic transformation. We do agree that this model can be estimated via a linear regression of the form $\ln S = \ln b_0 + b_1 \ln EL + \epsilon$. However, one should be careful when using the results from this regression to predict the spread S . In fact, taking the exponential and then the expectation yields $\mathbb{E}[S] = b_0 EL^{b_1} \mathbb{E}[e^\epsilon]$. Even though $\mathbb{E}[\epsilon] = 0$, it is not true that $\mathbb{E}[e^\epsilon] = 1$. Therefore, this term cannot be omitted from the transformation as is done in Galeotti et al. (2013) and Braun (2015). This may explain why these papers report very low predictive accuracy for the Major-Kreps model (see, e.g., Equation (8) and Figure 6 of Braun (2015)). In this work, we avoid such issues by using the nonlinear least squares method.

the Major-Kreps model and the HARA model seem to better reproduce the observed risk-return relation. In the context of the framework presented in Section 2.3, a nonlinear curve of the spread versus the expected loss indicates the presence of risk aversion. More specifically, the estimator of the γ coefficient in the HARA model is shown to be significantly greater than zero, which points to the case of decreasing absolute risk aversion. The bootstrap percentiles' 99% confidence interval of the estimator $\hat{\gamma}$ is $[0.254, 0.340]$. This finding is in agreement with the theoretical literature on risk aversion, which argues that increasing absolute risk aversion (i.e., risk assets are inferior goods) is an implausible property for an investor's utility function (see, e.g., Arrow (1971)).

To better identify the systematic pricing errors of each model, we now turn to Figure 4, in which the residuals of the regressions are illustrated. To detect residual patterns, we carry out a local linear regression on the residuals as a function of the expected loss variable using Bowman and Azzalini (2014)'s R library, which provides functions for implementing local polynomial kernel regression, as described in the appendix. It can be seen from Figure 4 that the HARA model displays the weakest relation between its residuals and the expected loss EL, especially in the region $EL \in [0, 400 \text{ bp}]$ in which most of the observations (330 out of 371) are located. The residuals of the other models seem to display significant patterns, indicating possible misspecifications. The next section pushes this investigation further by formally testing the adequacy of each specification.

4.4 Implementation details for the nonparametric specification tests

We now apply the nonparametric approach to test various CAT bond pricing specifications. In particular, we are interested in testing the linear specification since previous works claim that this model is well-suited for explaining CAT bond spreads (see, e.g., Lane and Mahul (2008), Bodoff and Gan (2012), Galeotti et al. (2013) and Braun (2015)).

We test the adequacy of each parametric specification using the two nonparametric approaches discussed in Section 3. The first test is based upon the idea that a well-specified parametric regression model should not predict significantly less well than a nonparametric one. As explained

previously, we can formally test the in-sample accuracy of a parametric specification using the statistic $\hat{t}_{\mathcal{X}}$ (see Equation (21)). The second test is based on the idea that residuals from a well-specified parametric regression must be patternless, i.e., the residuals should have zero expectation everywhere. In practice, we test for the presence of residual patterns using the statistic $\hat{r}_{\mathcal{X}}$ (see Equation (23)). An interesting feature of these nonparametric tests is that they allow one to test parametric specifications within a particular range \mathcal{X} of the independent variable, which is the expected loss EL in our case. In other words, we can test whether a functional specification is well-specified in a specific domain $EL \in \mathcal{X}$. In our case, reducing the scope to the region $EL \leq 600$ bp (i.e., $\mathcal{X} = [0, 600 \text{ bp}]$) could possibly increase the power of the tests since this region contains 357 observations whereas the region $EL > 600$ bp contains only 14 observations.²⁶ For this reason, we apply the specification tests in the full region (i.e., $\mathcal{X} = [0, +\infty)$) and also in the subregion $EL \leq 600$ bp (i.e., $\mathcal{X} = [0, 600 \text{ bp}]$).

For the nonparametric regressions involved in the computation of the test statistics $\hat{t}_{\mathcal{X}}$ and $\hat{r}_{\mathcal{X}}$, we use a local linear regression with a Gaussian kernel and the smoothing parameter obtained by means of the generalized cross-validation criterion.²⁷

To estimate the null distribution of each test statistic, we use 10,000 simulated samples of the form $\{S'_i = f(EL_i; \hat{\theta}) + \eta_i\}$, where $f(\cdot; \theta)$ is the function describing the parametric model, $\hat{\theta}$ is the estimator of the parametric model, EL_i is the expected loss of observation $\#i$, and η_i is an error term. In order to simulate the errors $\{\eta_i\}$, we first have to estimate the variance function $\sigma^2(EL)$ and the distribution of the error term. Various methods of performing these tasks are described below.

For the sake of robustness, we employ the three error variance function estimators presented in the appendix. In practice, these methods are implemented using Nadaraya-Watson kernel

²⁶The reason is that the nonparametric estimator is much more variable in the region $EL > 600$ bp so as to fit more closely the few data points in this region (see, e.g., Figure 4). The MSE of the nonparametric estimator is thus smaller than the parametric one because there are too few data points, but not necessarily because the parametric model is mis-specified. As a consequence, including the region $EL > 600$ bp in the specification tests can increase the test statistics ($\hat{t}_{\mathcal{X}}$ and $\hat{r}_{\mathcal{X}}$) for the wrong reason. The exact same thing happens to the null distributions of the test statistics (which are obtained by bootstrapping). The specification testing procedure remains technically correct, but the statistical power of the tests could be reduced.

²⁷This method is presented in the appendix. We implement it using Bowman and Azzalini (2014)'s R library.

regression (see Nadaraya (1964) and Watson (1964)) with an adaptive bandwidth.²⁸ The three estimators of the error standard deviation function are presented in Figure 5 for the case of the HARA model,²⁹ where we see that error variance increases with the expected loss until it reaches a level at which it stays roughly constant. This behavior seems to be in agreement with Figure 4, which illustrates the residuals of each regression model.

Regarding the distribution used to simulate the pricing errors $\{\eta_i\}$, we first consider a normal distribution $N(0, \sigma^2(\text{EL}))$, where EL is the expected loss of the CAT bond. The results of the nonparametric specification tests under the normal distribution approximation are presented in Table 2. To ensure robust findings, we also consider an alternative approach, which uses the bootstrapped distribution of the standardized parametric residuals. Since we do not want to assume homoskedasticity of the errors, we must first standardize the parametric residuals $\{\hat{\epsilon}_i\}$ by computing $\{z_i \equiv \hat{\epsilon}_i / \hat{\sigma}(\text{EL}_i)\}$, where $\hat{\sigma}^2(\cdot)$ is the error variance function estimator. We then bootstrap the z_i by sampling from their empirical distribution so as to obtain a bootstrapped sample $\{z_i^*\}$. The pricing errors are then obtained as $\{\eta_i = z_i^* \hat{\sigma}(\text{EL}_i)\}$. Note that this is appropriate as long as the distributions of the errors at different values of EL all come from the same scale family. In other words, the only difference between the distributions of the errors for different ELs must be their variances. The results of the nonparametric specification tests under the bootstrap distribution approximation are presented in Table 3.

4.5 Results of the nonparametric specification tests

We now turn to Tables 2 and 3, in which the results of the specification tests are presented under the normal and bootstrapped distribution approximations respectively. The tested specifications are the linear model (25), the quadratic model (26), the Fermat Capital model (27), the Major-Kreps model (28), the CARA model (15) and the HARA model (14). We require rejection of the null hypothesis under all three error variance estimators presented in the appendix in order to conclude that a model is robustly rejected. We can see in Tables 2 and 3 that the HARA model is

²⁸The adaptive bandwidth is calibrated locally to contain the 120 closest-neighbor observations.

