

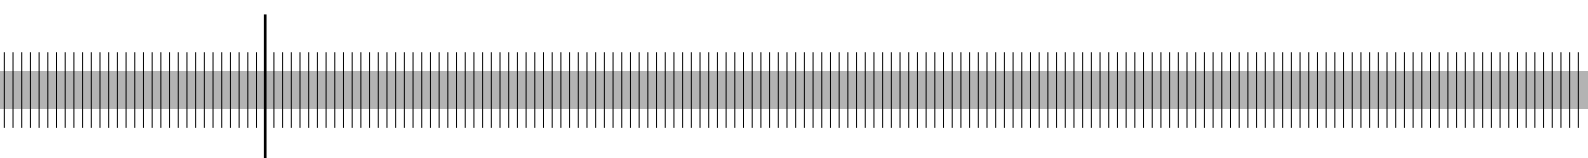
The pricing of correlated default risk: evidence from the credit derivatives market

Nikola Tarashev

(Bank for International Settlements)

Haibin Zhu

(Bank for International Settlements)



Discussion Paper
Series 2: Banking and Financial Studies
No 09/2008

Discussion Papers represent the authors' personal opinions and do not necessarily reflect the views of the Deutsche Bundesbank or its staff.

Editorial Board: Heinz Herrmann
Thilo Liebig
Karl-Heinz Tödter

Deutsche Bundesbank, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main,
Postfach 10 06 02, 60006 Frankfurt am Main

Tel +49 69 9566-1

Telex within Germany 41227, telex from abroad 414431

Please address all orders in writing to: Deutsche Bundesbank,
Press and Public Relations Division, at the above address or via fax +49 69 9566-3077

Internet <http://www.bundesbank.de>

Reproduction permitted only if source is stated.

ISBN 978-3-86558-408-3 (Printversion)

ISBN 978-3-86558-409-0 (Internetversion)

This paper was presented at the joint Deutsche Bundesbank / Basel Committee on Banking Supervision / Journal of Banking and Finance conference (December 2007) on “Interaction of Market and Credit Risk”. The views expressed in the paper are those of the authors and not necessarily those of the Bundesbank, the Basel Committee on Banking Supervision or the Journal of Banking and Finance.

Conference on the Interaction of Market and Credit Risk

6–7 December 2007, Berlin

Thursday, 6 December

8:30 – 9:00 Registration (Harnack Haus)

9:00 – 9:15 **Welcome Address by Hans Reckers (Deutsche Bundesbank)**

Session 1 Banking and Securitization

Chair: Myron Kwast (Federal Reserve Board)

9:15 – 10:15 **Recent Financial Market Developments**

Keynote address by E. Gerald Corrigan (Goldman Sachs)

10:15 – 11:05 **Banking and Securitization**

Wenying Jiangli (Federal Deposit Insurance Corporation)

Matthew Pritsker (Federal Reserve Board)

Peter Raupach (Deutsche Bundesbank)

Discussant: Deniz O. Igan (International Monetary Fund)

11:05 – 11:30 **Refreshments**

Session 2 Integrated Modelling of Market and Credit Risk I

Chair: Klaus Duellmann (Deutsche Bundesbank)

11:30 – 12:10 **Regulatory Capital for Market and Credit Risk Interaction: Is Current Regulation Always Conservative?**

Thomas Breuer (Fachhochschule Vorarlberg)

Martin Jandačka (Fachhochschule Vorarlberg)

Klaus Rheinberger (Fachhochschule Vorarlberg)

Martin Summer (Oesterreichische Nationalbank)

Discussant: Simone Manganelli (European Central Bank)

- 12:10 – 13:00 **An Integrated Structural Model for Portfolio Market and Credit Risk**
Paul H. Kupiec (Federal Deposit Insurance Corporation)

Discussant: Dan Rosen (R² Financial Technologies Inc.)
- 13:00 – 14:30 **Lunch**
- Session 3** **Integrated Modelling of Market and Credit Risk II**
Chair: Til Schuermann (Federal Reserve Bank of New York)
- 14:30 – 15:20 **The Integrated Impact of Credit and Interest Rate Risk on Banks: An Economic Value and Capital Adequacy Perspective**
Mathias Drehmann (European Central Bank)
Steffen Sorensen (Bank of England)
Marco Stringa (Bank of England)

