Affine Modeling of Credit Risk, Pricing of Credit Events and Contagion

Alain MONFORT*, Fulvio PEGORARO†,
Jean-Paul RENNE‡, and Guillaume ROUSSELLET§

First version: September 2017 *This version: February 2020*

Abstract

We propose a discrete-time affine pricing model that simultaneously allows for (i) the presence of systemic entities by departing from the no-jump condition on the factors' conditional distribution, (ii) contagion effects, (iii) and the pricing of credit events. Our affine framework delivers explicit pricing formulas for default-sensitive securities like bonds and credit default swaps (CDS). We estimate a euro-area multi-country version of the model and address economic questions related to the pricing of sovereign credit risk. We find that both frailty (common factors) and contagion phenomenons are important to account for the joint dynamics of credit spreads. Our results also provide evidence of credit-event pricing, which is at the source of substantial credit risk premiums, even for short maturities. Finally, we extract measures of depreciation-at-default from CDS denominated in different currencies.

JEL Codes: E43, G12.

Key-words: affine credit risk model, Gamma-zero distribution, no-jump condition, contagion, credit-event risk, sovereign credit risk and exchange rates.

*CREST, alain.monfort@ensae.fr

†ACPR (Direction d'Etude et Analyse des Risques) and CREST, pegoraro@ensae.fr

‡Corresponding author. University of Lausanne, HEC, jean-paul.renne@unil.ch

§Desautels Faculty of Management, McGill University, guillaume.roussellet@mcgill.ca

The authors are thankful to Torben ANDERSEN, Christian GOURIEROUX, Damir FILIPOVIC, Jing-Zhi HUANG, Loriano MANCINI, Julien HUGONNIER, Eric JONDEAU, Emanuel MOENCH, Nour MEDDAHI, Gerardo MANZO, Gustavo SCHWENKLER, Ricardo SCHECHTMAN, Patrick AUGUSTIN, Yaroslav MELNYK and Damien ACKERER for their useful comments. We are also grateful to seminar participants at EPFL Brownbag seminar 2016, Paris-Dauphine Finance seminar 2017, and CREST Financial Econometrics Seminar 2016, and to participants to TSE Financial Econometrics Conference 2016, ESEM (Geneva) 2016, NASM (Philadelphia) 2016, GSE Summer Forum 2016 (Barcelona), MFA 2017 (Chicago), EPFL School and Workshop on Dynamic Models in Finance 2017, 10^{th} annual SoFiE Conference (NYU Stern) 2017, IAAE Annual Conference 2017 (Sapporo), Advances in Fixed Income and Macro-Finance Research conference 2017 (Vancouver), Central Bank of Brazil XII Annual Seminar on Risk, Financial Stability and Banking (Sao Paulo) 2017 and C.R.E.D.I.T. 2017 (Venice). We are extremely thankful to Kay GIESEKE (the editor), an anonymous associate editor and two anonymous referees.

1 Introduction

The specification of no-arbitrage asset pricing models is concerned with the formulation of empirically realistic assumptions while maintaining a large degree of tractability. This trade-off is particularly problematic in credit-risk models, which require the modeling of the joint dynamics of risk factors (y_t) and of entity default indicators (d_t), along with their interplay reflecting financial and economic linkages between entities. In the tradition of Duffie and Singleton (1999), closed-form or semi closed-form pricing formulas for defaultable securities can be obtained in an affine intensity-based framework. In this class of models the vector y_t is an affine process in the risk-neutral world and both the default intensities and the risk-free short rate are affine functions of these factors.However, in order to get tractable pricing formula, existing studies usually resort to one or several of the following assumptions – each of them being debated in the theoretical or empirical literature.

First, the dynamics of y_t does not depend on the vector of default indicators d_t (i.e. d_t does not cause y_t in the Granger sense). This assumption usually referred to as the no-jump condition is made in particular in the doubly-stochastic Cox process framework used by e.g. Jarrow and Turnbull (1995), Lando (1998) or Duffie (2005).[1] If y_t contains macroeconomic variables, this condition implies that the modeled entities are not "systemic." While this is reasonable when the entities are firms of small size, it is less realistic when large banks, insurance companies, or supranational and sovereign entities are considered. The no-jump condition is, for instance, relaxed by Bai et al. (2015), who consider systemic firms whose credit events have economy-wide effects. Benzoni et al. (2015) show how Bayesian updating of beliefs triggered by defaults also invalidates the no-jump condition.

[1] A discussion of the no-jump assumption can be found in e.g. Duffie et al. (1996) and Duffie and Singleton (1999). Collin-Dufresne et al. (2004) propose formulas to value defaultable claims using expected risk-adjusted discounting provided that the expectation is taken under a new modified probability measure – which is different from the risk-neutral measure – that puts zero probability on paths where default occurs prior to the maturity, and is thus only absolutely continuous with respect to the risk-neutral probability measure.

Second, the default probabilities of different entities are independent given the path of y_t, hence there are no lagged or instantaneous contagion effects. In contrast, economic and financial linkages imply a significant amount of default clustering and dynamic contagion effects (see Jarrow and Yu 2001; Ait-Sahalia et al. 2014; Bai et al. 2015; Benzoni et al. 2015; Duan and Miao 2016; Azizpour et al. 2018).

Third, contrary to the factors driving the default intensity, the default event of any entity is usually not priced; that is, d_t is absent from the stochastic discount factor (SDF). However, the pricing of credit surprises have been shown to be an important driver of corporate bond returns and to be a possible explanation of the so-called credit spread puzzle (see Driessen, 2005, Huang and Huang, 2012, and Gourieroux et al., 2014).[2]

Fourth, the recovery payment in case of default is typically defined as a constant or predetermined fraction (recovery rate) of an exposure-at-default given by the zero-coupon bond price that would have prevailed in case of no default; this is the recovery of market value (RMV) convention of Duffie and Singleton (1999).[3] However, several studies point to the existence of stochastic recovery rates (e.g. Altman et al. 2005; Das 2009).

Hence, these four restrictive assumptions have been invalidated, one at a time, by theoretical and empirical studies. However, the literature still lacks a framework authorizing these channels being at play simultaneously. Empirically, this brings about the question of knowing which of these channels are the most relevant to explain the term structures of credit spreads and CDSs.

This paper introduces a general discrete-time affine positive credit-risk modeling framework able to simultaneously relax the above-listed assumptions of standard frameworks while maintaining tractable pricing formulas. The asset pricing model is based on the class of Vector Autoregressive

[2] The credit spread puzzle corresponds to the observation that corporate and sovereign bond spreads are seemingly higher than warranted by historical default rates (see e.g. D'Amato and Remolona 2003; Almeida and Philippon 2007; Gabaix 2012; Giesecke et al. 2011).

[3] Alternative modeling conventions are: the recovery of face value (RFV) convention and the recovery of Treasury (RT) convention. While the exposure-at-default is the face value of the considered bond under RFV; it is the value of the otherwise equivalent default-free bond under RT. Whatever the convention used, most existing studies consider constant recovery rates.

Gamma processes introduced by Monfort et al. (2017), generalizing the ARG process introduced by Gourieroux and Jasiak (2006). These non-negative processes belong to the affine class and some of their components can stay at zero for prolonged periods of time. In our model, the default event of each entity i is described by the latter type of components, called credit-event variable and denoted by $\delta_{i,t}$. The default date of any entity is defined as the first date at which $\delta_{i,t}$ becomes strictly positive. The other components of the multivariate process (y_t) are pricing, or risk, factors. Some components of y_t can be common factors, like a short rate, a non-negative transformation of macroeconomic variable, or a "frailty" factor – usually defined as an unobservable risk factor to which all defaultable entities are jointly exposed (e.g. Das et al. 2007).

Our approach accommodates contagion. Feedbacks between credit-event variables (δ_t) capture direct contagion effects, or "mutual excitation" effects (Ait-Sahalia et al. 2014). Indirect contagion (or systemic risk) can also be obtained if the credit-event variables affect some components of y_t which, in turn, influence other credit-event variables.

In the model, the SDF has a standard exponential-affine formulation. Importantly, though, the SDF depends not only on the factors y_t, but also on the credit-event variables δ_t. This formulation allows to price credit events while preserving the affine structure of our multivariate process under the associated risk-neutral measure, which is instrumental to obtain pricing tractability.

We close the model specification of default-sensitive securities' payoffs by assuming, for any entity, a stochastic recovery rate given by an exponential-affine function of (y_t, δ_t). The recovery payoff is defined as the product of this recovery rate and of the exposure-at-default. The three usual types of exposures-at-default are considered: recovery of market value (RMV), recovery of face value (RFV) and recovery of Treasury (RT) (see Brennan and Schwartz 1980; Duffie 1998; Jarrow and Turnbull 1995; Longstaff and Schwartz 1995; Duffie and Singleton 1999). When defaults are rare events, the identification of the driving factors of recovery rates is a challenging task (Pan and Singleton 2008). Our empirical investigation selects the simplest specification of the recovery rate and leaves aside a thorough investigation of the recovery rate parameterization.

We provide closed-form recursive formulas to price defaultable zero-coupon bonds and Credit Default Swaps (CDS), for any maturity. The availability of closed-form formulas hinges on the affine property of the state vector (y_t, δ_t) under the risk-neutral measure. The fact that the physical dynamics of the state vector is also affine is particularly useful when it comes to estimation. Indeed, the dynamics can be written in a convenient vector autoregressive form and moments are easily computed, which opens the door to standard estimation techniques (e.g. Maximum Likelihood, or Generalized Method of Moments, GMM). The Kalman filter can be used when some of the components of the state vector are unobserved.

We use our framework to shed new light on the pricing of euro-area sovereign credit risk. We jointly model the fluctuations of the sovereign CDS term structures of the four largest euro-area countries – France, Germany, Italy and Spain – and of Greece over the period January 2007 to July 2019. The five credit-event intensities are driven by a short rate, a frailty factor and country-specific factors, as in Ang and Longstaff (2013). Contrary to this latter paper, we make the SDF explicit, thereby opening the door to the computation of credit-risk premia – defined as the differences between observed CDS spreads and those that would prevail if agents were not risk-averse. Having an explicit SDF is also an important distinction with respect to the study by Ait-Sahalia et al. (2014), who also estimate affine term-structure models of euro-area sovereign CDS.[4]

We show that one common frailty factor and one country-specific factor for each country allow for a very good fit of CDS data. Our estimation detects space for contagion effects along with the frailty factor. Thus, both components are instrumental in explaining the dynamics of sovereign credit risk in the euro-area. We find that the frailty factor explains most of the CDS variability after mid-2012, coinciding with the ECB's *Outright Monetary Transactions* announcement. Moreover,

[4] Although Ait-Sahalia et al. (2014)'s dataset covers seven countries, they do not estimate a joint model and focus on small models involving two countries only. Besides, because they work only under the risk-neutral dynamics, they cannot examine sovereign risk premia.

we obtain sizable credit-risk premia along the whole maturity spectrum. Typically, credit-risk premia account for more than half of CDS spreads at the five year maturity for France, Germany, Italy and Spain. Both frailty and contagion channels are paramount to capture the magnitude of these premia empirically.

We also show how our framework offers the possibility to study quanto CDS spreads, which are the deviations between spreads of CDS on the same entity but denominated in different currencies. The quanto CDS written on a given defaultable entity contains information about the distribution of the exchange rate at the default time of this entity, i.e. about the expected depreciation-at-default (see e.g. Ehlers and Schonbucher 2004; Augustin et al. 2020). In order to price CDSs whose payoffs are denominated in euros or in U.S. dollars, we simply augment the model with a EURUSD exchange rate equation and allow for depreciationary effects of sovereign defaults. Let us stress that this could not be captured in standard frameworks where feedbacks from defaults to common factors are ruled out.

According to the estimated specification of the exchange rate – obtained by optimizing the fit of observed quanto CDS spreads – sovereign defaults in France, Germany, Greece, Italy and Spain would be followed by average euro depreciations of, respectively, 15%, 20%, 0%, 6% and 8%.[5] Our results further suggest that it is the fact that the exchange rate jumps upon default – and not the correlation between the exchange rate and the conditional default probability – that is key to explain the fluctuation of quanto CDS spreads.[6] These findings are consistent with the so-called "Twin Ds" phenomenon, whereby sovereign defaults are accompanied by dramatic devaluations (see Reinhart 2002; Na et al. 2018).

The remainder of the paper is organized as follows. Section 2 presents the general affine positive credit-risk modeling framework. Section 3 provides the associated explicit pricing formulas for

[5]Our exchange-rate-augmented model maps quanto CDS spreads to the expected impact of sovereign defaults on exchange rate, making it possible to back out these (expected) impacts even if defaults are not observed in the sample.

[6]This is in line with the findings of Ehlers and Schonbucher (2004) and of Brigo et al. (2015). The former paper is based on CDS data for Japanese multinational corporations, the latter exploits Italian sovereign CDS data.

defaultable bonds and CDSs. Section 4 develops the sovereign credit risk and Section 5 concludes. An online appendix provides proofs, technical results and details about the calibration of our sovereign credit-risk model.

2 A General Affine Positive Credit Risk Modeling Framework

2.1 Notations and Statistical Assumptions

We consider an economy with n defaultable entities indexed by i, firms or countries for instance. Each entity is associated with an indicator of default $d_{i,t}$, such that $d_{i,t} = 1$ if i is affected by a credit event at time t, and $d_{i,t} = 0$ otherwise. We represent the arrival of a credit event for entity i through a non-negative random process denoted by $\delta_{i,t}$ jumping upwards (see Section 2.2).

Our economy is also governed by a set of N_y common factors denoted by y_t. Under proper parameterization, some of the components of y_t may play the role of entity-specific factors (as will be illustrated by Section 4).We also use the notations $\delta_t = (\delta_{1,t}, \dots, \delta_{n,t})$ and $w_t = (y_t,\, \delta_t)$.

We denote by $\mathcal{F}_t^*$ the collection of present and past common factors $(y_t,\, y_{t-1},\, \dots)$, while $\mathcal{D}_{i,t}$ denotes the collection of all present and past entity-i credit-event variables $(\delta_{i,t}, \delta_{i,t-1}, \dots)$ and $\mathcal{D}_t = \cup_{i=1}^n \mathcal{D}_{i,t}$. That is, $\mathcal{D}_t$ is the entire history of all present and past credit-event variables. The entire information set available to investors is thus given by $\mathcal{F}_t = \mathcal{D}_t \cup \mathcal{F}_t^*$.

2.2 Default Time Modeling

Assumption 1 *The k^{th} default date $\tau_i^{(k)}$ of entity i is defined as:*

$$\tau_i^{(k)} = \inf \left\{ t > \tau_i^{(k-1)} : \{\delta_{i,t} > 0\} \cap \{\delta_{i,t-1} = 0\} \right\} ,$$

where $\tau_i^{(0)} = 0$ and $\delta_{i,t} \geqslant 0$ a.s. and is called credit-event variable.

This definition accommodates non-absorbing default states. As highlighted by Guo et al. (2009), the credit event affecting a firm triggers in reality a period of resolution that can possibly lead this entity to insolvency after the default date or to remain solvent (see also Kraft and Steffensen 2007). In the latter case, $\delta_{i,t}$ comes back to zero after it has jumped to positive values. Our assumed dynamics for the credit-event variables will feature such a mechanism.[7]

In the following, we introduce the main distributional assumptions for δ_t (Assumptions 2 and 3) and y_t (Assumptions 4 and 5).