²⁹The results for the other models are similar and available from the authors upon request.

the only model not robustly rejected at the significance level of 5% by any of the nonparametric specification tests. In fact, the p -values of the test statistics $\hat{t}_{\mathcal{X}}$ and $\hat{r}_{\mathcal{X}}$ are well above 0.05. In contrast, the p -values for the linear model, the Fermat Capital model and the CARA model are all below 0.01, implying that these specifications can be rejected at the significance level of 1%. The adequacy of the linear model for explaining CAT bond spreads, as claimed in the previous literature, is thus strongly rejected by the nonparametric tests. The quadratic model is also rejected at the 5% level, but only when the scope of the tests is reduced to the region $EL \leq 600$ bp (i.e., $\mathcal{X} = [0, 600 \text{ bp}]$). Finally, the Major-Kreps model is also robustly rejected at the 5% significance level by the $\hat{t}_{\mathcal{X}}$ test, and at the 6% significance level by the $\hat{r}_{\mathcal{X}}$ test.³⁰

5 Summary and conclusion

Most empirical works on CAT bonds aim to identify the premium calculation models that best describe the empirical data. However, as pointed out by Braun (2015), the previous literature relies almost exclusively on actuarial or pure econometric models. In particular, the utility-based pricing approaches that have been put forward in the literature on state-contingent securities have not yet been tested empirically in the context of CAT bonds. We work towards filling this gap by deriving a utility-based specification for pricing CAT bonds under hyperbolic absolute risk aversion. This model, which we refer to as the HARA model is straightforward to test empirically and is underpinned by the theory of equilibrium pricing.

Another issue that we address is related to the methods used to determine the adequacy of CAT bond pricing specifications. Most of the previous works (see, e.g., Bodoff and Gan (2012), Galeotti et al. (2013) and Braun (2015)) investigate this question by comparing the specifications in terms of in-sample fit and predictive power. Although this approach may indicate which of these models is most appropriate, it does not provide evidence of the validity of the selected model. In other words, one model may outperform all the others tested, yet not appropriately describe

³⁰The results presented here were also investigated within subsamples of the data, using a Monte Carlo subsamples approach. To be more specific, we applied the nonparametric specification tests to a large number (i.e., 1,000) of randomly selected subsamples. These results, which are difficult to summarize in a few words, support the evidence presented in Tables 2 and 3 and are available from the authors upon request.

the mechanism that generated the data. We propose, here, two specification tests that allow us to test simultaneously all possible departures from the model of interest, by using nonparametric estimation techniques. The first test is based upon the idea that a well-specified parametric regression model should not predict significantly less well than a nonparametric one, which will in general closely resemble the true model if the sample size is large enough. The second test is based on the idea that residuals from a well-specified parametric regression should be centered at zero over all values of the independent variable, implying that a nonparametric regression of the residuals on the independent variable should be identically zero everywhere. To the best of our knowledge, this is the first time that nonparametric regressions have been proposed for specification testing in the context of CAT bond premium calculation models.

Using an empirical data set that consists of 371 CAT bond tranches issued between June 1997 and March 2015, we investigate the adequacy of the HARA model by using both the traditional approach and the nonparametric approach mentioned above. We also test popular specifications that have been put forward in the previous literature. In particular, the linear model is rejected at the significance level of 1%. Among the tested models, we find that the HARA model is the only one not rejected at the significance level of 5%. These results provide the first empirical evidence that risk aversion is well suited to explain the risk-return relationship observed in the market for CAT bonds. The theoretical framework and the statistical methods we propose allow us to formulate and test hypotheses about the aggregate utility function of CAT investors, and suggest that it exhibits decreasing absolute risk aversion. We believe that a generalization of our framework will contribute to a broader understanding of the role that risk aversion plays in the pricing of CAT bonds, especially as more empirical data will become available. In future research, it could also be insightful to apply the contingent claims approaches that have been proposed in the literature on insurance-linked securities to estimate CAT bond spreads.

Appendix

Local polynomial kernel regression

In general, a nonparametric regression model in which y is the dependent variable and x is the explanatory one can be expressed in the following form:

$$y_i = m_{x_i}(x_i) + \epsilon_i, \quad i = 1, \dots, N,$$

where $m_{x_i}(\cdot)$ is the regression function associated with x_i , ϵ_i is an error term with zero mean and N is the number of observations.

In local polynomial kernel regression, the regression function $m_{x_i}(\cdot)$ associated with point x_i is a polynomial of degree $p \geq 1$ that can be expressed as (see, e.g., Chapter 5 of Wasserman (2007))

$$m_{x_i}(x) = \beta_{x_i,0} + \beta_{x_i,1}(x_i - x) + \beta_{x_i,2}(x_i - x)^2 + \dots + \beta_{x_i,p}(x_i - x)^p.$$

The choice $p = 0$ can also be implemented in practice, but is generally referred to as Nadaraya-Watson kernel regression (see Nadaraya (1964) and Watson (1964)). For each x_i , the associated regression function estimator is obtained as the solution to a kernel-weighted least squares problem:

$$\hat{m}_{x_i}(\cdot) = \arg \min_{\beta_{x_i,0}, \dots, \beta_{x_i,p}} \left\{ \sum_{j=1}^N (y_j - m_{x_i}(x_j))^2 \frac{1}{h} K\left(\frac{x_i - x_j}{h}\right) \right\}, \quad (29)$$

where h is the smoothing parameter and $K(\cdot)$ is a kernel, i.e., a function symmetric around 0 with the following properties:

$$\int K(u)du = 1, \quad \int uK(u)du = 0, \quad \int u^2K(u)du > 0.$$

In practice, several types of kernel functions are commonly used. For instance, the Gaussian

kernel and Epanechnikov (1969)'s kernel are respectively given by

$$K(u) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}u^2}, \quad K(u) = \frac{3}{4}(1-u^2)\mathbf{1}_{\{|u|\leq 1\}},$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function.

As for the choice of the smoothing parameter h , many criteria exist (see Härdle (1990)). For instance, the generalized cross-validation criterion is the one used by default in the R package of Bowman and Azzalini (2014), and consists of selecting the value of h that minimizes

$$\text{GCV}(h) = \frac{N}{N-\nu} \sum_{i=1}^N (y_i - \hat{y}_i)^2, \quad \hat{y}_i = \hat{m}_{x_i}(x_i),$$

where ν is the trace of the smoothing matrix M , defined such that $(\hat{y}_1, \dots, \hat{y}_N)^\top = M(y_1, \dots, y_N)^\top$.

Variance function estimation

This section of the appendix provides theoretical background for the variance function estimation methods presented in Table 4. The purpose of such methods is to estimate the variance function $\sigma^2(\cdot)$ of the error term in a generic parametric regression model of the form

$$y_i = f(x_i; \theta) + \sigma(x_i)\epsilon_i, \quad i = 1, \dots, N, \quad (30)$$

where $f(\cdot; \theta)$ is a function with (unknown) parameters θ , the $\{\epsilon_i\}$ are i.i.d. variables from a distribution that has zero mean and unit variance, and N is the number of observations.