Discussant: Jose A. Lopez (Federal Reserve Bank of San Francisco)
- 15:20 – 16:10 **An Economic Capital Model Integrating Credit and Interest Rate Risk**
Piergiorgio Alessandri (Bank of England)
Mathias Drehmann (European Central Bank)

Discussant: Andrea Sironi (Bocconi University)
- 16:10 – 16:40 **Refreshments**
- 16:40 – 18:00 **Panel discussion**
Moderator: Myron Kwast (Federal Reserve Board)
Panelists: Pierre Cailleteau (Moody's),
Christopher Finger (RiskMetrics),
Andreas Gottschling (Deutsche Bank),
David M. Rowe (SunGard)
- 20:00 **Conference Dinner (with Gerhard Hofmann, Deutsche Bundesbank)**

Friday, 7 December

- Session 4**
- Risk Measurement and Markets**
- Chair: Thilo Liebig (Deutsche Bundesbank)**
- 9:00 – 9:50 **A Value at Risk Analysis of Credit Default Swaps**
Burkhard Raunig (Oesterreichische Nationalbank)
Martin Scheicher (European Central Bank)
- Discussant: Alistair Milne (Cass Business School)
- 9:50 – 10:40 **The Pricing of Correlated Default Risk: Evidence From the Credit Derivatives Market**
Nikola Tarashev (Bank for International Settlements)
Haibin Zhu (Bank for International Settlements)
- Discussant: David Lando (Copenhagen Business School)
- 10:40 – 11:10 **Refreshments**
- 11:10 – 12:10 **Structural Models and the Linkage between Equity and Credit Markets**
Keynote Address by Hayne Leland (The University of California, Berkeley)
- Session 5A**
- Securitization, Regulation and Systemic Risk**
- Chair: Hayne Leland (The University of California, Berkeley)**
- 12:10 – 13:00 **Solvency Regulation and Credit Risk Transfer**
Vittoria Cerasi (Milano-Bicocca University)
Jean-Charles Rochet (Toulouse University)
- Discussant: Lorian Pelizzon (University of Venice)
- 13:00 – 14:30 **Lunch**
- 14:30 – 15:20 **Determinants of Banks' Engagement in Loan Securitization**
Christina E. Bannier (Frankfurt School of Finance and Management)
Dennis N. Hänsel (Goethe University Frankfurt)
- Discussant: Gabriel Jimenez (Bank of Spain)

15:20 – 16:10 **Systemic Bank Risk in Brazil: An Assessment of Correlated Market, Credit, Sovereign and Inter-Bank Risk in an Environment with Stochastic Volatilities and Correlations**

Theodore M. Jr. Barnhill (The George Washington University)

Marcos Rietti Souto (International Monetary Fund)

Discussant: Mathias Drehmann (European Central Bank)

Session 5B **Credit Dependencies and Market Risk**

Chair: Kostas Tsatsaronis (BIS)

12:10 – 13:00 **Interaction of Market and Credit Risk: An Analysis of Inter-Risk Correlation and Risk Aggregation**

Klaus Böcker (UniCredit Group)

Martin Hillebrand (Sal. Oppenheim)

Discussant: Rüdiger Frey (University of Leipzig)

13:00 – 14:30 **Lunch**

14:30 – 15:20 **Market Conditions, Default Risk and Credit Spread**

Dragon Tang (Kennesaw State University)

Hong Yan (University of South Carolina)

Discussant: Til Schuermann (Federal Reserve Bank of New York)

15:20 – 16:10 **The Effect of Seniority and Security Covenants on Bond Price Reactions to Credit News**

David D. Cho (State University of New York at Buffalo)

Hwagyun Kim (Texas A&M University)

Jungsoon Shin (State University of New York at Buffalo)

Discussant: Joerg Rocholl (European School of Management and Technology in Berlin)

16:10 – 16:30 **Final Remarks by Philipp Hartmann (European Central Bank)**

16:30 – 17:00 **Refreshments**

Deutsche Bundesbank
Conferences, 70-4

Tel: +49 69 9566-2923/-2864, Fax: +49 69 9566-4027

E-mail: conferences@bundesbank.de

The Pricing of Correlated Default Risk: Evidence from the Credit Derivatives Market*

Nikola Tarashev[†]
Haibin Zhu[‡]

Abstract

In order to analyze the pricing of portfolio credit risk – as revealed by tranche spreads of a popular credit default swap (CDS) index – we extract *risk-neutral* probabilities of default (PDs) and *physical* asset return correlations from single-name CDS spreads. The time profile and overall level of index spreads validate our PD measures. At the same time, the physical asset return correlations are too low to account for the spreads of index tranches and, thus, point to a large correlation risk premium. This premium, which covaries negatively with current realized correlations and positively with future realized correlations, sheds light on market perceptions of and attitude towards correlation risk.