Assumption 2 *Conditionally on $(\mathcal{F}_t^*, \mathcal{D}_{t-1})$, each credit-event variable $\delta_{i,t}$, $i \in \{1, \dots, n\}$, is independently drawn from a Gamma-zero distribution with intensity $\lambda_{i,t}^{\mathbb{P}}$. More precisely, there exists a Poisson distributed mixing variable $\mathsf{P}_{i,t}$ such that:*

$$\left(\mathsf{P}_{i,t} \,\middle|\, \mathcal{F}_t^*, \mathcal{D}_{t-1}\right) \overset{\mathbb{P}}{\sim} \mathcal{P}\left(\lambda_{i,t}^{\mathbb{P}}\right) \qquad \text{and} \qquad \left(\delta_{i,t} \,\middle|\, \mathsf{P}_{i,t}\right) \overset{\mathbb{P}}{\sim} \Gamma_{\mathsf{P}_{i,t}}\left(\mu_{\delta_i}\right) , \tag{1}$$

where $\mu_{\delta_i} > 0$ is the scaling parameter and $\mathsf{P}_{i,t}$ is the degree of freedom parameter at date t, of the Gamma distribution $\Gamma_{\mathsf{P}_{i,t}}(\mu_{\delta_i})$. The associated Gamma-zero distribution is denoted $\Gamma_0(\lambda_{i,t}^{\mathbb{P}}, \mu_{\delta_i})$.

According to Equation (1), the one-period-ahead survival probability of entity i, given $(\mathcal{F}_t^*, \mathcal{D}_{t-1})$, is the probability that the Poisson mixing variable is equal to 0, i.e. $e^{-\lambda_{i,t}^{\mathbb{P}}}$. (Indeed, when the mixing variable is equal to zero, the Gamma distribution collapses to a Dirac mass at zero.) Hence, the conditional probability of default is approximately equal to $\lambda_{i,t}^{\mathbb{P}}$ when this variable – the physical credit-event intensity – is small. We defer the specification of $\lambda_{i,t}^{\mathbb{P}}$ to the next section.

The following proposition summarizes the properties of an autonomous Autoregressive Gamma-zero process (featuring a Gamma-zero conditional distribution); more details can be found in Monfort et al. (2017), who introduce this process.

Proposition 2.1 *Let us assume that the random process (ℓ_t) is a $ARG_0(\alpha_\ell, \beta_\ell, \mu_\ell)$ process of order one. The conditional distribution of ℓ_{t+1}, given $\underline{\ell_t} = (\ell_t, \ell_{t-1}, \dots)$, is the Gamma-zero distribution:*

$$(\ell_{t+1} | \underline{\ell_t}) \sim \Gamma_0(\alpha_\ell + \beta_\ell\, \ell_t, \mu_\ell) \quad \text{for} \quad \alpha_\ell \geq 0,\, \mu_\ell > 0,\, \beta_\ell > 0\,.$$

[7] In the case of sovereign debts, Asonuma and Trebesch (2016) find that 62% of debt restructuring episodes observed between 1978 and 2010 occurred post-default with an average duration of five years. The remaining 38% of restructuring episodes are preemptive, that is with the restructuring implemented prior to a credit event. The associated average duration is of one year.

The conditional Laplace transform $\varphi_{\ell,t}(u\,;\,\alpha_\ell,\beta_\ell,\mu_\ell)$ of the $ARG_0(\alpha_\ell,\beta_\ell,\mu_\ell)$ process is given by:

$$\varphi_{\ell,t}(u\,;\,\alpha_\ell,\beta_\ell,\mu_\ell) \quad := \quad \mathbb{E}\left[\exp\left(u\,\ell_{t+1}\right) \mid \underline{\ell_t}\right] = \exp\left[\frac{u\,\mu_\ell}{1-u\,\mu_\ell}(\alpha_\ell+\beta_\ell\,\ell_t)\right], \quad \textit{for} \quad u < \frac{1}{\mu_\ell}\,.$$

Proposition 2.1 shows that the conditional Laplace transform of ℓ_{t+1} is exponential-affine in ℓ_t, thus formalizing the affine nature of the process.

In the context of the the ARG_0 process considered by Proposition 2.1, the affine property is obtained by specifying the intensity – that is also the parameter of a Poisson distribution (see Equation 1) – as an affine function of ℓ_t. In the next subsection, we consider the multivariate extension of the ARG_0 process. As in the univariate case, the affine nature of the state vector will be obtained by making the credit-event intensity $\lambda^{\mathbb{P}}_{i,t}$ linearly depend on the state variables.

2.3 Credit-Event Intensity Specification

Assumption 3 *For any entity i, the physical credit-event intensity is given by the following affine function of y_t and δ_{t-1}:*

$$\lambda^{\mathbb{P}}_{i,t} = \alpha_{\lambda_i} + \beta^{(y)\prime}_{\lambda_i} y_t + \mathbf{C}'_i\delta_{t-1}\,, \tag{2}$$

where α_{λ_i} is a scalar, $\beta^{(y)}_{\lambda_i}$ is a size-N_y vector, and $\mathbf{C}_i$ has n non-negative entries, such that $\lambda^{\mathbb{P}}_{i,t} \geqslant 0$ a.s. In the case where y_t is non-negative, $\lambda^{\mathbb{P}}_{i,t} \geqslant 0$ is guaranteed by $\alpha_{\lambda_i} \geqslant 0$, $\beta^{(y)}_{\lambda_i} \geqslant 0$ and $\mathbf{C}_i \geqslant 0$.

Depending on whether $\mathbf{C}_i = \mathbf{0}$, the intensity processes are either $\mathcal{F}^*_t$- or $\mathcal{F}^*_t \cup \mathcal{D}_{t-1}$-adapted. The case where at least one component $\mathbf{C}_{i,e}$ ($i \neq e$) is different from zero is important since it allows our framework to feature *direct contagion* between at least two entities, or mutually exciting processes. To see this, remember that a credit event happens for entity e at date $t-1$ if its credit-event variable $\delta_{e,t-1}$ jumps from zero to a strictly positive value. If $\mathbf{C}_{i,e}$ is positive, the intensity $\lambda^{\mathbb{P}}_{i,t}$ increases at date t, thus generating a higher default probability for entity i. Our model can then reproduce cross-excitation, a crucial feature in the credit-risk modeling literature (see e.g. Giesecke and Zhu 2013; Ait-Sahalia et al. 2014; Ait-Sahalia et al. 2015).[8,9]

Because of the potential persistence of y_t, the credit-event variable of entity i can remain positive for several periods after the first jump of $\delta_{i,t-1}$ before going back to zero. Through contagion effects, this may increase other credit-event intensities persistently, having a long-lasting impact on the price of their defaultable securities (see following sections).

To close the model, we have to specify the dynamics of y_t and characterize its conditional distribution. Before proposing a specific conditional distribution (in Subsection 2.4), we define the general context under which we get an affine state vector (y_t, δ_t) – thus closed-form pricing formulas:

Assumption 4 *Given $\mathcal{F}_{t-1}$, the process $\{y_t\}$ has an exponential-affine Laplace transform:*

$$\varphi^{\mathbb{P}}_{y_{t-1}}(u_y) \quad := \quad \mathbb{E}\left[\exp(u_y' y_t) \,\middle|\, \mathcal{F}_{t-1}\right] = \exp\left[A_y^{(y)}(u_y)' y_{t-1} + A_y^{(\delta)}(u_y)'\delta_{t-1} + B_y(u_y)\right],$$

where u_y is an argument of size N_y.

Assumption 4 effectively breaks down the no-jump condition whenever the image of the loading function $A_y^{(\delta)}(u_y)$ contains values different from zero. In other words, the credit events of any entity $\delta_{i,t}$ can have an impact on the common factors dynamics through Granger-causality, and the process y_t is not autonomous. We call this mechanism indirect contagion (or systemic risk) since the default of a single entity can for instance have an impact on one of the components of y_t representing the state of the economy, which will in turn feedback onto higher credit-event intensities.

Proposition 2.2 *Under Assumptions 2, 3 and 4, the stochastic process $\{w_t\}$ is affine under the historical probability measure $\mathbb{P}$. That is, the conditional Laplace transform of w_t given $\mathcal{F}_{t-1}$ is an*

[8] Compared with the model of Ait-Sahalia et al. (2014), our assumption is that any jump of $\delta_{i,t}$ automatically triggers a credit event while they assume that a jump has a certain probability to generate a default (see their Equation 5). While such a mechanism could be introduced in our formulation, we leave it aside for simplicity.

[9] Our model also allows for self-excitation – in the Hawkes sens – if $\mathbf{C}_{i,i} > 0$. Note however that a situation where $\mathbf{C}_{i,i} > 0$ for some i's but $\mathbf{C}_{i,e} = 0$ for $i \neq e$ is not consistent with contagion phenomena (unless the model features indirect contagion, see the discussion below Assumption 5).

exponential-affine function of w_{t-1}. Formally:

$$\varphi^{\mathbb{P}}_{w_{t-1}}(u_w) \quad := \quad \mathbb{E}\left[\exp(u_y' y_t + u_\delta' \delta_t) \,\middle|\, \mathcal{F}_{t-1}\right] = \exp\left[A_w^{(y)}(u_w)' y_{t-1} + A_w^{(\delta)}(u_w)' \delta_{t-1} + B_w(u_w)\right],$$

where $u_w = (u_y, u_\delta)$, and the functions $A_w^{(y)}$, $A_w^{(\delta)}$ and B_w are given by:

$$\begin{aligned}
A_w^{(y)}(u_w) &= A_y^{(y)}\left(u_y + \beta_\lambda^{(y)} \frac{u_\delta \odot \mu_\delta}{\mathbf{1} - u_\delta \odot \mu_\delta}\right), \; A_w^{(\delta)}(u_w) = A_y^{(\delta)}\left(u_y + \beta_\lambda^{(y)} \frac{u_\delta \odot \mu_\delta}{\mathbf{1} - u_\delta \odot \mu_\delta}\right) + \mathbf{C}\, \frac{u_\delta \odot \mu_\delta}{\mathbf{1} - u_\delta \odot \mu_\delta}, \\
B_w(u_w) &= B_y\left(u_y + \beta_\lambda^{(y)} \frac{u_\delta \odot \mu_\delta}{\mathbf{1} - u_\delta \odot \mu_\delta}\right) + \alpha_\lambda' \frac{u_\delta \odot \mu_\delta}{\mathbf{1} - u_\delta \odot \mu_\delta},
\end{aligned}$$

where $\beta_\lambda^{(y)}$ is the $(N_y \times n)$ matrix whose columns are $\beta_{\lambda_i}^{(y)}$, $\mathbf{C}$ is the $(n \times n)$ matrix whose columns are $\mathbf{C}_i$, and α_λ is the vector of all α_{λ_i}. The operator $\odot$ is the (Hadamard) element-by-element product and the $\frac{\cdot}{\cdot}$ operator is taken element-by-element.

Proof See Online Appendix **??**. ■

2.4 Proposed Affine Vector Autoregressive Gamma Dynamics

Assumption 4 leaves some freedom about the specification of the dynamics of y_t. Assumption 5 describes a particularly convenient choice that we develop further in the empirical application (see Chen and Filipovic, 2007, for a continuous-time approach).

Assumption 5 *The factors $y_{j,t}$, $j \in \{1, \dots, N_y\}$, have non-central Gamma dynamics given by:*

$$\left(\mathsf{P}_{y_j,t} \,\middle|\, \mathcal{F}_{t-1}\right) \overset{\mathbb{P}}{\sim} \mathcal{P}\left(\beta_{y_j}^{(y)\prime} y_{t-1} + \mathbf{I}_j' \delta_{t-1}\right) \qquad \text{and} \qquad \left(y_{j,t} \,\middle|\, \mathsf{P}_{y_j,t}\right) \overset{\mathbb{P}}{\sim} \Gamma_{\nu_{y_j} + \mathsf{P}_{y_j,t}}\left(\mu_{y_j}\right). \tag{3}$$

All parameters of Equation (3) are non-negative and of adapted dimension, and ν_{y_j} is the deterministic component of the degree of freedom. In the case where $\nu_{y_j} > 0$, $y_{j,t}$ is strictly positve a.s. Moreover, conditionally on $\mathcal{F}_{t-1}$, the scalar components of y_t are independent.

Two relevant characteristics stand out from this specification. First, the vectors of parameters $\mathbf{I}_j$ represent the transmission channel of what we have called systemic risk in the previous subsection. If entity e defaults at $t-1$, its credit-event variable $\delta_{e,t-1}$ jumps to a positive value, which increases the conditional mean of the common factor $y_{j,t}$ as long as Granger-causality is allowed by $\mathbf{I}_{j,e} > 0$. Parameters $\mathbf{I}_{j,e} > 0$ *and* $\beta^{(y)}_{\lambda_i}$ then open the way to an *indirect contagion* channel from the systemic entity e (featuring $\delta_{e,t-1} > 0$) to another entity i via the common factor $y_{j,t}$. (The notation $\mathbf{I}_j$ makes reference to the indirect nature of the contagion channel; in the same way, $\mathbf{C}_i$ was referring to direct contagion.) This feature results from the recursive specification of the two sets of variables: while Equation (2) implies that y_t instantaneously causes δ_t, Equation (3) shows that y_t depends on δ_{t-1}.

Under Assumption 5, risk factors y_t are non-negative. This may represent a problem if one wants to include observable risk factors with negative values. In some cases, this problem may be circumvented by employing non-negative transformations of the considered variables.

The state process $w_t = (y_t, \delta_t)$ resulting from Assumptions 2, 3 and 5 is called recursive Vector Autoregressive Gamma (VARG) process. The following proposition complements Proposition 2.2 in the VARG context. Specifically, it gives the forms of functions functions $A^{(y)}_y$, $A^{(\delta)}_y$ and B_y resulting from Assumption 5.

Proposition 2.3 *Under Assumptions 2, 3 and 5, the Laplace transform of w_t, conditional on $\mathcal{F}_{t-1}$, is exponential-affine in w_{t-1} and given by Proposition 2.2, with functions $A^{(y)}_y$, $A^{(\delta)}_y$ and B_y given by:*

$$A^{(y)}_y(u_y) \quad = \quad \beta^{(y)}_y \frac{u_y \odot \mu_y}{\mathbf{1} - u_y \odot \mu_y}\,,\ A^{(\delta)}_y(u_y) = \mathbf{I}\,\frac{u_y \odot \mu_y}{\mathbf{1} - u_y \odot \mu_y}\,,\ B_y(u_y) = -\nu'_y \log\left[\mathbf{1} - u_y \odot \mu_y\right]\,,$$

where $\beta^{(y)}_y$ is the $(N_y \times N_y)$ matrix whose columns are $\beta^{(y)}_{y_j}$, $\mathbf{I}$ is the $(n \times N_y)$ matrix whose columns are $\mathbf{I}_j$, and μ_y and ν_y are the vectors stacking together the individual elements with the same notations. The operator $\odot$ is the Hadamard element-by-element product, and the $\log(\cdot)$ and $\div$ operators are taken element-by-element.

Proof See Online Appendix **??**. ■

As for standard affine processes, our assumed dynamics for w_t has convenient properties in terms of conditional cumulants, stationarity conditions and predictions. These properties directly derive from the semi-strong VAR representation of the state-vector dynamics (see Online Appendix **??**).[10]

2.5 The Stochastic Discount Factor and Credit-Risk Pricing

We assume the existence of a representative investor who prices all assets in the economy such that no-arbitrage holds.[11] We focus on the following financial instruments: risk-free debt, debt issued by all defaultable entities i, and credit-risk derivatives on all entities.