The first approach presented in Table 4 is described in Wasserman (2007) and estimates $\ln \sigma^2(x)$ by regressing the log squared residuals on the explanatory variable x . To understand why, we simply have to use Equation (30) to show that

$$\ln(y_i - f(x_i; \theta))^2 = \ln \sigma(x_i)^2 + \ln \epsilon_i^2 \quad \Rightarrow \quad \mathbb{E}[\ln(y_i - f(x_i; \theta))^2] = \ln \sigma(x_i)^2 + \mathbb{E}[\ln \epsilon_i^2].$$

The second method presented in Table 4 estimates $\sigma^2(x)$ by regressing the squared residuals on

the explanatory variable x . This is a frequently used approach (see, e.g., Yu and Jones (2004)) which we can justify by rewriting Equation (30) and using the property $\mathbb{E}[\epsilon_i^2] = 1$ to show that

$$(y_i - f(x_i; \theta))^2 = \sigma(x_i)^2 \epsilon_i^2 \quad \Rightarrow \quad \mathbb{E}[(y_i - f(x_i; \theta))^2] = \sigma(x_i)^2.$$

The third method presented in Table 4 is a kernel-weighted version of Rice (1984)'s difference-based estimator, and has the benefit of being independent of the function $f(\cdot; \theta)$ used for the regression model. This approach assumes that the $\{x_i\}$ are ordered and that the functions $f(\cdot; \theta)$ and $\sigma(\cdot)$ are sufficiently smooth to allow the approximations $f(x_{i+1}; \theta) \approx f(x_i; \theta)$ and $\sigma(x_i) \approx \sigma(x_{i+1})$. Starting from Equation (30), we can then use these approximations together with the properties $\mathbb{E}[\epsilon_i^2] = \mathbb{E}[\epsilon_{i+1}^2] = 1$ and $\mathbb{E}[\epsilon_{i+1}\epsilon_i] = 0$ to show that

$$(y_{i+1} - y_i)^2 \approx (\sigma(x_i)\epsilon_{i+1} - \sigma(x_i)\epsilon_i)^2 \quad \Rightarrow \quad \mathbb{E}[(y_{i+1} - y_i)^2] \approx 2\sigma(x_i)^2.$$

Monte Carlo cross-validation study

Following Galeotti et al. (2013) and Braun (2015), the out-of-sample pricing performance of the econometric specifications is also investigated. We use the same empirical data set as before (see Figure 2), which consists of 371 CAT bond tranches issued between June 1997 and March 2015.

We consider a Monte Carlo cross-validation experiment in which each repetition consists of randomly selecting a subsample of size N_C that is used to calibrate the parametric model. We then measure the out-of-sample accuracy of the model using the remaining $N_T = 371 - N_C$ observations. To do so, we compute the mean absolute error (MAE), the root mean squared error (RMSE) and the out-of-sample R^2 (denoted R_{OS}^2), which are given by

$$\text{MAE} = \sum_{i=1}^{N_T} \frac{|S_i - \hat{S}_i|}{N_T}, \quad \text{RMSE} = \sqrt{\sum_{i=1}^{N_T} \frac{(S_i - \hat{S}_i)^2}{N_T}}, \quad R_{OS}^2 = 1 - \frac{\sum_{i=1}^{N_T} (S_i - \hat{S}_i)^2}{\sum_{i=1}^{N_T} (S_i - \bar{S})^2}, \quad (31)$$

where N_T is the number of observations in the test sample, \bar{S} is the mean spread in the calibration

sample and \hat{S}_i is the predicted spread for observation $\#i$ in the test sample.

The results of the Monte Carlo cross-validation experiment are presented in Table 5 for calibration samples of size $N_C \in \{50, 100, 200\}$, and show that the best out-of-sample performance is achieved with the HARA model. This further confirms the results of the nonparametric specification tests.

References

- Arrow, K. J. (1971). *Essays in the theory of risk-bearing*. Markham Publishing Co.
- Arrow, K. J. and Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 22(33):265–290.
- Azzalini, A., Bowman, A. W., and Härdle, W. (1989). On the use of nonparametric regression for model checking. *Biometrika*, 76(1):1–11.
- Back, K. (2010). *Asset pricing and portfolio choice theory*. Oxford University Press.
- Bodoff, N. M. and Gan, Y. (2012). An analysis of the market price of Cat bonds. *Variance*, 6(2):161–177.
- Bowman, A. W. and Azzalini, A. (2014). *R package sm: Nonparametric smoothing methods* (version 2.2-5.4). University of Glasgow, UK and Università di Padova, Italia.
- Braun, A. (2015). Pricing in the primary market for Cat bonds: New empirical evidence. *Journal of Risk and Insurance*. Article first published online : DOI: 10.1111/jori.12067.
- Breeden, D. T. and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651.
- Burnecki, K., Janczura, J., and Weron, R. (2011). Building loss models. In Cizek, P., Härdle, W. K., and Weron, R., editors, *Statistical Tools for Finance and Insurance*, chapter 9. Springer Berlin Heidelberg.
- Chang, C. W., Chang, J. S. K., and Lu, W. (2010). Pricing catastrophe options with stochastic claim arrival intensity in claim time. *Journal of Banking & Finance*, 34(1):24–32.
- Constantinides, G. M. (1982). Intertemporal asset pricing with heterogeneous consumers and without demand aggregation. *Journal of Business*, 55(2):253–267.
- Cox, S. H. and Pedersen, H. W. (2000). Catastrophe risk bonds. *North American Actuarial Journal*, 4(4):56–82.

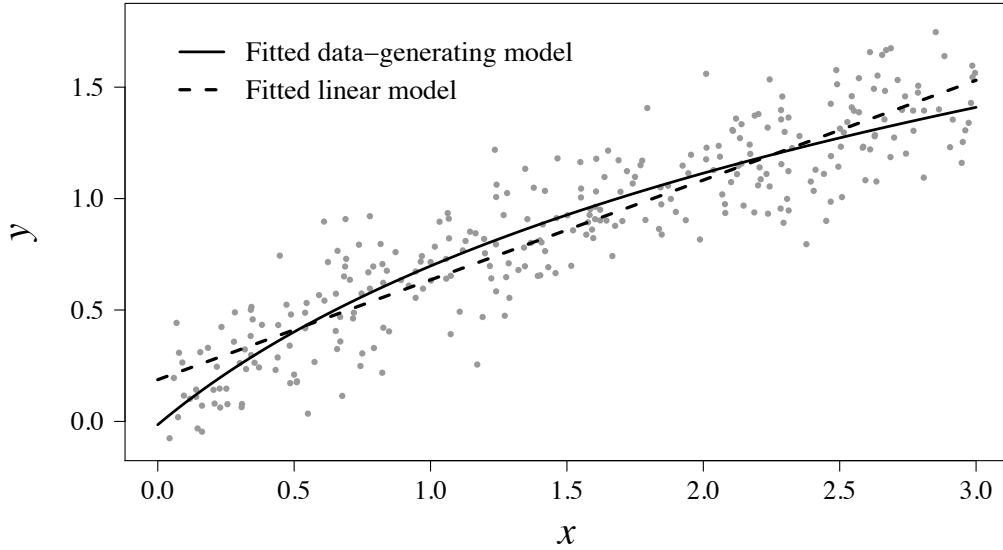
- Cummins, J. D. and Trainar, P. (2009). Securitization, insurance, and reinsurance. *Journal of Risk and Insurance*, 76(3):463–492.
- Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap methods and their application*, volume 1. Cambridge University Press.
- Dieckmann, S. (2010). By force of nature: Explaining the yield spread on catastrophe bonds. Working Paper, University of Pennsylvania.
- Efron, B. (1992). Jackknife-after-bootstrap standard errors and influence functions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 54(1):83–127.
- Efron, B. and Tibshirani, R. J. (1994). *An introduction to the bootstrap*. CRC press.
- Egami, M. and Young, V. R. (2008). Indifference prices of structured catastrophe (CAT) bonds. *Insurance: Mathematics and Economics*, 42(2):771–778.
- Elzhov, T. V., Mullen, K. M., Spiess, A.-N., and Bolker, B. (2013). *minpack.lm: R interface to the Levenberg-Marquardt nonlinear least-squares algorithm found in MINPACK, plus support for bounds*. R package version 1.1-8.
- Epanechnikov, V. A. (1969). Non-parametric estimation of a multivariate probability density. *Theory of Probability and its Applications*, 14(1):153–158.
- Galeotti, M., Gürtler, M., and Winkelvos, C. (2013). Accuracy of premium calculation models for CAT bonds - An empirical analysis. *Journal of Risk and Insurance*, 80(2):401–421.
- Gürtler, M., Hibbeln, M., and Winkelvos, C. (2014). The impact of the financial crisis and natural catastrophes on CAT bonds. *Journal of Risk and Insurance*. Article first published online : DOI: 10.1111/jori.12057.
- Härdle, W. (1990). *Applied nonparametric regression*, volume 27. Cambridge University Press.
- Jarrow, R. A. (2010). A simple robust model for Cat bond valuation. *Finance Research Letters*, 7(2):72–79.