The following assumption pertains to the short-term risk-free rate and the SDF.

Assumption 6 *The risk-free one-period yield between $t-1$ and t, denoted by r_{t-1}, is given by an affine function of the factors:*

$$r_{t-1} = \xi_0 + \xi_y' y_{t-1} + \xi_\delta' \delta_{t-1}\,, \tag{4}$$

where ξ_0, ξ_y, and ξ_δ are respectively a scalar, a N_y-dimensional vector and a n-dimensional vector. The one-period SDF is denoted $M_{t-1,t}$ and given by:

$$M_{t-1,t} = \exp\left(-r_{t-1} + \theta_y' y_t + \mathbf{S}'\delta_t - \log\left[\varphi^{\mathbb{P}}_{w_{t-1}}(\theta_y,\, \mathbf{S})\right]\right)\,, \tag{5}$$

where $(\theta_y,\, \mathbf{S})$ are the risk-correction parameters, or "prices of risk", and the last term in (5) guarantees to satisfy the no-arbitrage condition $\mathbb{E}\left(M_{t-1,t}|\mathcal{F}_{t-1}\right) = e^{-r_{t-1}}$.

Although our SDF formulation is primarily motivated by computational reasons, its exponential-affine formulation can stem from structural models where agents feature CRRA or Epstein-Zin preferences and where consumption growth is affine in w_t.[12] In particular, these structural ap-

[10] In particular, once the semi-strong VAR representation is known, the state vector is strictly stationary iff the eigenvalues of the auto-regressive matrix (denoted by M_1 in Online Appendix **??**) are strictly lower than one in modulus.

[11] In the discrete-time context, it can be shown that under the assumptions of (a) existence of uniqueness of a price, (b) price linearity and continuity and (c) absence of arbitrage opportunity, there exists a unique positive SDF. This derives from a conditional version of the Riesz representation theorem (see e.g. Hansen and Richard 1987).

[12] If the representative agent features time-separable CRRA preferences and if consumption growth between dates $t-1$ and t is denoted by Δc_t, then it is easily shown that the SDF $M_{t-1,t}$ is proportional to $\exp(-\gamma\Delta c_t)$, where γ is the relative risk aversion parameter. Eraker 2008 proposes an approach to solve for an approximated exponential affine SDF when the representative agent features Epstein-Zin preferences and when Δc_t linearly depends on an affine process. Bai et al. (2015) also obtain a pricing kernel that depends, in an exponential affine way, on credit-event variables in the context of a production economy populated by firms (see their Equations 25).

proaches imply that $\mathbf{S}_i > 0$ if the credit event associated with entity i coincide with a drop in consumption, which is for instance consistent with empirical evidence on the effect of sovereign defaults (see e.g. Reinhart and Rogoff 2011; Mendoza and Yue 2012; Trebesch and Zabel 2017).[13]

The price of risk θ_y, which is standard in the context of credit-risk modeling, drives a wedge between the physical and risk-neutral moments of the factors driving the credit-event intensities. The vector of parameters $\mathbf{S}$ allows us to relax the widely-used assumption according to which credit events do not enter the representative investor's SDF – and are therefore not a priced source of risk. In line with Gourieroux et al. (2014), we call this mechanism credit-event pricing (or surprise pricing, hence the notation $\mathbf{S}$).

Now that the physical dynamics and the SDF are known, we can price assets whose payoffs depend on future values of the state vector. Specifically, the date-t price of an asset providing the payoff P_{t+1} on date $t+1$ is given by $\mathbb{E}(M_{t,t+1}P_{t+1}|\mathcal{F}_t)$. This price can also be written $\mathbb{E}^{\mathbb{Q}}(e^{-r_t}\,P_{t+1}|\mathcal{F}_t)$, where the one-period-ahead change of measure from $\mathbb{P}$ (the physical measure) to $\mathbb{Q}$ (the risk-neutral measure) is given by $M_{t,t+1}/\mathbb{E}(M_{t,t+1}|\mathcal{F}_t)$. Determining the risk-neutral measure often facilitates pricing. We explore the risk-neutral dynamics resulting from our assumptions below.

Proposition 2.4 *Under Assumptions 2, 3, 4 (or 5) and 6, the stochastic process $\{w_t\}$ has an exponential-affine conditional Laplace transform given $\mathcal{F}_{t-1}$ under the risk-neutral measure:*

$$\varphi^{\mathbb{Q}}_{w_{t-1}}(u_w) \quad := \quad \mathbb{E}^{\mathbb{Q}}\left[\exp(u_y'y_t + u_\delta'\delta_t)\,\middle|\,\mathcal{F}_{t-1}\right] = \exp\left[A^{\mathbb{Q}(y)}_w(u_w)'y_{t-1} + A^{\mathbb{Q}(\delta)}_w(u_w)'\delta_{t-1} + B^{\mathbb{Q}}_w(u_w)\right],$$

where the loadings functions are given by:

$$\begin{aligned} A^{\mathbb{Q}(\ell)}_w(u_w) &= A^{(\ell)}_w\big(u_y+\theta_y,\, u_\delta+\mathbf{S}\big) - A^{(\ell)}_w\big(\theta_y,\, \mathbf{S}\big)\,,\ \ \ell\in\{y,\delta\}\,,\\ B^{\mathbb{Q}}_w(u_w) &= B_w\big(u_y+\theta_y,\, u_\delta+\mathbf{S}\big) - B_w\big(\theta_y,\, \mathbf{S}\big)\,, \end{aligned}$$

[13] The fact that consumption drops coincide with sovereign defaults is also consistent with the disaster-risk literature, where sovereign default can be triggered by exogenous disasters having dramatic recessionary effects (see e.g. Barro 2006; Gabaix 2012).

where $A^{(\ell)}_w(u_w)$, $\ell = \{y, \delta\}$, *and* $B_w(u_w)$ *are defined in Proposition 2.2.*

Proof Straightforward application of the Esscher transform. ■

Proposition 2.4 underlines the wide use of affine processes in asset-pricing models: the exponential-affine specification of the SDF preserves the affine property when we move from the physical to the risk-neutral measure. This leads to closed-form pricing formulas not only for securities paying off an exponential-affine function of w_t, but also featuring payoffs of the form $\mathbb{1}_{\{\delta_{i,t}>0\}}$ (see Section 3).

While Proposition 2.4 perfectly characterizes the risk-neutral dynamics of w_t, it does not make the risk-neutral conditional distributions of δ_t and of y_t explicit. Remarkably, the risk-neutral conditional distributions of δ_t and of y_t remain of the same type as their physical counterpart. This is formalized in the next two propositions, which echo the assumptions made on the (physical) conditional distributions of δ_t and of y_t, respectively by Assumptions 2 and 5.

Proposition 2.5 *Under Assumptions 2, 3, 4 (or 5) and 6, and conditionally on $\mathcal{F}^*_t$, $\mathcal{D}_{t-1}$, the credit-event variables $\delta_{i,t}$ are Gamma-zero distributed under the risk-neutral probability measure $\mathbb{Q}$. In particular, there exists a risk-neutral credit-event intensity process $\lambda^{\mathbb{Q}}_{i,t}$ adapted to $\mathcal{F}^*_t$, $\mathcal{D}_{t-1}$, such that:*

$$\left(\mathsf{P}_{i,t}\,\middle|\,\mathcal{F}^*_t,\mathcal{D}_{t-1}\right) \overset{\mathbb{Q}}{\sim} \mathcal{P}\left(\lambda^{\mathbb{Q}}_{i,t}\right) \qquad \text{and} \qquad \left(\delta_{i,t}\,\middle|\,\mathsf{P}_{i,t}\right) \overset{\mathbb{Q}}{\sim} \Gamma_{\mathsf{P}_{i,t}}\left(\mu^{\mathbb{Q}}_{\delta_i}\right), \tag{6}$$

where $\mu^{\mathbb{Q}}_{\delta_i} = \dfrac{\mu_{\delta_i}}{1-\mathbf{S}_i\mu_{\delta_i}}$ *and the risk-neutral credit-event intensity is given by:*

$$\lambda^{\mathbb{Q}}_{i,t} = \alpha^{\mathbb{Q}}_{\lambda_i} + \beta^{\mathbb{Q}(y)'}_{\lambda_i} y_t + \mathbf{C}^{\mathbb{Q}'}_i \delta_{t-1} = \frac{\lambda^{\mathbb{P}}_{i,t}}{1-\mathbf{S}_i\mu_{\delta_i}}, \tag{7}$$

with $\alpha^{\mathbb{Q}}_{\lambda_i} = \frac{1}{1-\mathbf{S}_i\mu_{\delta_i}}\alpha_{\lambda_i}$, $\beta^{\mathbb{Q}(y)}_{\lambda_i} = \frac{1}{1-\mathbf{S}_i\mu_{\delta_i}}\beta^{(y)}_{\lambda_i}$, *and* $\mathbf{C}^{\mathbb{Q}}_i = \frac{1}{1-\mathbf{S}_i\mu_{\delta_i}}\mathbf{C}_i$.

Proof Straightforward application of the Esscher transform to the distribution associated with the conditional Laplace transform $\varphi_{\delta_{i,t}}(u;\,\lambda^{\mathbb{P}}_{i,t},\,\mu_{\delta_i}) = \mathbb{E}\left[\exp(u\delta_{i,t})\middle|\mathcal{F}^*_t,\mathcal{D}_{t-1}\right]$. ■

Proposition 2.5 emphasizes an important property for the risk-neutral credit-event intensities. For standard credit-risk models, credit-event prices of risk are null ($\mathbf{S} = 0$), implying that physical and risk-neutral intensities are identical functions of all risk factors (see Equation 7). In other words, $\lambda^{\mathbb{Q}}_{i,t} = \lambda^{\mathbb{P}}_{i,t}$ if credit events are not priced by the representative investor. In that case, credit-risk premia arise only because the conditional moments of the risk factors w_t are different when taken under the risk-neutral measure or under the historical one. Instead, when the credit-event risk is priced ($\mathbf{S}_i > 0$), the risk-neutral intensities become proportional to the physical ones, i.e. $\lambda^{\mathbb{Q}}_{i,t} = \frac{\lambda^{\mathbb{P}}_{i,t}}{1-\mathbf{S}_i \mu_{\delta_i}}$, as in Jarrow et al. (2005) and Driessen (2005) (see also Duffie, 2005, and references therein). This creates an additional wedge between the physical and risk-neutral moments of the credit-event variables $\delta_{i,t}$.

Proposition 2.6 *Assume that the historical dynamics of w_t is described by the recursive VARG process of Assumptions 2, 3 and 5, and the SDF specification of Assumption 6. Under the risk-neutral measure, conditionally on $\mathcal{F}_{t-1}$, the components of y_t are independent and follow a non-central Gamma distribution (as under the physical measure). More precisely, for $y_{j,t}$, $j \in \{1, \dots, N_y\}$:*

$$\left(\mathsf{P}_{y_j,t} \,\middle|\, \mathcal{F}_{t-1}\right) \overset{\mathbb{Q}}{\sim} \mathcal{P}\left(\beta^{\mathbb{Q}(y)'}_{y_j} y_{t-1} + \mathbf{I}^{\mathbb{Q}'}_{j} \delta_{t-1}\right) \qquad \text{and} \qquad \left(y_{j,t} \,\middle|\, \mathsf{P}_{y_j,t}\right) \overset{\mathbb{Q}}{\sim} \Gamma_{\nu_{y_j} + \mathsf{P}_{y_j,t}}\left(\mu^{\mathbb{Q}}_{y_j}\right), \tag{8}$$

where the risk-neutral parameters are given by:

$$\beta^{\mathbb{Q}(y)}_{y_j} = \beta^{(y)}_{y_j} \frac{1}{1 - \mu_{y_j} \widetilde{\theta}_{y_j}}, \; \mu^{\mathbb{Q}}_{y_j} = \frac{\mu_{y_j}}{1 - \mu_{y_j} \widetilde{\theta}_{y_j}}, \quad \text{and} \quad \mathbf{I}^{\mathbb{Q}}_{j} = \mathbf{I}_j \frac{1}{1 - \mu_{y_j} \widetilde{\theta}_{y_j}},$$

with

$$\widetilde{\theta}_y = \theta_y + \beta^{(y)}_{\lambda} \frac{\mathbf{S} \odot \mu_\delta}{\mathbf{1} - \mathbf{S} \odot \mu_\delta}. \tag{9}$$

Proof Straightforward application of Propositions 2.5 and 2.4.

Propositions 2.5 and 2.6 show that the risk-neutral distribution of the state vector w_t thus remains recursive VARG. The recursive structure also underlines another important effect on the prices of risk applied to the common factors. Because y_t instantaneously feeds back on the credit-event variables δ_t, it is associated with two sources of priced risk in the modified risk-adjustment parameter $\widetilde{\theta}_y$. The representative investor can dislike upwards movements of y_t itself, for instance because it represents worsened economic conditions; this is represented by θ_y. The second term of Equation 9 shows that the investor also prices movements in y_t because of the instantaneous impact of the common factors y_t on the default intensities via Equation (2). It is thus natural to see $\mathbf{S}$, the prices of risk associated with credit-event variables, appearing in the pricing of the common factors y_t.

2.6 The Reverse-Order Multi-Horizon Laplace Transform

Before exploring the pricing properties of the framework introduced in the previous sections, we explain hereafter why and how multi-horizon Laplace transforms can be computed in a fast way when their arguments feature a so-called "reverse-order structure". This approach allows for an efficient computation of all default-sensitive asset prices we consider, i.e. any security whose cashflows can be expressed as exponential-affine combinations of w_t, or featuring terms of the form $\mathbb{1}_{\{\delta_{i,t}>0\}}$. We first introduce the general definition of the multi-horizon Laplace transform and we then discuss the particular reverse-order case.