- Lai, V. S., Parcollet, M., and Lamond, B. F. (2014). The valuation of catastrophe bonds with exposure to currency exchange risk. *International Review of Financial Analysis*, 33(2):243–252.
- Lane, M. N. (2000). Pricing risk transfer transactions. *ASTIN Bulletin*, 30(2):259–293.
- Lane, M. N. and Beckwith, R. G. (2015). Crawling along or coming off bottom? Annual review for the four quarters, Q2 2014 to Q1 2015. *Lane Financial LLC Trade Notes*. March 31.
- Lane, M. N. and Mahul, O. (2008). Catastrophe risk pricing: An empirical analysis. Working Paper, The World Bank.
- Lee, J. P. and Yu, M. T. (2002). Pricing default-risky CAT bonds with moral hazard and basis risk. *Journal of Risk and Insurance*, 69(1):25–44.
- Lei, D. T., Wang, J. H., and Tzeng, L. Y. (2008). Explaining the spread premiums on catastrophe bonds. Working Paper.
- Li, Q. and Racine, J. S. (2007). *Nonparametric econometrics: Theory and practice*. Princeton University Press.
- Lipschutz, S., Spiegel, M., and Liu, J. (2012). *Schaum's outline of mathematical handbook of formulas and tables*. McGraw-Hill Education, 4th edition.
- Ma, Z. G. and Ma, C. Q. (2013). Pricing catastrophe risk bonds: A mixed approximation method. *Insurance: Mathematics and Economics*, 52(2):243–254.
- Major, J. A. and Kreps, R. E. (2002). Catastrophe risk pricing in the traditional market. In Lane, M. N., editor, *Alternative Risk Strategies*, chapter 10. Risk Books.
- Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413.
- Mocklow, D., DeCaro, J., and McKenna, M. (2002). Catastrophe bonds. In Lane, M. N., editor, *Alternative Risk Strategies*, chapter 3. Risk Books.

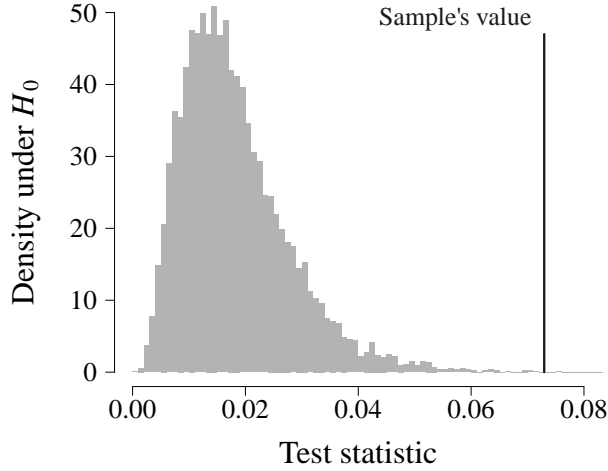
- Moré, J. J. (1978). The Levenberg-Marquardt algorithm: Implementation and theory. In *Numerical analysis*, pages 105–116. Springer.
- Mossin, J. (1968). Optimal multiperiod portfolio policies. *Journal of Business*, 41(2):215–229.
- Myers, R. H. (1990). *Classical and modern regression with applications*, volume 2. Duxbury Press.
- Nadaraya, E. A. (1964). On estimating regression. *Theory of Probability and its Applications*, 9(1):141–142.
- Nowak, P. and Romaniuk, M. (2013). Pricing and simulations of catastrophe bonds. *Insurance: Mathematics and Economics*, 52(1):18–28.
- Papachristou, D. (2011). Statistical analysis of the spreads of catastrophe bonds at the time of issue. *Astin Bulletin*, 41(1):251–277.
- Perrakis, S. and Boloorforoosh, A. (2013). Valuing catastrophe derivatives under limited diversification: A stochastic dominance approach. *Journal of Banking & Finance*, 37(8):3157–3168.
- R Core Team (2015). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. R package version 0.99.441.
- Rice, J. (1984). Bandwidth choice for nonparametric regression. *The Annals of Statistics*, 12(4):1215–1230.
- Rubinstein, M. (1974). An aggregation theorem for securities markets. *Journal of Financial Economics*, 1(3):225–244.
- Têtu, A., Lai, V. S., and Soumaré, I. (2015). Hedging flood losses using Cat bonds. *Asia-Pacific Journal of Risk and Insurance*, 9(2):149–184.
- Tibshirani, R. (1988). Variance stabilization and the bootstrap. *Biometrika*, 75(3):433–444.
- Tower Watson (2010). Catastrophe bonds evolve to address credit risk issues. <http://www.towerswatson.com>.

- Vaugirard, V. E. (2003). Pricing catastrophe bonds by an arbitrage approach. *The Quarterly Review of Economics and Finance*, 43(1):119–132.
- Wang, S. (2000). A class of distortion operators for pricing financial and insurance risks. *Journal of Risk and Insurance*, 67(1):15–36.
- Wang, S. (2004). Cat bond pricing using probability transforms. *Geneva Papers: Etudes et Dossiers*, 278:19–29.
- Wasserman, L. (2007). *All of nonparametric statistics*. Springer.
- Watson, G. S. (1964). Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A*, 26(4):359–372.
- Yu, K. and Jones, M. C. (2004). Likelihood-based local linear estimation of the conditional variance function. *Journal of the American Statistical Association*, 99(465):139–144.
- Zheng, J. X. (1996). A consistent test of functional form via nonparametric estimation techniques. *Journal of Econometrics*, 75(2):263–289.
- Zhu, W. (2011). Ambiguity aversion and an intertemporal equilibrium model of catastrophe-linked securities pricing. *Insurance: Mathematics and Economics*, 49(1):38–46.

(a) Simulated data to motivate nonparametric specification tests



(b) Testing the linear model



(c) Testing the data-generating model

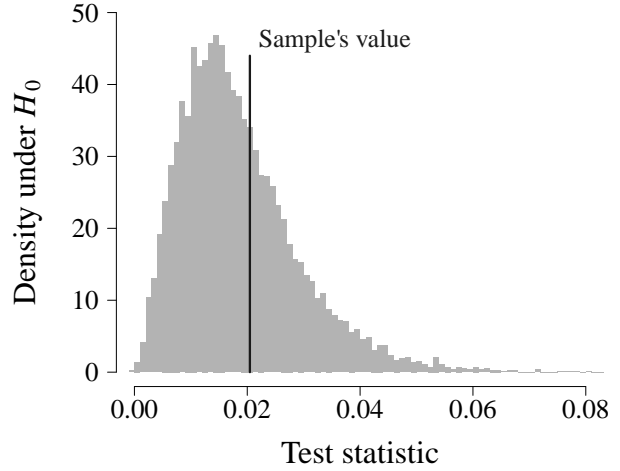


Figure 1: Numerical example to motivate the nonparametric specification tests used in this work. A data sample simulated from the nonlinear model (16) is illustrated in **graph (a)** together with the fitted linear model (dashed line) and the fitted nonlinear model (solid line) of Equation (17). The adjusted R^2 of the two regressions are quite close (0.816 for the linear model and 0.825 for the nonlinear one) and the nonlinearity in the data is very subtle. Both regressions are then tested in the full domain (i.e., $\mathcal{X} = (-\infty, +\infty)$) using the nonparametric test of Section 3.2 with 10,000 bootstrap draws. The sample's test statistic and its distribution under the null hypothesis are presented in **graph (b)** for the linear model, and in **graph (c)** for the nonlinear one. From this, we estimate a p -value of 0.0002 for the linear model and a p -value of 0.3483 for the nonlinear one. The (incorrect) linear model is thus discarded in favor of the correct one.