Proposition 2.7 *Let us consider a horizon $h \in \mathbb{N}$ and a set of arguments $(u_1, \ldots, u_h)$, where each vector u_j is of dimension $N_y + n$. The h-period-ahead risk-neutral multi-horizon Laplace transform of the affine process $\{w_t\}$, for arguments $(u_1, \ldots, u_h)$, is given by:*

$$\begin{aligned}\varphi^{\mathbb{Q}}_{w_t}(u_1, \ldots, u_h) &:= \mathbb{E}^{\mathbb{Q}}\left[\exp\left(\sum_{j=1}^{h} u_j' w_{t+j}\right) \Big| \mathcal{F}_t\right] \\ &= \exp\Big(\mathcal{A}_h(u_1, \ldots, u_h)' w_t + \mathcal{B}_h(u_1, \ldots, u_h)\Big),\end{aligned} \tag{10}$$

where the loadings functions $\mathcal{A}_h$ and $\mathcal{B}_h$ are defined for any arguments $(v_1, \ldots, v_h)$ through the following recursive system (for $h > 0$), initialized with $\mathcal{A}_0 = \mathbf{0}$ and $\mathcal{B}_0 = 0$:

$$\left\{\begin{aligned}\mathcal{A}_h(v_1, \ldots, v_h) &= A^{\mathbb{Q}}_w(v_1 + \mathcal{A}_{h-1}(v_2, \ldots, v_h)) \\ \mathcal{B}_h(v_1, \ldots, v_h) &= B^{\mathbb{Q}}_w(v_1 + \mathcal{A}_{h-1}(v_2, \ldots, v_h)) + \mathcal{B}_{h-1}(v_2, \ldots, v_h),\end{aligned}\right. \tag{11}$$

Proof See Online Appendix **??**. ■

For two sets of arguments $U_{h_1} = (u_1, \ldots, u_{h_1})$ and $U_{h_2} = (u_1, \ldots, u_{h_2})$, with $h_1 < h_2$, the respective h_1-step and h_2-step recursions (11) have to be run separately. More precisely, the multiple calls to functions $A_w^{\mathbb{Q}}$ and $B_w^{\mathbb{Q}}$ (Equation 11), that are necessary to evaluate the multi-horizon Laplace transform $\varphi_{w_t}^{\mathbb{Q}}(u_1, \ldots, u_{h_1})$, are of no use to compute $\varphi_{w_t}^{\mathbb{Q}}(u_1, \ldots, u_{h_2})$ even if the first h_1 vectors of U_{h_1} and U_{h_2} are the same, because the arguments to be called in $A_w^{\mathbb{Q}}$ and $B_w^{\mathbb{Q}}$ are never the same. For example, while the computation of $\varphi_{w_t}^{\mathbb{Q}}(u_1)$ involves $A_w^{\mathbb{Q}}(u_1)$ and $B_w^{\mathbb{Q}}(u_1)$, the computation of $\varphi_{w_t}^{\mathbb{Q}}(u_1, u_2)$ involves $A_w^{\mathbb{Q}}(u_2)$, $B_w^{\mathbb{Q}}(u_2)$, $A_w^{\mathbb{Q}}(u_1 + A_w^{\mathbb{Q}}(u_2))$ and $B_w^{\mathbb{Q}}(u_1 + A_w^{\mathbb{Q}}(u_2))$. If $u_1 \neq u_2$, there is no overlap in the computations.

From a computational point of view, the situation is more favorable when the different sets of considered arguments are (also) nested, but organized in "reverse order"; that is when we want to compute $\varphi_{w_t}^{\mathbb{Q}}(u_h, \ldots, u_1)$, with h growing. Indeed, simply changing the order of the arguments of $\mathcal{A}_h$ and $\mathcal{B}_h$ in Equation (11), we see that:

$$\left\{ \begin{array}{lcl} \mathcal{A}_h(u_h, \ldots, u_1) & = & A_w^{\mathbb{Q}}(u_h + \mathcal{A}_{h-1}(u_{h-1}, \ldots, u_1)) \\ \mathcal{B}_h(u_h, \ldots, u_1) & = & B_w^{\mathbb{Q}}(u_h + \mathcal{A}_{h-1}(u_{h-1}, \ldots, u_1)) + \mathcal{B}_{h-1}(u_{h-1}, \ldots, u_1), \end{array} \right. \tag{12}$$

which shows that $\mathcal{A}_h(u_h, \ldots, u_1)$ and $\mathcal{B}_h(u_h, \ldots, u_1)$ – which determine $\varphi_{w_t}^{\mathbb{Q}}(u_h, \ldots, u_1)$ – are directly deduced from $\mathcal{A}_{h-1}(u_{h-1}, \ldots, u_1)$ and $\mathcal{B}_{h-1}(u_{h-1}, \ldots, u_1)$ – which themselves determine $\varphi_{w_t}^{\mathbb{Q}}(u_{h-1}, \ldots, u_1)$. Looking back at the example of the previous paragraph, the computation of $\varphi_{w_t}^{\mathbb{Q}}(u_1)$ and $\varphi_{w_t}^{\mathbb{Q}}(u_2, u_1)$ both involve $A_w^{\mathbb{Q}}(u_1)$ and $B_w^{\mathbb{Q}}(u_1)$ as a first step.[14]

The following subsections make an intensive use of conditional multi-horizon Laplace transforms applied on arguments satisfying the reverse-order property. Moreover, the involved reverse-order structure is systematically characterized by only two vectors u and v (say), i.e. $\varphi_{w_t}^{\mathbb{Q}}(u, \ldots, u, v)$.

[14] In other words, the computations of $\varphi_{w_t}^{\mathbb{Q}}(u_1)$, $\varphi_{w_t}^{\mathbb{Q}}(u_2, u_1)$, $\ldots$, $\varphi_{w_t}^{\mathbb{Q}}(u_h, \ldots, u_1)$ involve h calls of functions $A_w^{\mathbb{Q}}$ and $B_w^{\mathbb{Q}}$, against $h(h+1)/2$ calls when the arguments do not feature the reverse-order structure, that is for instance when we have to compute $\varphi_{w_t}^{\mathbb{Q}}(u_1)$, $\varphi_{w_t}^{\mathbb{Q}}(u_1, u_2)$, $\ldots$, $\varphi_{w_t}^{\mathbb{Q}}(u_1, \ldots, u_h)$.

For ease of presentation, we will adopt the following notation:

$$\varphi^{\mathbb{Q}}_{w_t(h)}(u,\, v) = \varphi^{\mathbb{Q}}_{w_t}(u,\, \ldots,\, u,\, v) \tag{13}$$

and, by abuse of notation, we will replace $\mathcal{A}_h(u, \ldots, u, v)$ and $\mathcal{B}_h(u, \ldots, u, v)$ by $\mathcal{A}_h(u, v)$ and $\mathcal{B}_h(u, v)$, respectively.

3 Defaultable Asset Pricing

3.1 Recovery Rate and Recovery Conventions

We denote by $B^*(t, h)$ the price of a default-free bond of residual maturity h at time t, and by $B_i(t, h)$ the price of a defaultable bond issued by entity i. By no-arbitrage, we have:

$$B^*(t,h) = \mathbb{E}^{\mathbb{Q}}\left[e^{-r_t-\ldots-r_{t+h-1}} \,\middle|\, \mathcal{F}_t\right] .$$

As is standard in affine models, this price is a closed-form exponential-affine function of w_t.

Proposition 3.1 *The price of the risk-free bond of any residual maturity h can be computed explicitly through the multi-horizon Laplace transform recursions as:*

$$\begin{aligned} B^*(t,h) &= \exp\left[-r_t - (h-1)\xi_0\right] \times \varphi^{\mathbb{Q}}_{w_t(h-1)}(-\xi, -\xi) \\ &= \exp\left\{-h\xi_0 + \left[\mathcal{A}_{h-1}(-\xi,\, -\xi) - \xi\right]' w_t + \mathcal{B}_{h-1}(-\xi,\, -\xi)\right\} , \end{aligned}$$

where $\xi = (\xi_y,\, \xi_\delta)$ is defined in Equation (4) and where functions $\mathcal{A}_{h-1}$ and $\mathcal{B}_{h-1}$ can be evaluated using system (12).

We now consider the case of defaultable bonds and, for ease of notation, $\tau_i^{(k)}$ defined in Assumption 1 will be denoted τ_i. We assume that when entity i defaults, all its outstanding bonds are terminated and provide a (possibly stochastic) recovery payment for each unit of face value. The following definition summarizes the potential recovery assumptions commonly adopted in the literature (see Brennan and Schwartz 1980; Duffie 1998; Duffie and Singleton 1999).

Definition 3.1 *The recovery payment for bond $B_i(\tau_i, h)$ in case of default of entity $i \in \{1, \ldots, n\}$ at time τ_i is given by the product of a recovery rate and a recovery payment, denoted by ϱ_{i,τ_i} and $\Pi_{i,\tau_i}(h)$, respectively. The following assumptions on the recovery payment are commonly made:*

- *RMV: Recovery of Market Value, $\Pi_{i,\tau_i}(h) = \widetilde{B}_i(\tau_i, h)$,*
- *RFV: Recovery of Face Value, $\Pi_{i,\tau_i}(h) = 1$,*
- *RT: Recovery of Treasury, $\Pi_{i,\tau_i}(h) = B^*(\tau_i, h)$,*

where $\widetilde{B}_i(\tau_i, h)$ would be the price of the defaultable bond of entity i at time τ_i if there had been no credit event.

Assumption 7 *The recovery rate is given by an exponential-affine function of the state process w_t, written as:*

$$\varrho_{i,\tau_i} = \exp\left(-\omega_{i,0} - \omega_i^{(y)\prime}\, y_{\tau_i} - \omega_i^{(\delta)}\, \delta_{i,\tau_i}\right) , \tag{14}$$

or, more compactly, $\varrho_{i,\tau_i} = \exp\left(-\omega_{i,0} - \omega_i^{(w)\prime}\, w_{\tau_i}\right)$, where all parameters are of adapted size.

Whenever we assume that the entire dynamics of w_t are defined by Assumptions 2, 3 and 5, it is easy to impose that the recovery rate is bounded between 0 and 1 by forcing all parameters in Equation (14) to be non-negative. This equation specifies a stochastic recovery rate whose time-varying magnitude may depend on common and entity-specific factors. An interesting particular case arises when we simply assume $\omega_{i,0} = 0$, $\omega_i^{(y)} = \mathbf{0}$ and $\omega_i^{(\delta)} = 1$:

$$\varrho_{i,\tau_i} = \exp\left(-\ \delta_{i,\tau_i}\right) . \tag{15}$$

This specification delivers a clean interpretation of the process $\delta_{i,t}$. Relation (15) indeed formalizes the idea of a stochastic recovery rate equal to one as long as $\delta_{i,t} = 0$, and leaving the unitary upper bound at the default time τ_i with a reduction whose magnitude depends on the size of the credit-event variable on the default date (δ_{i,τ_i}). The size of the jump in $\delta_{i,t}$ represents the "severity" of the credit event. It jointly determines the effective recovery rate on bonds, the systemic impact on common factors (through $\mathbf{I}$ parameters) and the increase in other entities credit-event intensities λ_t through direct contagion effects (parameters $\mathbf{C}$).

We focus hereafter on the RMV and RFV conventions. The RT case is presented in the online Appendix **??**. We consider the recovery rate specification (15) for the RMV convention. The RFV case is treated under the more general specification (14).

3.2 Defaultable Bond Pricing

We first establish a general expression for the price of a defaultable bond – without specifying the recovery convention. Assume that entity i has not defaulted at time t. It is useful to rewrite the default indicator formalizing a default event happening at time $t+k$ as:

$$\mathbb{1}_{\{\delta_{i,t:t+k-1}=\mathbf{0}\}} \times \mathbb{1}_{\{\delta_{i,t+k}>0\}} = \mathbb{1}_{\{\delta_{i,t:t+k-1}=\mathbf{0}\}} - \mathbb{1}_{\{\delta_{i,t:t+k}=\mathbf{0}\}}. \tag{16}$$

The bond trading at price $B_i(t,h)$ on date t provides its holder with a single payoff between dates $t+1$ and $t+h$: this payoff is either $\varrho_{i,t+k}\,\Pi_{i,t+k}(h-k)$ (settled on date $t+k$) if the credit event happens at time $t+k$ ($k \leqslant h$) or 1 (settled on date $t+h$) if the default does not happen during the life of the bond ($t+h<\tau_i$). Accordingly, the price of the bond has to satisfy:

$$\begin{aligned} B_i(t,h) &= \sum_{k=1}^{h} \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\sum_{\ell=0}^{k-1} r_{t+\ell}\right)\left[\mathbb{1}_{\{\delta_{i,t:t+k-1}=\mathbf{0}\}} - \mathbb{1}_{\{\delta_{i,t:t+k}=\mathbf{0}\}}\right]\varrho_{i,t+k}\Pi_{i,t+k}(h-k)\,\Big|\,\mathcal{F}_t\right] \\ &+ \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\sum_{\ell=0}^{h-1} r_{t+\ell}\right)\mathbb{1}_{\{\delta_{i,t:t+h}=\mathbf{0}\}}\,\Big|\,\mathcal{F}_t\right]. \end{aligned} \tag{17}$$

Looking at (17), it is not obvious that the conditional expectations will be obtained in closed-form, nor is it clear that multi-horizon Laplace transforms will be useful to compute such quantities. The following Lemma will prove to be crucial for the derivation of closed-form pricing formula[15]

Lemma 3.1 *Let Z_1 be a random variable valued in $\mathbb{R}^d$ ($d \geq 1$) and Z_2 be a random variable valued*

[15]This lemma can be seen as a generalization of Lemma 2.1 in Monfort et al. (2017) (see also Chen and Filipovic 2007).

in $\mathbb{R}^+ = [0, +\infty)$. *Suppose that* $\mathbb{E}[\exp(u_1' Z_1 - u_2 Z_2)]$ *exists for a given* u_1 *and* $u_2 \geq 0$. *Then, we have:*

$$\mathbb{E}\left[\exp(u_1' Z_1)\, \mathbb{1}_{\{Z_2=0\}}\right] = \lim_{u_2 \to +\infty} \mathbb{E}\left[\exp(u_1' Z_1 - u_2\, Z_2)\right] .$$

Proof See Online Appendix **??**. ■

Replacing Z_1 by sets of future w_t and Z_2 by sums of $\delta_{t+\ell+1}$, we obtain some of the expectations appearing in Equation (17) as limits of the multi-horizon Laplace transform (whose computation benefits from a reverse order structure, see Subsection 2.6). Depending on the recovery assumption, the pricing formulas can simplify further and are summarized by the following propositions.

Proposition 3.2 *Under the RMV assumption (see Definition 3.1), assuming that the recovery rate is given by Equation (15), then the no-arbitrage price of the defaultable bond satisfies:*

$$B_i(t,h) = \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\sum_{\ell=0}^{h-1}(r_{t+\ell} + \delta_{i,t+\ell+1})\right) \Big| \mathcal{F}_t\right], \tag{18}$$

and can be computed as follows:

$$\begin{aligned} B_i(t,h) &= \exp\left[-r_t - (h-1)\xi_0\right] \times \varphi^{\mathbb{Q}}_{w_t(h)}\left(-\xi - \mathrm{e}_{\delta_i}, -\mathrm{e}_{\delta_i}\right) \\ &= \exp\left\{-h\xi_0 + \left[\mathcal{A}_h\left(-\xi - \mathrm{e}_{\delta_i}, -\mathrm{e}_{\delta_i}\right) - \xi\right]' w_t + \mathcal{B}_h\left(-\xi - \mathrm{e}_{\delta_i}, -\mathrm{e}_{\delta_i}\right)\right\} , \end{aligned} \tag{19}$$

where the vector e_{δ_i} *is a selection vector such that* $\mathrm{e}_{\delta_i}' w_t = \delta_{i,t}$ *and where functions* $\mathcal{A}_h$ *and* $\mathcal{B}_h$ *can be evaluated using system (11).*

Proof See online Appendix **??**. ■

Equation (18) is a key result of the paper. It shows that our framework still leads to the familiar no-arbitrage bond pricing formula based on the default-adjusted short rate $(r_{t+\ell} + \delta_{i,t+\ell+1})$, in spite of the fact that the credit events are sources of risk that are priced. This result can be seen as a discrete-time generalization of the RMV setting proposed by Duffie and Singleton (1999).

In the case where no default is observed throughout a sample, $\delta_{i,t}$ is uniformly equal to zero for all observed dates. However, the price of defaultable bonds do not collapse to that of risk-free bonds as long as the distribution of future $\delta_{i,t}$'s, conditional on $\mathcal{F}_t$, is not concentrated at zero.