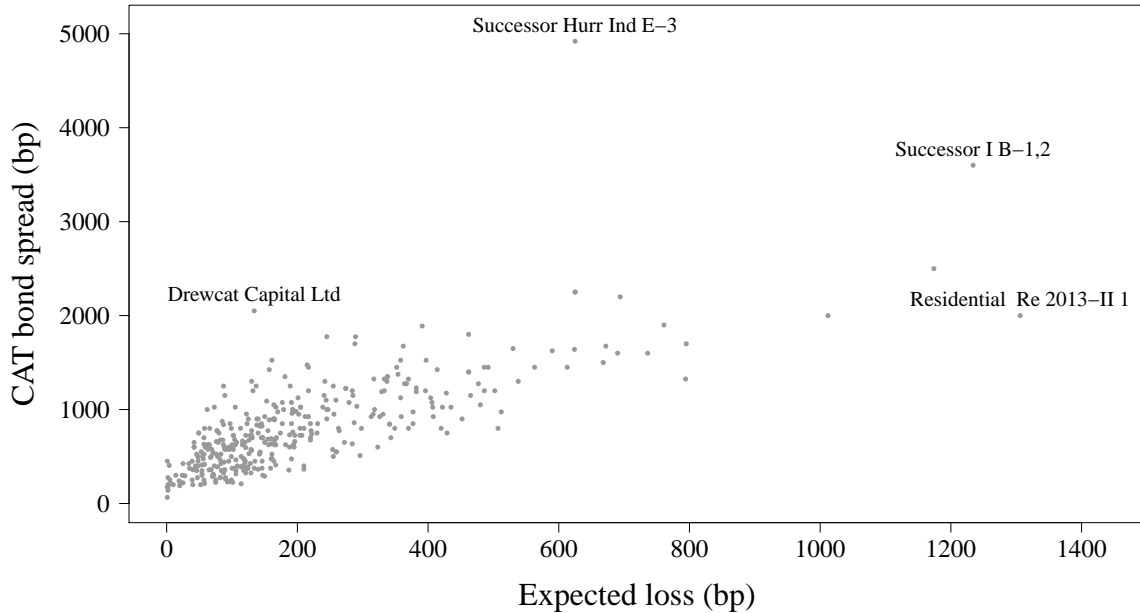


Figure 2: Illustration of the empirical data set, which consists of 375 CAT bond tranches issued between June 1997 and March 2015. When fitting the regression models of Section 4.2, four influential observations (the labeled data points) are detected from their Cook’s distance, externally studentized residual, and/or standardized jackknife value. For instance, when fitting the linear model, which is a popular model in the existing literature on CAT bonds, the Cook’s distances of these observations are (from left to right) 0.028, 0.867, 0.226 and 0.653, and the corresponding externally studentized residuals are 4.434, 11.962, 2.217, and -3.543. We are left with 371 observations after these influential observations are removed from the sample.

Table 1: Regression results for the HARA model proposed in this work and some alternative specifications. The empirical data consist of 371 CAT bond tranches that were issued between June 1997 and March 2015. For each coefficient, the estimator’s value is presented with its bootstrapped standard error (in parenthesis), estimated from 10,000 nonparametric bootstrap replications of the original sample. The adjusted R^2 and the mean squared error MSE_p (in bp^2) are also presented and measure goodness of fit. We can see that the highest in-sample accuracy is achieved with the HARA model.

	Linear $b_0+b_1\text{EL}$	Quadratic $b_0+b_1\text{EL}+b_2\text{EL}^2$	Fermat Capital $\text{EL}+b_0\sqrt{\text{EL}-\text{EL}^2}$	Major-Kreps $b_0\text{EL}^{b_1}$	CARA $\text{EL}^{1-\frac{1}{\alpha}}$	HARA $\text{EL}(1+\frac{\gamma}{\alpha}\ln\text{EL})^{-\frac{1}{\gamma}}$
\hat{b}_0	0.039 (0.002)	0.032 (0.002)	0.436 (0.011)	0.667 (0.059)		
\hat{b}_1	1.977 (0.092)	2.589 (0.223)		0.528 (0.023)		
\hat{b}_2		-8.543 (3.256)				
$\hat{\alpha}$					2.773 (0.039)	3.392 (0.061)
$\hat{\gamma}$						0.302 (0.015)
MSE_p	65974	63964	67155	64355	68712	63430
Adj. R^2	0.624	0.634	0.618	0.633	0.609	0.639

Note: The linear, quadratic and Fermat Capital models are estimated by means of the ordinary least squares approach, whereas the Major-Kreps, CARA and HARA models are estimated using the Levenberg-Marquardt nonlinear least squares approach (see, e.g., Moré (1978)).