Proposition 3.3 *Under the RFV assumption, assuming the general recovery rate of Equation (14), the price of the defaultable bond is given by the following sum of exponentially-affine functions:*

$$B_i(t,h) = \lim_{u\to+\infty} e^{-r_t}\left[\sum_{k=1}^{h} e^{-\omega_{i,0}-ue'_{\delta_i}w_t-(k-1)\xi_0}\left(\varphi^{\mathbb{Q}}_{w_t(k)}\left[-\xi - ue_{\delta_i},\, -\omega_i^{(w)}\right] - \varphi^{\mathbb{Q}}_{w_t(k)}\left[-\xi - ue_{\delta_i},\, -ue_{\delta_i} - \omega_i^{(w)}\right]\right)\right.$$
$$\left. + \; e^{-(h-1)\xi_0 - ue'_{\delta_i}w_t}\varphi^{\mathbb{Q}}_{w_t(h)}\left[-\xi - ue_{\delta_i},\, -ue_{\delta_i}\right]\right], \tag{20}$$

where $\varphi^{\mathbb{Q}}_{w_t(k)}(u,v) = \exp\{\mathcal{A}_k(u,v)'w_t + \mathcal{B}_k(u,v)\}$ is the conditional multi-Laplace transform that can be evaluated using system (12).

Proof See online Appendix **??**. ■

In the following section, we will use the pricing formula (20) to obtain the closed-form CDS pricing result. For the sake of precision, we will use the notation $B_i^{\text{RFV}}\left(t, h; \omega_{i,0}, \omega_i^{(w)}\right)$ to refer to this defaultable bond price under RFV with the recovery rate defined as in Equation (14).

3.3 Credit Default Swap (CDS) Valuation

Let us now consider the problem of CDS pricing. A CDS is a derivative contract where a protection buyer accepts to regularly pay a fixed rate called CDS premium (or spread) to a protection seller as long as the underlying entity does not suffer a credit event. In case of default, the contract terminates and the protection seller provides the loss-given-default on a reference bond to the protection buyer.[16] We denote by $\mathcal{S}_i(t,h)$ this CDS spread, set at date t with maturity $t+h$. We

[16] This description of the CDS contract is stylized and neglects particular institutional features such as the auction process when a credit event is triggered, the cheapest-to-deliver premium, potential counterparty and liquidity risk embedded in these contracts, or the fact that CDSs can be quoted in a different currency than the underlying bond. Although our assumptions may appear simplistic, they are in line with most of the reduced-form CDS term structure literature.

assume in the following that the notional is equal to one. The CDS spread is such that the present value of the payments made by the protection buyer (the fixed leg) is equal to present value of the payment made by the protection seller in case of default (the floating leg).

As far as the fixed leg is concerned, if entity i has not defaulted at date $t+k$ ($k \leq h$), the cash flow on this date is $\mathcal{S}_i(t,h)$, and is independent of k. The present value of the fixed-leg payments is denoted by $\text{PB}_i(t,h)$ and is given by:

$$\text{PB}_i(t,h) = \mathcal{S}_i(t,h) \sum_{k=1}^{h} \mathbb{E}^{\mathbb{Q}} \left[\exp\left(-\sum_{\ell=0}^{k-1} r_{t+\ell} \right) \mathbb{1}_{\{\delta_{i,t:t+k}=\mathbf{0}\}} \,\middle|\, \mathcal{F}_t \right]. \tag{21}$$

Under the RFV convention, the protection seller will make a payment of $(1-\varrho_{i,t+k})$ (the Loss-Given-Default) at date $t+k$ in case of default over the time interval $]t+k-1, t+k]$. The present value of this promised payment is given by:

$$\text{PS}_i(t,h) = \sum_{k=1}^{h} \mathbb{E}^{\mathbb{Q}} \left[\exp\left(-\sum_{\ell=0}^{k-1} r_{t+\ell} \right) (1-\varrho_{i,t+k}) \left[\mathbb{1}_{\{\delta_{i,t:t+k-1}=\mathbf{0}\}} - \mathbb{1}_{\{\delta_{i,t:t+k}=\mathbf{0}\}} \right] \,\middle|\, \mathcal{F}_t \right]. \tag{22}$$

Both expressions can be easily obtained using the same methods as for defaultable bonds. The pricing result is expressed in the following proposition.

Proposition 3.4 *The no-arbitrage CDS spread $\mathcal{S}_i(t,h)$, negotiated at date t and associated with a credit default swap (CDS) maturing in h periods, is such that the present values of the protection buyer and seller are equal, thus given by:*

$$\mathcal{S}_i(t,h) = \frac{B_i^{\text{RFV}}(t,h;\, 0,\, \mathbf{0}) - B_i^{\text{RFV}}\left(t,h;\, \omega_{i,0},\, \omega_i^{(w)}\right)}{\lim\limits_{u\to+\infty} \sum_{k=1}^{h} e^{-k\xi_0 - (\xi + u e_{\delta_i})' w_t} \times \varphi^{\mathbb{Q}}_{w_t(k)}\left(-\xi - u e_{\delta_i}, -u e_{\delta_i}\right)}, \tag{23}$$

where B_i^{RFV} is given by Equation (20) in Proposition 3.3.

Proof See online Appendix **??**. ■

As for bond prices under the RFV convention, CDS spreads are explicit but not given by an exponential-affine function of w_t. (Equation (23) expresses them as ratios of sums of exponentially-affine functions.) It should be noted that computing bond prices and CDS spreads partly involves the same Laplace transforms, thus reducing the overall number of recursions that need to be performed to price all assets. The applications presented in Section 4 consider CDS spreads denominated both in domestic and foreign currencies. The latter case is presented in Appendix **??**.

4 Applications

This section explores the pricing of sovereign credit risk from an empirical perspective. Using euro-area data, we notably explore the influence of allowing for the pricing of credit events, we compare frailty and contagion channels, and we extract measures of depreciation-at-default from CDS denominated in different currencies. The online Appendix **??** presents simulated-based evidence that the different credit-risk channels considered in our framework are indeed identifiable through maximum likelihood filtering methods.

4.1 A Model for European Sovereign Credit Risk

In this section, we exploit the framework presented above to study the pricing of sovereign credit risk. We focus on five euro-area countries: the four largest ones, that are France, Germany, Italy and Spain − accounting for 75% of the 19-country euro area GDP − and Greece, which defaulted on March 9, 2012.[17] In spite of the high credit quality of the first four countries, the associated sovereign CDS spreads have reached relatively high levels over the last twelve years, especially during the so-called euro-area sovereign debt crisis initiated in late 2009.

The framework we propose can be seen as an extension of Ait-Sahalia et al. (2014) along

[17] On March 9, 2012, the International Swaps and Derivatives Association (ISDA) decided the payment on Greek CDSs. The ISDA indeed considered that the Greek legislation that forced losses on all private creditors constituted a credit event.

the following dimensions: (a) whereas the models estimated in the latter study involve pairs of countries, ours jointly accounts for five economies, (b) by specifying the SDF, we explicitly model investors' risk preferences, opening the door to risk-premium analysis and the extraction of physical probabilities of default and (c) our model allows for both a common factor and country-specific ones, while Ait-Sahalia et al. (2014)'s models entail only country-specific factors.

We estimate different versions of the model, which allows us to revisit three credit-risk issues in the sovereign context. First, we examine the influence of allowing for credit-event pricing, that is when sovereign defaults directly affect the SDF (Subsection 4.5). Second, we discuss the differences resulting from allowing for frailty and/or contagion in the model (Subsection 4.6). Third, we extend the model in order to investigate quanto CDS spreads, that are spread differentials between euro- and dollar-denominated CDSs (Subsection 4.7).

We consider $n = 5$ economies. There are three types of components in vector y_t: the first is the short-term rate r_t; the second, denoted by z_t, is a frailty factor influencing all countries; the last n components of y_t, gathered in $x_t = (x_{1,t}, \dots, x_{n,t})$, are country-specific factors, in the sense that $x_{i,t}$ intervenes only in the default intensity of country i. The historical default intensities are given by:

$$\lambda^{\mathbb{P}}_{i,t} = \beta^{(x)}_{\lambda,i} x_{i,t} + \mathbf{C}'_i \delta_{t-1}, \tag{24}$$

where $\beta^{(x)}_{\lambda,i}$ is a scalar and where vector $\mathbf{C}_i$ is of the form $c_i \boldsymbol{\kappa}_c$, $i = 1, \dots, n$, c_i being a non-negative scalar and $\boldsymbol{\kappa}_c$ being an n-dimensional vector of country weights (summing to one). To simplify the analysis, we do not allow for potential indirect contagion here. More precisely, we set the $\mathbf{I}_j$ parameters appearing in Equation (3) to zero.[18]

Equation (24) is consistent with the general formulation (2), with $\alpha_{\lambda_i} = 0$ and $\beta^{(y)}_{\lambda_i} = (0, 0, \beta^{(x)}_{\lambda,i} e_i)$, where e_i denotes the i^{th} column of the $n \times n$ identity matrix. Conditionally on $\mathcal{F}_{t-1}$, the components of y_t are independent and we have:[19]

[18] Preliminary estimations pointed towards the non-significance of such parameters.

[19] Our short-term interest riskless rate specification does not accommodates negative yields, which is at odds with

$$
\begin{aligned}
r_t\big|\mathsf{P}_{r,t},\mathcal{F}_{t-1} &\sim \gamma_{\mathsf{P}_{r,t}}(\mu_r) &&\text{where}\quad \mathsf{P}_{r,t}\big|\mathcal{F}_{t-1} \sim \mathcal{P}\left(\alpha_r+\beta_r r_{t-1}\right),\\
z_t\big|\mathsf{P}_{z,t},\mathcal{F}_{t-1} &\sim \Gamma_{\nu_z+\mathsf{P}_{z,t}}(1) &&\text{where}\quad \mathsf{P}_{z,t}\big|\mathcal{F}_{t-1} \sim \mathcal{P}\left(\beta_z^{(z)} z_{t-1}\right),\\
x_{i,t}\big|\mathsf{P}_{x_i,t},\mathcal{F}_{t-1} &\sim \Gamma_{\mathsf{P}_{\nu_x+x_i,t}}(1) &&\text{where}\quad \mathsf{P}_{x_i,t}\big|\mathcal{F}_{t-1} \sim \mathcal{P}\left(\beta_x^{(z)} z_{t-1}+\beta_x^{(x)} x_{i,t-1}\right).
\end{aligned}
\tag{25}
$$

Last, the one-period SDF is given by:

$$
M_{t-1,t} = \exp\left(-r_{t-1}+\theta_r r_t+\theta_z z_t+\boldsymbol{\kappa}_M'(\theta_x x_t+\mathbf{S}\delta_t)-\psi^{\mathbb{P}}_{w,t-1}(\theta_w)\right), \tag{26}
$$

where $\boldsymbol{\kappa}_M$ is a vector of country weights summing to one and where $\theta_w = (\theta_r, \theta_z, \theta_x\boldsymbol{\kappa}_M, \mathbf{S}\boldsymbol{\kappa}_M)$. The previous formulation is consistent with Equation (5).

4.2 Data

The data are monthly and cover the period from January 2007 to July 2019 (ends of month). CDS spreads and proxies of risk-free zero-coupon yields are extracted from Thomson Reuters Datastream.We remove CDS spreads that do not fluctuate for three consecutive months, for this indicates low trading volumes. (In Datastream, in the absence of quotes, the last-observed ones are repeated.) We also remove (Greek) CDS spreads that are higher than 20,000 basis points.[20]

For CDS spreads and risk-free yields, we consider maturities of 1, 2, 3, 5 and 10 years. We therefore have 35 measurement equations: 25 correspond to CDS spreads, 5 correspond to risk-free zero-coupon yields and 5 correspond to the $\delta_{i,t}$. The latter are all null except for one instance: for March 2012, when the Greek sovereign default took place, the Greek credit-event variable $\delta_{i,t}$ is set to $-\log(0.22)$, consistently with an observed recovery rate of 22% (see Coudert and Gex 2013).

the data since the short-term interest rates is at -0.40% at the end of the sample. This deviation is however absorbed by the measurement error, whose standard deviation is of about 30 basis points (see row σ_{RF} of Table ??), and we believe this does not affect our credit risk results. That could be solved for by imposing that $r_t = \underline{r}+\widetilde{r}_t$, where $\underline{r}<0$ is a lower bound and $\widetilde{r}_t$ has gamma dynamics as in Equation (25).

[20]The reason why CDS spreads can potentially be above 10,000 basis points (100%) is that the payments of the premium leg are usually made on a quarterly basis (the payment being equal to the annualized spread divided by 4). The CDS spread can therefore be equal to up to 40,000 basis points if the default is almost certain in the coming month and if the recovery rate is expected to be close to zero.

4.3 Estimation Strategy

Most of the parameters are estimated by maximizing the (approximate) likelihood function obtained as a by-product of the extended Kalman filter (see online Appendix ?? for details and references regarding this type of estimation technique). To facilitate or discipline the estimation, some parameters are calibrated or constrained.

First, in order to diminish the number of parameters to be estimated, we assume that the n components of κ_M (Equation 26) are functions of countries' Gross Domestic Product (GDP). Specifically, denoting by GDP_i the GDP of country i, we assume that κ_M is proportional to $[GDP_1^\ell, \dots, GDP_n^\ell]$, where ℓ is a parameter to be estimated.[21] Second, we set $\mu_{\delta_i} = 0.6$, which makes our model consistent with the 1983-2015 average of sovereign-default recovery rates (see online Appendix ??, making use of the Moody's 2016, dataset). Third, in the spirit of Cochrane and Saa-Requejo (2000), we impose an upper bound for the sample average of the maximum one-year Sharpe ratio. As advocated by Cochrane and Saa-Requejo (2000), this bound is set to 1. It is worth noting that the fact that maximum Sharpe ratios are available in close form in an affine model is instrumental to make this approach feasible (see online Appendix ?? for details).[22]

Table ?? reports parameter estimates for eight versions of the model, whose parameterizations are summarized in Table 1. Models (1) is the "complete" version of the model, allowing for pricing of credit events (i.e. $\mathbf{S} \neq 0$), contagion (i.e. $\mathbf{C}_i \neq 0$) and a frailty component z_t. The latter component is absent from Models (5) to (8). Models (2), (4), (6) and (8) do not allow for contagion. Finally, $\mathbf{S}$ is set to zero in Models (3), (4), (7) and (8).

The penultimate row of Table ?? reports the sample average of the maximum Sharpe ratios; the bound appears to be binding for half of the estimated models, including the complete one (Model (1)). The maximum values of the log-likelihood functions are reported in the last row of

[21] We use 2018 GDPs, as measured by Eurostat.

[22] Preliminary estimations of the model – without this third restriction – yielded to extreme and unreasonable maximum Sharpe ratios, which was reflected in extremely large credit-risk premia. The last phenomenon echoes findings by Duffee (2010), who documents that maximum Sharpe ratios are often far too large when one estimates unconstrained no-arbitrage yield curve models.

Table **??**; the highest value is naturally obtained for Model (1), the most complete model.

4.4 Estimation Results (Complete Model)

Let us focus on the model featuring credit-event pricing, contagion and a frailty component (Model (1)).

Figure 1 illustrates the model fit by comparing observed CDS spreads (black crosses) to the model-implied ones (grey solid line). It appears that the model is able to capture a large share of credit spreads' fluctuations, across time and maturities. On the same figure, black solid lines represent those counterfactual CDS spreads that would be observed if the prices of risk θ_r, θ_z, θ_x and $\mathbf{S}$ (see Equation (26)) were equal to zero. This characterizes a counterfactual world where investors are not risk averse. By definition, credit-risk premia are the differentials between the latter spreads – dubbed "$\mathbb{P}$ CDS spreads" – and model-implied ones – the "$\mathbb{Q}$ CDS spreads". Figure 1 therefore confirms that credit-risk premia are substantial for all maturities, including short ones (12 months). As discussed in the next subsection, allowing for credit-event pricing (i.e. $\mathbf{S} > 0$) is instrumental to obtain sizable short term credit-risk premia.

The existence of credit-risk premia translates into differences between physical and risk-neutral probabilities of default (Figure 2). Figure 2 shows that, at the five-year horizon, the average ratios between risk-neutral and physical default probabilities $\mathbb{Q}/\mathbb{P}$ go from 1.3 (Greece) to 3.5 (France and Germany). These ratios are intermediary for Italy and Spain, with respective values of 2.6 and 2.9. The previous observations are suggestive of a negative relationship between the credit-riskiness of a country and the $\mathbb{Q}/\mathbb{P}$ ratios. This finding echoes results from the corporate credit-risk literature, according to which the part of spreads accounted for by credit loss expectations reflects a smaller fraction of yield spreads for investment-grade bonds than for lower credit-quality bonds (see e.g. Table 1 of D'Amato and Remolona, 2003, or Huang and Huang, 2012).

Estimated factors are shown in the online Appendix (Figure **??**). Unsurprisingly given its construction, the frailty factor z_t is strongly correlated to the CDS spreads, with an average correlation of 57% across countries and maturities. Most of CDSs' variability is explained by country-specific factors $x_{i,t}$ up to mid-2012. These factors reproduce the idiosyncratic country differences in credit risk observed during the first half of the sample. From mid-2012 onward, the frailty factor z_t takes over explaining the term structures of all countries' CDSs whereas all factors $x_{i,t}$ virtually converge to zero and barely fluctuate. This interestingly coincides with the *Outright Monetary Transaction* announcement in August 2012, guaranteeing potentially unlimited but conditional support to euro-area countries from the ECB. Unreported results also show that this factor relates to the euro-area unemployment rate and to the European economic policy uncertainty indicator computed by Baker et al. (2016).

4.5 The Pricing of Sovereign Credit Events

A limited number of credit-risk studies explicitly distinguish the risk of credit spread changes – if no default occurs – from the risk of the default event itself. As shown by Jarrow et al. (2005), the latter form of risk, called credit-event risk, cannot be priced if default jumps are conditionally independent across an infinite number of entities. As a consequence, after having found evidence of credit-event pricing in the context of large U.S. private firms, Driessen (2005) concludes that default jumps are not conditionally independent across the considered firms, or that not enough corporate bonds are traded to fully diversify the default jump risk. By contrast, Bai et al. (2015) find that when contagion is introduced within a general equilibrium framework for an economy comprising a large number of firms, credit-event risk premia have an upper bound of a few basis points.[23]

Because the number of sovereign entities is far smaller than the number of private borrowers,

[23] Bai et al. (2015) focus on the returns of the asset value of firms and do not explicitly consider the prices of medium-term to long-term financial instruments. In their model, each firm is associated with a jump process. The asset value of firm i falls when its own jump process is activated or, to a lesser extent, when it is the case for the jump process of another firm (which is how contagion is modeled). Because the jump intensities are constant, this model does not feature self- nor cross-excitation.

Jarrow et al. (2005)'s conditions for the absence of default event pricing are a priori not satisfied in the sovereign credit-risk context. And, as a matter of fact, our econometric results point to the existence of sovereign default event pricing. Indeed, the differences between the maximum log-likelihoods obtained for the versions of the model where $\mathbf{S}$ is restricted to be null and the respective versions where $\mathbf{S}$ is free – e.g. Model (3) vs Model (1), or Model (4) vs Model (2) – are well above critical values based on χ^2 distributions.

What are the economic implications of sovereign credit-event pricing? Recall that if $\mathbf{S} = 0$, then the physical and risk-neutral default intensities are the same (see Subsection 2.5). Since default intensities are closely related to the one-period credit spreads, having $\mathbf{S} = 0$ tends to contain credit-risk premia for short maturities. This is illustrated by Figure 3, which displays, for Model (1) and for two different dates, the model-implied term structures of CDS spreads (black solid lines), together with counterfactual term structures that are obtained after setting $\mathbf{S}$ to zero, all else being equal (dotted line). We also show the spreads that would prevail if all prices of risk were equal to zero (grey solid line). It appears that, for short-maturities, credit-risk premia are essentially accounted for by the credit-event price of risk $\mathbf{S}$.

The influence of credit-event pricing is further illustrated by Table 2, which reports the average shares of CDS spreads accounted for by credit-risk premia across the different estimated models and for two horizons: one short horizon of 6 months and a longer horizon of 5 years. In this table, figures in bold font indicate the models where $\mathbf{S}$ is allowed to be strictly positive. As expected, even for short horizons, credit-risk premia are substantial for those models allowing for credit-event pricing. For instance, the 6-month credit-risk premia amount to about 50% of CDS spreads in Model (1), against 10% in Model (3), the only difference between these two models being that $\mathbf{S} = 0$ (before estimation of the remaining parameters) in the latter. Differences are lower for larger maturities.

These results suggest that one may substantially underestimate short-term credit-risk premia when using a credit-risk model that does not allow for credit-event pricing. This, in turn, may lead to overestimation of short-term probabilities of default.

4.6 Frailty and Contagion

Exploring the default history of U.S. industrial firms between 1979 and 2004, Das et al. (2007) show that one cannot reconcile observed default clustering to standard credit-risk models where the default intensities depend solely on observable macro-finance variables. Default clustering can however be accounted for by "frailty," by which many firms can be jointly exposed to one or more unobservable risk factors (see Duffie et al. 2009). Another potential source of clustering is contagion, through which the default by one entity has a direct impact on the health of other firms. Azizpour et al. (2018) reject the hypothesis that U.S. corporates' default times are correlated only because of frailty-like mechanisms.

The previous studies build on the availability of rich datasets of corporate defaults. Sovereign defaults are rarer, making it particularly challenging to disentangle frailty from contagion in the sovereign context.[24] An advantage of the sovereign case is however the availability of various financial instruments conveying information regarding the market perception of this risk along both time and maturity dimensions. In the aftermath of the euro-area sovereign debt crisis, several studies have documented the correlation of sovereign credit spreads (see Beirne and Fratzscher 2013; Ludwig 2014; Lucas et al. 2014; Caporin et al. 2018). Nevertheless, very few studies have explicitly considered frailty or contagion mechanisms to account for the dynamics of the whole term structure of sovereign spreads. And when they do, these studies consider only one or the other phenomenon, but not both at the same time. Typically, while Ait-Sahalia et al. (2014) allow for contagion but not frailty in their no-arbitrage model; the inverse holds true for Ang and Longstaff (2013).

[24] Both types of channel have been mentioned as potential sources of sovereign default clustering. Longstaff et al. (2011) suggest that co-movements in sovereign credit risk may reflect a strong influence of global macroeconomic factors, which is rather suggestive of the frailty mechanism. Bai et al. (2012) put the emphasis on feedback mechanisms between credit and liquidity risks that may give rise to contagious spillover effects among sovereign entities. Benzoni et al. (2015) explore an alternative contagion-like mechanism, whereby agents updating fragile beliefs they have about the state of the economy.

Mechanically, the values of the maximum log-likelihoods obtained with models featuring a frailty factor – Models (1), (2), (3) and (4) – are substantially larger than those associated with models with no frailty – respectively Models (5), (6), (7) and (8). Unfortunately, because the parameters of the dynamics of z_t are not identifiable under the null hypothesis of no frailty, the distribution of the likelihood-ratio test statistic is non-standard in this situation (see e.g. Hansen 1992). By contrast, we can test whether the contagion parameters, that are the c_i's and the $\boldsymbol{\kappa}_i$'s are jointly statistically significant by employing likelihood-based tests. The test statistic being well beyond the critical values at any confidence level, we are led to reject the hypothesis of no contagion pricing.

Notwithstanding this econometric evidence, what are the economic implications of contagion? To address this question, we resort to the following exercise. For each of the models featuring contagion – Models (1), (3), (5) and (7) – we compute the average decreases in spreads obtained after killing contagion effects – setting c_i parameters to zero – all else being equal. For Models (1), (3), (5) and (7), we get average decreases of respectively 7%, 11%, 54% and 50%. Recalling that the last two models feature no frailty component, these results suggest that the (economic) importance of the contagion channel critically depends on the existence of a frailty component in the model.

4.7 Sovereign Defaults and Exchange Rates

The previous results are based on euro-denominated CDSs. However, CDS protection on many international corporations and on sovereign entities are available in euros and in U.S. dollars. (Data on both types of CDSs are collected by Thomson Reuters Datastream.) While most of European sovereign bonds are denominated in euros, a large share of European CDS are denominated in dollars. This is because the latter provides a better protection against a potential severe depreciation of the bond's currency in the case of a sovereign credit event (see Fontana and Scheicher 2010;

Augustin et al. 2020).[25] Indeed, the notional of a euro-denominated CDS is fixed in euros and that of a dollar-denominated CDS is fixed in dollars. Therefore, a euro depreciation leads to an increase of the notional of the dollar-denominated CDS expressed in euros. Formally, consider two CDS negotiated at date t: the first is a maturity-h euro-denominated CDS and the second is a dollar-denominated one with the same maturity. At inception, we consider that both CDS have identical face values, say N euros for the former and $N\exp(-s_t)$ dollars for the latter – denoting by s_t the log of the exchange rate between the domestic and the foreign currency. Assume that entity i defaults before the maturity of the contract and that the default triggers a euro depreciation: $s_{\tau_i} - s_t > 0$. Then, the payoff of the protection leg is higher for the dollar-denominated CDS than for the euro-denominated one. Indeed, we have $N(1-\varrho_{i,\tau_i})\exp(s_{\tau_i} - s_t) > N(1-\varrho_{i,\tau_i})$. ($\varrho_{i,\tau_i}$ is the recovery rate defined in Assumption 7.) Therefore, if the default by a euro-area member state is expected to be accompanied by a euro depreciation, dollar-denominated CDSs should have higher spreads than euro-denominated ones. The data are consistent with this view: the quanto CDS spreads„ defined by the deviation between a dollar-denominated CDS spread and a euro-denominated one, are mostly positive (see crosses in Figure 4).

In the following, we show that, once the previous model is augmented with the EURUSD exchange rate, it can capture the main fluctuations of the term structure of the quanto CDS spreads for the countries into consideration.

We assume that Δs_t, the one-period change in the logarithm of the real exchange rate, is given by:

$$\Delta s_t = \chi + v_t + u_\delta' \delta_t, \tag{27}$$

where v_t is an additional autonomous component of y_t, namely $y_t = (r_t, z_t, x_t', v_t)'$. If the elements of u_δ are positive, a sovereign default implies a depreciation of the euro with respect to the U.S. dollar.[26] Note that, in our framework, the depreciation upon default dies out soon after the

[25]The study of the potential liquidity differences between euro-denominated and dollar-denominated bonds, mentioned e.g. in Credit Suisse (2010), is beyond the scope of this paper.

[26]Because $v_t \geq 0$, χ has to be negative enough to allow for possible large euro appreciation (assuming the elements

default since the credit-event variables only persist through their intensities. We posit a Gamma distribution for the (i.i.d.) v_t shocks; the associated scale and shape parameters are determined in such a way as to match the sample mean and variance of observed changes in real exchange rate.[27] Estimated values of the elements of u_δ are obtained by minimizing a weighted sum of squared deviations between the observed and the model-implied quanto CDS spreads. The weights are the inverses of the sample means of squared quanto CDS spreads. (There is one weight for each country-maturity pair.)

According to the results, on average, sovereign defaults in France, Germany, Greece, Italy and Spain would respectively trigger euro depreciations of 15%, 20%, 0%, 6% and 8%.[28] As expected, these results suggest in particular that defaults by France and Germany, the two largest economies of this reduced euro-area, would have stronger impacts on the EURUSD exchange rate. Figure 4 compares observed and model-implied quanto CDS spreads. Let us stress that this result is obtained without introducing a novel latent variable (v_t is indeed observed for $t \in [1, T]$). Except for Greece, the fit is surprisingly good. For the latter country, note that quanto CDSs correspond to a small part of observed CDSs (around 5%), and one may suspect that liquidity issues outweigh the exchange-rate-related spread differentials. Leaving Greece aside, this simple model extension accounts for two thirds of the variances of observed quanto CDS spreads, on average across countries and maturities.

Figure 4 also displays the (model-implied) quanto CDS spreads that would be observed if agents were risk-neutral. These spreads are obtained by applying the pricing formulas of Proposition 3.4 under the physical measure or, equivalently, after having set the prices of risk to zero. The

of u_δ are non-negative). We set $\chi = -0.5$. This implies that the lowest possible change in the real exchange rate is of about 40% (in one month), which seems to constitute a reasonable lower bound. The results are insensitive to the value of this parameter. For instance, replacing this value by -1 or -0.10 yields virtually identical results.

[27] Since Δs_t and δ_t are observed, v_t shocks are available.

[28] These effects are deduced as follows from the components of u_δ (denoted by $u_{\delta,i}$). Note that as long as the intensity $\lambda^{\mathbb{P}}_{i,t}$ is small, conditional on having a default by country i on date t (i.e. conditional on $\mathsf{P}_{i,t} > 0$), the probability of having $\mathsf{P}_{i,t} = 1$ is close to one. It therefore comes that the distribution of $\delta_{i,t}$ on a default date is approximately $\Gamma_1(0.6)$ since $\mu_{\delta,i} = 0.6$ (see Subsection 4.3). The expected depreciation upon default is the expectation of $\exp(s_t) - 1$ conditional on $\delta_{i,t} > 0$. We thus approximate this by the expectation of $\exp(u_{\delta,i}X) - 1$ where $X \sim \Gamma_1(0.6)$; this conditional expectation directly results from the knowledge of Laplace transform of a gamma distribution.

differences between $\mathbb{Q}$ and $\mathbb{P}$ quanto CDS spreads can be interpreted as risk premia. Our results indicate that risk premia account for an important share of total quanto CDS spreads, especially for long maturities. This is consistent with the fact that quanto CDSs provide positive payoffs to the protection buyer in particularly bad states of the world (sovereign defaults).

Equation (27) assumes that credit risk affects the exchange rate through the credit-event variables only. We have considered a more general specification where Δs_t is also allowed to depend on z_t and x_t. That is, a term $u_z z_t + u_x' x_t$ is added on the right-hand side of Equation (27). This augmented flexibility hardly allows for an improvement of the fit. In addition, we have considered a specification where the term $u_z z_t + u_x' x_t$ is maintained in the specification of Δs_t, but where the credit-event variables are removed (i.e. $u_\delta = 0$). The fit resulting from this alternative specification is very poor.[29] Altogether, these results suggest that it is the relationship between the exchange rate and the credit events *per se*, and less between the exchange rate and conditional default probabilities – driven by $(z_t, x_t')'$ – that is key to explain the fluctuation of quanto CDS spreads. These results are in line with those of Ehlers and Schonbucher (2004) and of Brigo et al. (2015).