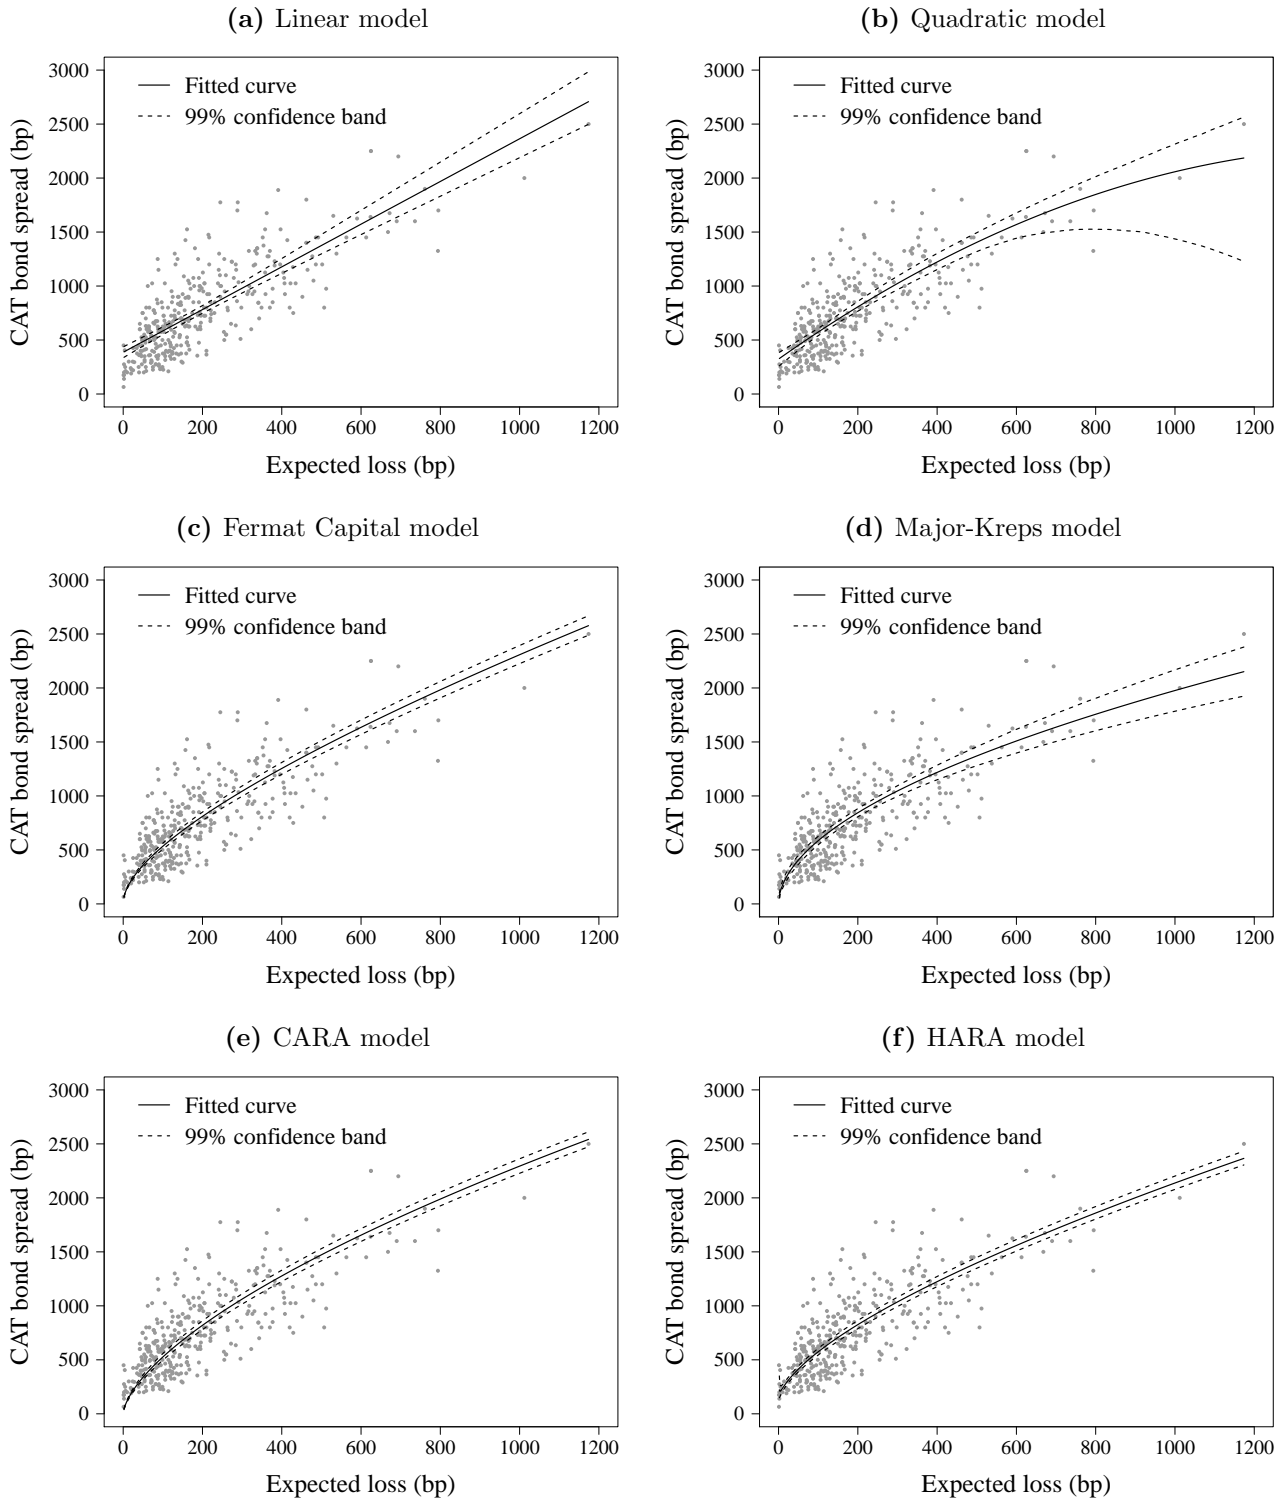


Figure 3: Illustration of the fitted models. For each graph, the solid line illustrates the fitted regression curve and the dotted lines illustrate its 99% bootstrapped confidence band. The linear, quadratic and Fermat Capital models are estimated by means of the ordinary least squares approach, whereas the Major-Kreps and HARA models are estimated using the Levenberg-Marquardt nonlinear least squares approach. The corresponding regression equations are presented in Table 1 and the residuals in Figure 4.