5 Conclusion

We present a general affine positive credit-risk model able to simultaneously relax restrictive assumptions often employed in the reduced-form credit-risk literature while preserving tractability in the pricing of default-sensitive securities. Building on the non-negative affine Gamma-zero process, the model accommodates the presence of systemic risk (i.e. potential feedbacks from defaults to common risk factors), contagion, credit-event pricing and stochastic recovery rates. We provide explicit formulas to price defaultable bonds and CDS, for different recovery-rate conventions.

We exploit this framework to investigate the pricing of sovereign credit risk using euro-area

[29] For our the most complete specification (Model (1)), the ratios of mean squared pricing errors to mean squared quanto CDS spreads are of 23%, 21%, 14% and 12% on average across maturities for Germany, France, Italy and Spain, respectively. For the specification where credit-event variables cannot affect the exchange rate, the same ratios are 72%, 70%, 57% and 62%, respectively.

data. We show that one common factor and one country-specific factor for each country allow for a very good fit of CDS data. The estimation detects contagion effects, even when allowing for a frailty factor. Moreover, we find sizable credit-risk premia along the whole maturity spectrum. Typically, credit-risk premia account for more than half of CDS spreads at the five year maturity for France, Germany, Italy and Spain. Our findings also highlight the importance of credit-event pricing to allow for non-trivial short-term credit-risk premia. A simple extension of the model finally allows us to extract measures of expected (EURUSD) depreciations-at-default by jointly modeling term structures of sovereign CDSs denominated in euros and in U.S. dollars.

References

Ait-Sahalia, Y., J. Cacho-Diaz, and R. J. A. Laeven. 2015. "Modeling Financial Contagion Using Mutually Exciting Jump Processes". *Journal of Financial Economics* 117:585–606.

Ait-Sahalia, Y., R. J. A. Laeven, and L. Pelizzon. 2014. "Mutual excitation in Eurozone sovereign CDS". *Journal of Econometrics* 183:151–167.

Almeida, H., and T. Philippon. 2007. "The Risk-Adjusted Cost of Financial Distress". *Journal of Finance* 62:2557–2586.

Altman, E., B. Brady, A. Resti, and A. Sironi. 2005. "The Link between Default and Recovery Rates: Theory, Empirical Evidence, and Implications". *The Journal of Business* 78 (6): 2203–2228.

Ang, A., and F. Longstaff. 2013. "Systemic sovereign credit risk: Lessons from the U.S. and Europe". *Journal of Monetary Economics* 60 (5): 493–510.

Asonuma, T., and C. Trebesch. 2016. "Sovereign Debt Restructurings: Preemptive or Post-Default". *Journal of the European Economic Association* 14 (1): 175–214.

Augustin, P., M. Chernov, and D. Song. 2020. "Sovereign credit risk and exchange rates: Evidence from CDS quanto spreads". *Journal of Financial Economics.*

Azizpour, S., K. Giesecke, and G. Schwenkler. 2018. "Exploring the sources of default clustering". *Journal of Financial Economics* 129 (1): 154–183.

Bai, J., P. Collin-Dufresne, R. Goldstein, and J. Helwege. 2015. "On bounding credit-event risk premia". *Review of Financial Studies* 28 (9): 2608–2642.

Bai, J., C. Julliard, and K. Yuan. 2012. *Eurozone sovereign bond crisis: Liquidity or fundamental contagion.* Unpublished working paper. Federal Reserve Bank of New York.

Baker, S. R., N. Bloom, and S. J. Davis. 2016. "Measuring economic policy uncertainty". *The Quarterly Journal of Economics* 131 (4): 1593–1636.

Barro, R. 2006. "Rare disasters and asset markets in the twentieth century". *The Quarterly Journal of Economics* 121 (3): 823–866.

Beirne, J., and M. Fratzscher. 2013. "The pricing of sovereign risk and contagion during the European sovereign debt crisis". *Journal of International Money and Finance* 34 (C): 60–82.

Benzoni, L., P. Collin-Dufresne, R. Goldstein, and J. Helwege. 2015. "Modeling credit contagion via the updating of fragile beliefs". *Review of Financial Studies* 28 (7): 1960–2008.

Brennan, M., and E. Schwartz. 1980. "Analyzing Convertible Bonds". *Journal of Financial and Quantitative Analysis* 15:907–929.

Brigo, D., N. Pede, and A. Petrelli. 2015. *Multi currency credit default swaps: quanto effects and FX devaluation jumps.* Working Paper. Imperial College London.

Caporin, M., L. Pelizzon, F. Ravazzolo, and R. Rigobon. 2018. "Measuring sovereign contagion in Europe". *Journal of Financial Stability* 34 (C): 150–181.

Chen, L., and D. Filipovic. 2007. "Credit Derivatives in an Affine Framework". *Asia-Pacific Financial Markets* 14:123–140.

Cochrane, J. H., and J. Saa-Requejo. 2000. "Beyond arbitrage: Good-deal asset price bounds in incomplete markets". *Journal of Political Economy* 108 (1): 79–119.

Collin-Dufresne, P., R. Goldstein, and J. Hugonnier. 2004. "A General formula for valuing defaultable securities". *Econometrica* 72 (5): 1377–1407.

Coudert, V., and M. Gex. 2013. "Why the Greek CDS settlement did not lead to the feared meltdown". *Financial Stability Review*, no. 17: 135–150.

Credit Suisse. 2010. *Sovereign CDS Primer*. Fixed Income Research.

D'Amato, J., and E. M. Remolona. 2003. "The credit spread puzzle". *BIS Quarterly Review.*

Das, P., Sanjiv R.and Hanouna. 2009. "Implied recovery". *Journal of Economic Dynamics and Control* 33 (11): 1837–1857.

Das, S. R., D. Duffie, N. Kapadia, and L. Saita. 2007. "Common failings: How corporate defaults are correlated". *Journal of Finance* 62 (1): 93–117.

Driessen, J. 2005. "Is the default event risk priced in corporate bonds?" *Review of Financial Studies* 18 (1): 165–195.

Duan, J.-C., and W. Miao. 2016. "Default correlations and large-portfolio credit analysis". *Journal of Business & Economic Statistics* 34 (4): 536–546.

Duffee, G. 2010. *Sharpe Ratios in Term Structure Models.* Economics Working Paper Archive 575. The Johns Hopkins University,Department of Economics.

Duffie, D. 2005. "Credit risk modeling with affine processes". *Journal of Banking and Finance* 29:2751–2802.

Duffie, D., A. Eckner, G. Horel, and L. Saita. 2009. "Frailty correlated default". *Journal of Finance* 64 (5): 2089–2123.

Duffie, D., M. Schroder, and C. Skiadas. 1996. "Recursive valuation of the defaultable securities and the timing of resolution of uncertainty". *The Annals of Applied Probability* 6 (4): 1075–1090.

Duffie, D., and K. Singleton. 1999. "Modeling term structures of defaultable bonds". *The Review of Financial Studies* 12 (4): 687–720.

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

Ehlers, P., and P. Schonbucher. 2004. *The influence of FX risk on credit spreads*. Working Paper Series. ETH Zurich.

Eraker, B. 2008. "Affine general equilibrium models". *Management Science* 54 (12): 2068–2080.

Fontana, A., and M. Scheicher. 2010. "An Analysis of Euro Area Sovereign CDS and their Relation with Government Bonds." *Journal of Banking and Finance* 62:126–140.

Gabaix, X. 2012. "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance". *The Quarterly Journal of Economics* 127 (2): 645–700.

Giesecke, K., F. Longstaff, S. Schaefer, and I. Strebulaev. 2011. "Corporate bond default risk: A 150-year perspective". *Journal of Financial Economics* 102 (2): 233–250.

Giesecke, K., and S. Zhu. 2013. "Transform Analysis for Point Processes and Applications in Credit Risk". *Mathematical Finance* 23 (4): 742–762.

Gourieroux, C., and J. Jasiak. 2006. "Autoregressive gamma processes". *Journal of Forecasting* 25 (2): 129–152.

Gourieroux, C., A. Monfort, and J.-P. Renne. 2014. "Pricing Default Events: Suprise, Exogeneity and Contagion". *Journal of Econometrics* 182.

Guo, X., R. A. Jarrow, and Y. Zeng. 2009. "Modeling the Recovery Rate in a Reduced Form Model". *Mathematical Finance* 19 (1).

Hansen, B. E. 1992. "The likelihood ratio test under nonstandard conditions: Testing the Markov switching model of GNP". *Journal of Applied Econometrics* 7 (S): 61–82.

Hansen, L. P., and S. F. Richard. 1987. "The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models". *Econometrica* 55 (3): 587–613.

Huang, J.-Z., and M. Huang. 2012. "How much of the corporate Treasury yield spread is due to credit risk?" *The Review of Asset Pricing Studies* 2 (2): 153–202.

Jarrow, R. A., D. Lando, and F. Yu. 2005. "Default risk and diversification: Theory and empirical implications". *Mathematical Finance* 15 (1): 1–26.

Jarrow, R. A., and S. Turnbull. 1995. "Pricing Derivatives on Financial Securities Subject to Credit Risk". *Journal of Finance* 50 (1): 53–86.

Jarrow, R. A., and F. Yu. 2001. "Counterparty risk and the pricing of defaultable securities". *Journal of Finance* 56 (5): 1765–1799.

Kraft, H., and M. Steffensen. 2007. "Bankruptcy, Counterparty Risk, and Contagion". *Review of Finance* 11:209–252.

Lando, D. 1998. "On Cox processes and credit risky securities". *The Review of Derivatives Research* 2:99–120.

Longstaff, F., J. Pan, L. Pedersen, and K. Singleton. 2011. "How sovereign is sovereign credit risk?" *American Economic Journal, Macroeconomics* 3:75–103.

Longstaff, F., and E. Schwartz. 1995. "A Simple Approach to Valuing Risky Fixed and Floating Rate Debt". *Journal of Finance* 50 (3): 789–821.

Lucas, A., B. Schwaab, and X. Zhang. 2014. "Conditional euro area sovereign default risk". *Journal of Business & Economic Statistics* 32 (2): 271–284.

Ludwig, A. 2014. "A unified approach to investigate pure and wake-up-call contagion: Evidence from the Eurozone's first financial crisis". *Journal of International Money and Finance* 48:125–146.

Mendoza, E. G., and V. Z. Yue. 2012. “A general equilibrium model of sovereign default and business cycles”. *The Quarterly Journal of Economics* 127 (2): 889–946.

Monfort, A., F. Pegoraro, J.-P. Renne, and G. Roussellet. 2017. “Staying at Zero with Affine Processes: An Application to Term Structure Modelling”. *Journal of Econometrics* 201, no. 2 (): 348–366.

Moody’s. 2016. *Sovereign Default and Recovery Rates, 1983-2015*. Moody’s Investor Service Data Report.

Na, S., S. Schmitt-Grohe, M. Uribe, and V. Z. Yue. 2018. “The twin Ds: Optimal default and devaluation”. *American Economic Review* 108 (7): 1773–1819.

Pan, J., and K. Singleton. 2008. “Default and recovery implicit in the term structure of sovereign CDS spreads”. *The Journal of Finance* 63:2345–2384.

Reinhart, C. M. 2002. “Default, currency crises, and sovereign credit ratings”. *World Bank Economic Review* 16 (2): 151–170.

Reinhart, C. M., and K. S. Rogoff. 2011. “The forgotten history of domestic debt”. *Economic Journal* 121 (552): 319–350.

Trebesch, C., and M. Zabel. 2017. “The output costs of hard and soft sovereign default”. *European Economic Review* 92:416–432.

Tables and Figures

Table 1: Summary of model parameterizations

Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Frailty	✓	✓	✓	✓	✗	✗	✗	✗
Contagion	✓	✗	✓	✗	✓	✗	✓	✗
Credit-event pricing	✓	✓	✗	✗	✓	✓	✗	✗

Table 2: Shares of CDS spreads corresponding to credit-risk premia

Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A – Horizon = 6 months							
Germany	**0.510**	**0.409**	0.118	0.158	**0.289**	**0.587**	0.086	0.122
France	**0.518**	**0.409**	0.112	0.158	**0.268**	**0.571**	0.068	0.125
Italy	**0.503**	**0.409**	0.110	0.157	**0.209**	**0.558**	0.067	0.130
Spain	**0.502**	**0.409**	0.105	0.157	**0.208**	**0.540**	0.062	0.128
Greece	**0.388**	**0.400**	0.087	0.141	**0.142**	**0.458**	0.043	0.109
	Panel B – Horizon = 60 months							
Germany	**0.784**	**0.714**	0.614	0.729	**0.752**	**0.733**	0.542	0.606
France	**0.767**	**0.701**	0.585	0.716	**0.727**	**0.714**	0.510	0.603
Italy	**0.701**	**0.637**	0.489	0.642	**0.670**	**0.671**	0.491	0.548
Spain	**0.724**	**0.661**	0.520	0.667	**0.656**	**0.670**	0.467	0.584
Greece	**0.466**	**0.498**	0.260	0.425	**0.475**	**0.452**	0.353	0.220

Note: This table reports the sample averages of the shares of CDS spreads corresponding to credit-risk premia. Credit-risk premia are defined as the difference between model-implied CDS spreads and the counterfactual CDS spreads obtained after having set the prices of risk to zero (the prices of risk are the components of θ_w, see Equation 26). Figures in bold font are for models where $\mathbf{S} > 0$.