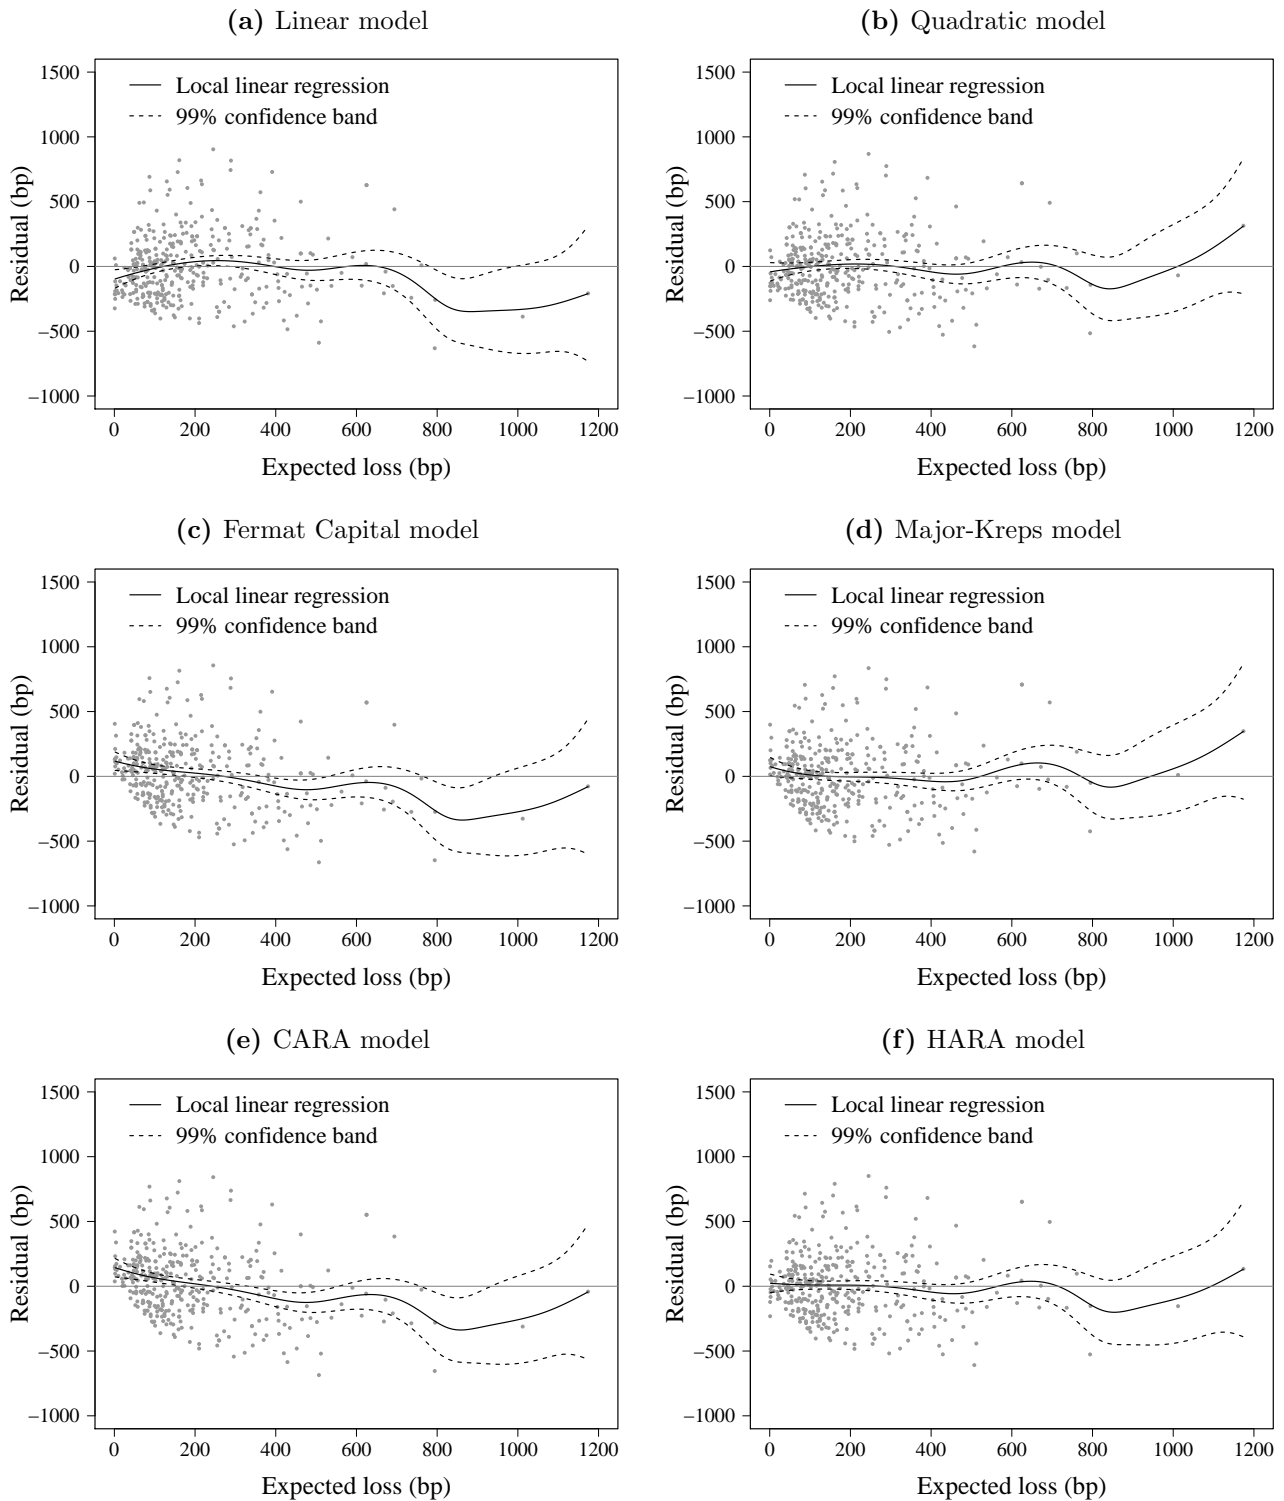


Figure 4: Residuals of the fitted models. For each graph, the solid line presents a local linear regression of the parametric residuals on the expected loss, and the dotted lines illustrate its 99% confidence band. It can be seen that the HARA model displays the weakest relation between its residuals and the expected loss EL, especially in the region $EL \in [0, 400 \text{ bp}]$ in which most of the observations are located. The residuals of the other models seem to display significant patterns, indicating the presence of mis-specifications.




Figure 5: Estimated standard deviation of the error for the HARA model obtained using the three methods presented in the appendix. Note that the results obtained using other pricing specifications are similar to those presented here. We can see that the error variance increases with the expected loss until it reaches a level at which it stays roughly constant. This behavior seems to be in agreement with Figure 4, which shows the residuals of each regression model.

Table 2: Nonparametric specification tests under the assumption that the pricing error follows a normal distribution $N(0, \sigma^2(\text{EL}))$, where EL is the expected loss of the CAT bond. The in-sample accuracy is tested using the statistic $\hat{t}_{\mathcal{X}}$ (see Equation (21)), and the presence of residual patterns is tested for using the statistic $\hat{r}_{\mathcal{X}}$ (see Equation (23)). The region of the expected loss variable EL in which the specifications are tested is denoted by \mathcal{X} . The **top panel** presents the results of tests carried out over the full region (i.e., $\mathcal{X} = [0, +\infty)$), and the **bottom panel** presents the results of tests carried out in the subregion $\text{EL} \leq 600$ bp (i.e., $\mathcal{X} = [0, 600 \text{ bp}]$). The p -values are estimated from 10,000 bootstrap draws under each of the three error variance estimators presented in the appendix. We require rejection of the null under all three error variance estimation methods to reject a parametric model. This ensures that the rejection of a specification is robust to the choice of the variance function estimator. We can see that the HARA model is the only model not rejected at the significance level of 5% by any of the nonparametric specification tests.

	Linear	Quadratic	Fermat Capital	Major-Kreps	CARA	HARA
Specification tests applied to the full region ($\mathcal{X} = [0, +\infty)$)						
$\hat{t}_{\mathcal{X}}$	0.0605	0.0282 n.r.	0.0794	0.0344	0.1045	0.0195 n.r.
$\hat{p}\text{-Val}_{\text{LSR}}$	0.0081	0.0839	0.0002	0.0041	0.0000	0.2907
$\hat{p}\text{-Val}_{\text{SR}}$	0.0036	0.0820	0.0003	0.0113	0.0000	0.2803
$\hat{p}\text{-Val}_{\text{FD}}$	0.0022	0.0640	0.0002	0.0073	0.0000	0.2493
$\hat{r}_{\mathcal{X}}$	0.0605	0.0302 n.r.	0.0831	0.0354 n.r.	0.1084	0.0225 n.r.
$\hat{p}\text{-Val}_{\text{LSR}}$	0.0081	0.0680	0.0002	0.0463	0.0000	0.3943
$\hat{p}\text{-Val}_{\text{SR}}$	0.0036	0.0653	0.0002	0.0716	0.0000	0.3526
$\hat{p}\text{-Val}_{\text{FD}}$	0.0022	0.0526	0.0001	0.0536	0.0000	0.3176
Specification tests applied to the subregion $\text{EL} \leq 600$ bp ($\mathcal{X} = [0, 600 \text{ bp}]$)						
$\hat{t}_{\mathcal{X}}$	0.0405	0.0171	0.0671	0.0167	0.0945	0.0091 n.r.
$\hat{p}\text{-Val}_{\text{LSR}}$	0.0009	0.0384	0.0000	0.0014	0.0000	0.1576
$\hat{p}\text{-Val}_{\text{SR}}$	0.0002	0.0321	0.0000	0.0027	0.0000	0.1614
$\hat{p}\text{-Val}_{\text{FD}}$	0.0002	0.0234	0.0000	0.0010	0.0000	0.1455
$\hat{r}_{\mathcal{X}}$	0.0405	0.0186	0.0707	0.0173 n.r.	0.0982	0.0118 n.r.
$\hat{p}\text{-Val}_{\text{LSR}}$	0.0009	0.0315	0.0000	0.0543	0.0000	0.2313
$\hat{p}\text{-Val}_{\text{SR}}$	0.0002	0.0261	0.0000	0.0549	0.0000	0.2015
$\hat{p}\text{-Val}_{\text{FD}}$	0.0002	0.0179	0.0000	0.0413	0.0000	0.1731

Note: The subscript following $\hat{p}\text{-Val}$ indicates the method used to estimate the variance of the error term ϵ as a function of the expected loss EL. For example, $\hat{p}\text{-Val}_{\text{LSR}}$ means that the p -value is computed under the assumption that $\text{var}(\epsilon | \text{EL}) = \hat{\sigma}_{\text{LSR}}^2(\text{EL})$, where $\hat{\sigma}_{\text{LSR}}^2(\cdot)$ is the variance estimator based on log squared residuals, which is described in the appendix together with the two other estimators.

n.r. Not robustly rejected at the 5% level, i.e., $\max\{\hat{p}\text{-Val}_{\text{LSR}}, \hat{p}\text{-Val}_{\text{SR}}, \hat{p}\text{-Val}_{\text{FD}}\} > 0.05$.

Table 3: Nonparametric specification tests using the bootstrap distribution of the standardized parametric residuals to simulate pricing errors under the null (see note \star below for details). The in-sample accuracy is tested using the statistic $\hat{t}_{\mathcal{X}}$ (see Equation (21)), and the presence of residual patterns is tested for using the statistic $\hat{r}_{\mathcal{X}}$ (see Equation (23)). The region of the expected loss variable EL in which the specifications are tested is denoted by \mathcal{X} . The **top panel** presents the results of tests carried out over the full region (i.e., $\mathcal{X} = [0, +\infty)$), and the **bottom panel** presents the results of tests run in the subregion $\text{EL} \leq 600$ bp (i.e., $\mathcal{X} = [0, 600 \text{ bp}]$). The p -values are estimated from 10,000 bootstrap draws under each of the three error variance estimators presented in the appendix. We require rejection of the null under all three error variance estimation methods to reject a parametric model. This ensures that the rejection of a specification is robust to the choice of the variance function estimator. We can see that the HARA model is the only model not rejected at the significance level of 5% by any of the nonparametric specification tests.