Figure 1: Observed vs model-implied CDSs

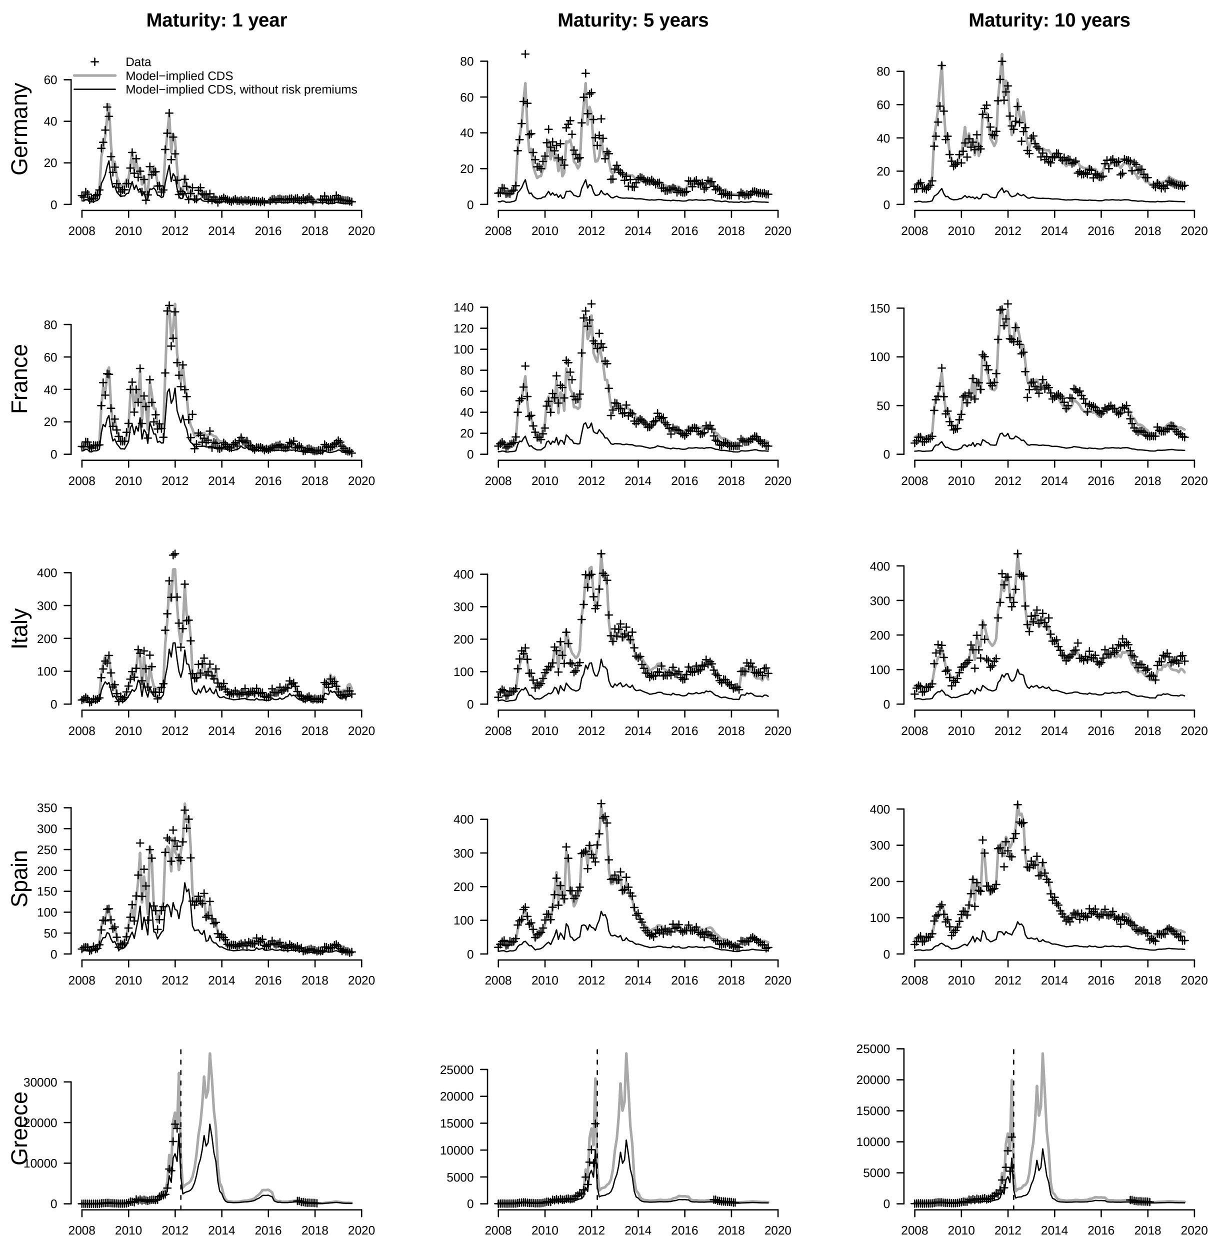

Note: The gray lines correspond to the model-implied CDS spreads, expressed in basis points. The data span the period from January 2007 to July 2019 at the monthly frequency. The thin black line corresponds to (model-implied) $\mathbb{P}$ CDS spreads, that are the spreads that would be observed if agents were not risk averse. The $\mathbb{P}$ CDS spreads are obtained by applying the CDS pricing formulas after having set the prices of risk (θ_x, θ_y, θ_r and $\mathbf{S}$) to zero. For Greece: the vertical dashed bar indicates the default period (March 2012).

Figure 2: Model-implied probabilities of default

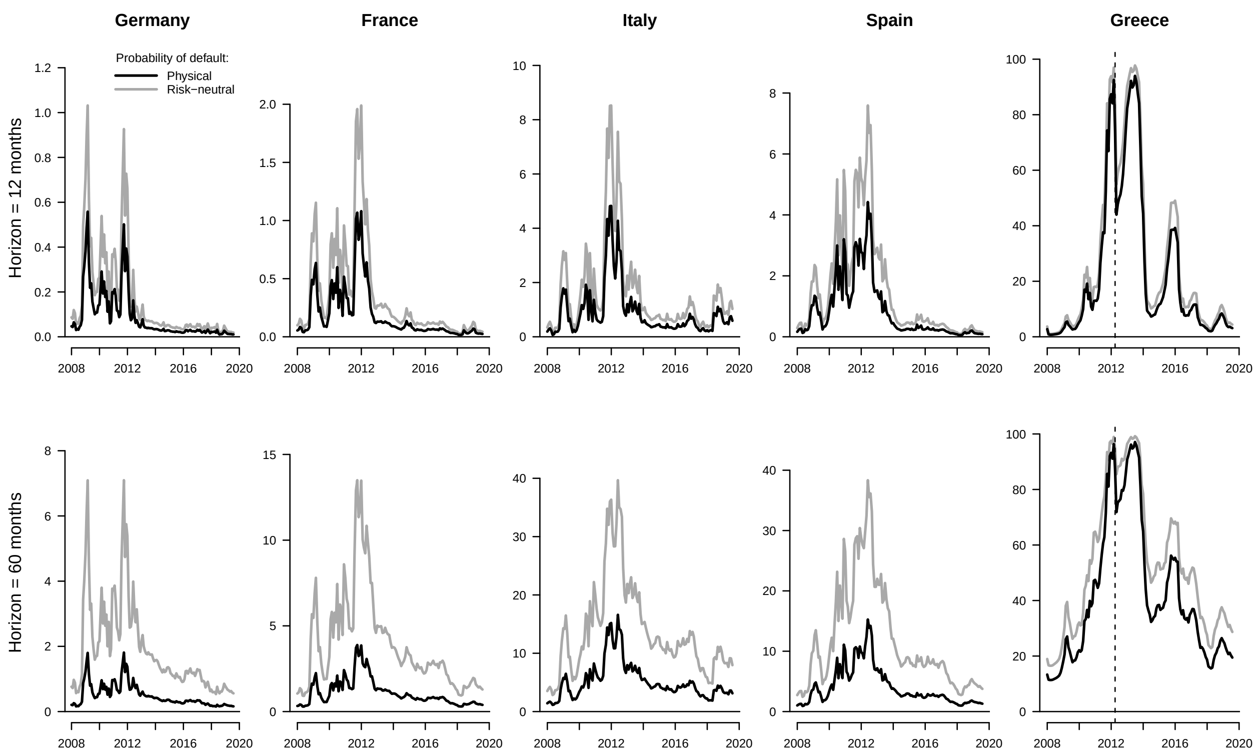

Note: This figure displays model-implied conditional probabilities of default for two horizons: 12 months and 60 months. The grey (respectively black) line corresponds to the physical (respectively risk-neutral) probability of default. The difference between the two curves reflects credit-risk premia. For Greece: the vertical dashed bar indicates the default period (March 2012).

Figure 3: Model-implied term structures of probabilities of default

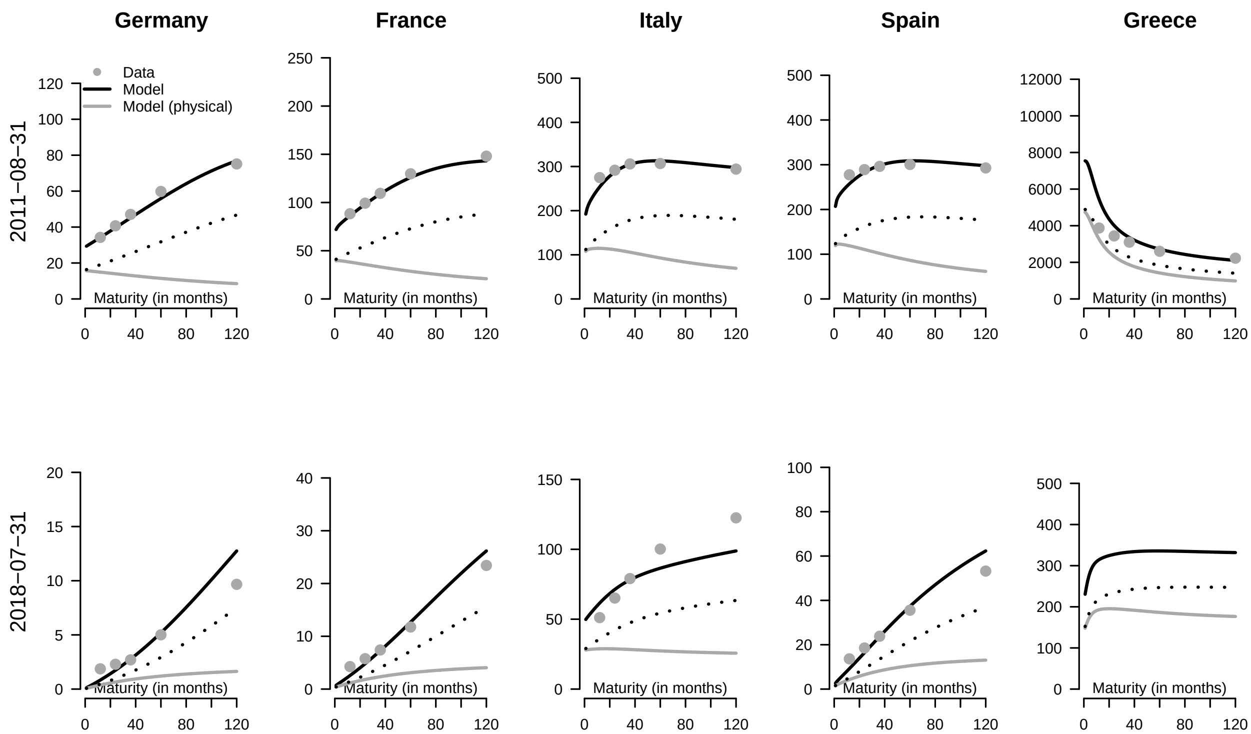

Note: This figure displays model-implied term structures of CDS spreads (solid black lines) together with observed CDS spreads (grey dots) for two dates (under the complete model, i.e. Model (1) in Table **??**). The solid grey line represents the CDS spreads that would prevail if agents were risk-neutral or, equivalently, if the prices of risk (θ_w, see Equation 26) were null. The spread differentials between the solid black line and the grey solid line therefore correspond to credit-risk premia. The dashed line corresponds to the CDS spreads that would prevail – everything else equal – if the credit-event pricing parameter $\mathbf{S}$ was equal to zero. The figure shows in particular that short-term credit-risk premia are small in the latter case.

Figure 4: Quanto CDS

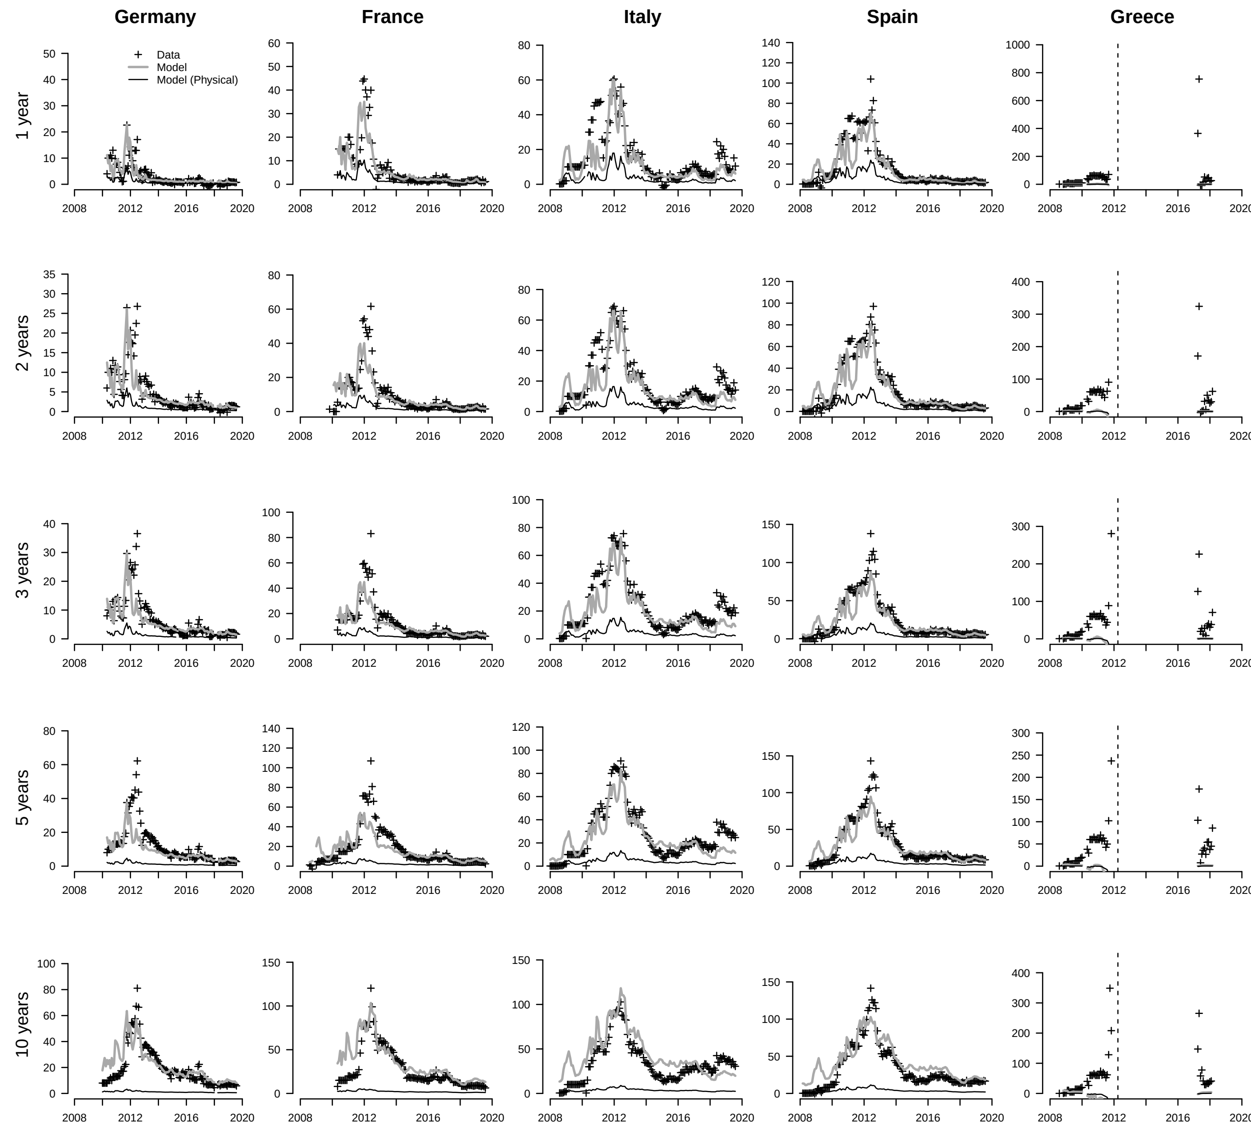

Note: This figure compares observed and model-implied quanto CDS spreads (expressed in basis points). quanto CDS spreads are defined as the differences between dollar-denominated CDS premia and their euro-denominated counterparts. For some countries and maturities, Datastream-extracted CDS premia are the same for the euro- and dollar-denominated CDS; in these cases, the data are removed from the estimation sample. Data points are also removed when CDS premia do not change for three months in a row (which indicates illiquidity). The thin solid line corresponds to the (model-implied) quanto CDS spreads that would be observed if agents were risk-neutral. For Greece: the vertical dashed bar indicates the default period (March 2012).