	Linear	Quadratic	Fermat Capital	Major-Kreps	CARA	HARA
Specification tests applied to the full region ($\mathcal{X} = [0, +\infty)$)						
$\hat{t}_{\mathcal{X}}$	0.0605	0.0282 n.r.	0.0794	0.0344	0.1045	0.0195 n.r.
$\hat{p}\text{-Val}_{\text{LSR}}$	0.0090	0.0871	0.0000	0.0130	0.0000	0.3427
$\hat{p}\text{-Val}_{\text{SR}}$	0.0033	0.0843	0.0003	0.0249	0.0000	0.3025
$\hat{p}\text{-Val}_{\text{FD}}$	0.0016	0.0649	0.0000	0.0151	0.0000	0.2677
$\hat{r}_{\mathcal{X}}$	0.0605	0.0302 n.r.	0.0831	0.0354 n.r.	0.1084	0.0225 n.r.
$\hat{p}\text{-Val}_{\text{LSR}}$	0.0090	0.0717	0.0001	0.0499	0.0000	0.3772
$\hat{p}\text{-Val}_{\text{SR}}$	0.0033	0.0686	0.0005	0.0776	0.0000	0.3434
$\hat{p}\text{-Val}_{\text{FD}}$	0.0016	0.0508	0.0003	0.0583	0.0000	0.3065
Specification tests applied to the subregion $\text{EL} \leq 600$ bp ($\mathcal{X} = [0, 600 \text{ bp}]$)						
$\hat{t}_{\mathcal{X}}$	0.0405	0.0171	0.0671	0.0167	0.0945	0.0091 n.r.
$\hat{p}\text{-Val}_{\text{LSR}}$	0.0008	0.0432	0.0000	0.0126	0.0000	0.1931
$\hat{p}\text{-Val}_{\text{SR}}$	0.0004	0.0372	0.0000	0.0162	0.0000	0.1747
$\hat{p}\text{-Val}_{\text{FD}}$	0.0004	0.0292	0.0000	0.0094	0.0000	0.1525
$\hat{r}_{\mathcal{X}}$	0.0405	0.0186	0.0707	0.0173 n.r.	0.0982	0.0118 n.r.
$\hat{p}\text{-Val}_{\text{LSR}}$	0.0008	0.0338	0.0000	0.0552	0.0000	0.2245
$\hat{p}\text{-Val}_{\text{SR}}$	0.0004	0.0296	0.0000	0.0568	0.0000	0.1986
$\hat{p}\text{-Val}_{\text{FD}}$	0.0004	0.0222	0.0000	0.0432	0.0000	0.1703

Note: The subscript following $\hat{p}\text{-Val}$ indicates the method used to estimate the variance of the error term ϵ as a function of the expected loss EL. For example, $\hat{p}\text{-Val}_{\text{LSR}}$ means that the p -value is computed under the assumption that $\text{var}(\epsilon | \text{EL}) = \hat{\sigma}_{\text{LSR}}^2(\text{EL})$, where $\hat{\sigma}_{\text{LSR}}^2(\cdot)$ is the variance estimator based on log squared residuals, which is described in the appendix together with the two other estimators.

\star We first standardize the parametric residuals $\{\hat{\epsilon}_i\}$ by computing $\{z_i \equiv \hat{\epsilon}_i / \hat{\sigma}(\text{EL}_i)\}$, where $\hat{\sigma}^2(\cdot)$ is the error variance function estimator. We then bootstrap the z_i by sampling from their empirical distribution so as to obtain a bootstrapped sample $\{z_i^*\}$. The pricing errors are then obtained as $\{\eta_i = z_i^* \hat{\sigma}(\text{EL}_i)\}$. n.r. Not robustly rejected at the 5% level, i.e., $\max\{\hat{p}\text{-Val}_{\text{LSR}}, \hat{p}\text{-Val}_{\text{SR}}, \hat{p}\text{-Val}_{\text{FD}}\} > 0.05$.

Table 4: Variance function estimation methods used in this work. This table presents the steps for calculating three different estimators of the error variance function $\sigma^2(\cdot)$ in a generic parametric regression model of the form $y_i = f(x_i; \theta) + \sigma(x_i)\epsilon_i$, ($i = 1, \dots, N$), where $f(\cdot; \theta)$ is a function with (unknown) parameters θ , the $\{\epsilon_i\}$ are i.i.d. variables from a distribution that has zero mean and unit variance, and N is the number of observations.

Estimating variance from log squared residuals (LSR)

1. Estimate the parametric regression $y_i = f(x_i; \theta) + \sigma(x_i)\epsilon_i$ to obtain the estimator $\hat{\theta}$.
 2. Define $z_i = \ln(y_i - f(x_i; \hat{\theta}))^2$ for $i = 1, \dots, N$.
 3. Estimate the nonparametric regression $z_i = m_{x_i}(x_i) + \eta_i$, where $m_{x_i}(\cdot)$ is the regression function associated with point x_i and η_i is an error term with zero mean.
 4. Estimate the variance as $\hat{\sigma}_{\text{LSR}}^2(x) = e^{\hat{m}_x(x) - \delta}$, where $\delta \equiv \mathbb{E}[\ln \epsilon^2]$ depends on the assumption for the distribution of ϵ , e.g., $\mathbb{E}[\ln \epsilon^2] = -1.2704$ for the standard normal distribution.
-

Estimating variance from squared residuals (SR)

1. Estimate the parametric regression $y_i = f(x_i; \theta) + \sigma(x_i)\epsilon_i$ to obtain the estimator $\hat{\theta}$.
 2. Define $z_i = (y_i - f(x_i; \hat{\theta}))^2$ for $i = 1, \dots, N$.
 3. Estimate the nonparametric regression $z_i = m_{x_i}(x_i) + \eta_i$, where $m_{x_i}(\cdot)$ is the regression function associated with point x_i and η_i is an error term with zero mean.
 4. Estimate the variance as $\hat{\sigma}_{\text{SR}}^2(x) = \hat{m}_x(x)$.
-

Estimating variance from first differences (FD)

1. Make sure that the $\{x_i\}$ are ordered.
 2. Define $z_i = (y_{i+1} - y_i)^2$ for $i = 1, \dots, N - 1$ and define $z_N = (y_N - y_{N-1})^2$.
 3. Estimate the nonparametric regression $z_i = m_{x_i}(x_i) + \eta_i$, where $m_{x_i}(\cdot)$ is the regression function associated with point x_i and η_i is an error term with zero mean.
 4. Estimate the variance as $\hat{\sigma}_{\text{FD}}^2(x) = \frac{\hat{m}_x(x)}{2}$.
-

Table 5: Out-of-sample pricing performance in Monte Carlo cross-validation for the HARA model and some alternative specifications. Each repetition consists of randomly selecting a subsample of size N_C that is used to calibrate the parametric model and then to measure its out-of-sample accuracy on the remaining $N_T = 371 - N_C$ observations by computing the mean absolute error (MAE, in bp), the root mean squared error (RMSE, in bp) and the out-of-sample R^2 (see Equation (31)). The table presents the average results obtained from 10,000 such repetitions. It can be seen that the best performance is achieved with the HARA model. Note that the empirical data set consists of 371 CAT bond tranches that were issued between June 1997 and March 2015 and is illustrated in Figure 2.

	Linear	Quadratic	Fermat Capital	Major-Kreps	CARA	HARA
$N_C = 50$						
MAE	206.3	207.4	205.3	202.9	208.7	200.5
RMSE	264.1	272.2	262.6	260.8	265.6	260.3
R_{OS}^2	0.611	0.584	0.615	0.619	0.607	0.621
$N_C = 100$						
MAE	203.8	202.3	204.3	200.2	207.7	198.2
RMSE	260.7	261.5	261.1	257.4	264.0	256.2
R_{OS}^2	0.617	0.615	0.616	0.625	0.607	0.629
$N_C = 200$						
MAE	202.6	200.2	203.8	198.9	207.2	196.8
RMSE	258.9	257.1	260.0	255.7	263.1	253.4
R_{OS}^2	0.619	0.624	0.618	0.628	0.606	0.634