JEL Classification Numbers: G12, G13, C15

Keywords: Portfolio credit risk, Correlation risk premium, CDS index, Tranche spread, Copula.

*This paper was previously circulated under the title “The Pricing of Portfolio Credit Risk”. We are grateful to Claudio Borio, Joost Driessen, Song Han, Peter Hoerdahl, Nikunj Kapadia, Srichander Ramaswamy, Kostas Tsatsaronis, Hao Zhou and seminar participants at the Bank for International Settlements, the 2006 FDIC Derivatives Conference, the 2006 CREDIT conference in Venice, the 2006 Meeting of the European Econometric Society in Vienna, the 2007 Federal Reserve Board credit risk conference, and the Swiss Finance Institute conference on “Portfolio Management and Derivatives” for helpful comments. We also thank Marcus Jellinghaus for valuable help with the data. The views presented here are our own and do not necessarily represent those of the Bank for International Settlements.

[†]Nikola Tarashev: Research and Policy Analysis, Bank for International Settlements, Basel, Switzerland. Tel.: 41-61-280-9213. Fax: 41-61-280-9100. E-mail: nikola.tarashev@bis.org.

[‡]Haibin Zhu: Research and Policy Analysis, Bank for International Settlements, Basel, Switzerland. Tel.: 41-61-280-9164. Fax: 41-61-280-9100. E-mail: haibin.zhu@bis.org.

Non-technical summary

Portfolio credit risk has three key components: probability of default (PD), loss given default (LGD) and the probability distribution of joint defaults. With the rapid development of innovative products in structured finance, the third component has grown remarkably in importance. However, there is no consensus on how market participants estimate it.

In this paper, we first propose an approach to deriving the probability distribution of joint defaults on the basis of data from the credit default swap (CDS) market. This approach extracts *risk-neutral* PDs and *physical* asset return correlations from the levels and co-movements of single-name CDS spreads. Then, in a concrete application of our approach, we use these estimates to compute *predicted* tranche spreads of a popular CDS index – the Dow Jones CDX North America Investment-Grade Index – and compare them with *empirical* spreads from the CDS index market.

We find that predicted spreads fall short of accounting for their empirical counterparts tranche by tranche. In particular, observed spreads of protection against a large (small) number of defaults are much higher (lower) than predicted ones. Further analysis reveals that such discrepancies are likely to be driven by the existence of a correlation risk premium, defined as the *risk-neutral* asset return correlation, used for pricing CDS index tranches, minus the *physical* correlation underpinning predicted spreads. On average, the correlation risk premium amounts to 66% of the physical correlation.

Our analysis also sheds light on the joint intertemporal behavior of the correlation risk premium and the physical correlation. When the contemporaneous realization of physical correlations increases, the correlation risk premium tends to decrease, consistent with the market anticipating lower upward correlation risk or having greater appetite for correlation risk. In addition, a positive change in the correlation risk premium is associated with a future increase in physical correlations, suggesting that either the market tends to price correctly correlation risk or changes in the attitude towards such risk feed back into market outcomes down the road.

Nichttechnische Zusammenfassung

Das Portfoliokreditrisiko setzt sich aus drei Hauptkomponenten zusammen: der Ausfallwahrscheinlichkeit (probability of default, PD), der Verlustquote (loss given default, LGD) und der Wahrscheinlichkeitsverteilung für gemeinsame Ausfälle. Mit der rasanten Entwicklung innovativer Produkte im Bereich der strukturierten Finanzierung ist die Bedeutung der dritten Komponente zusehends gestiegen. Allerdings herrscht keine Einigkeit darüber, wie die Marktteilnehmer diese schätzen.

Im vorliegenden Arbeitspapier schlagen wir zunächst einen auf CDS-Marktdaten beruhenden Ansatz zur Ableitung der Wahrscheinlichkeitsverteilung für gemeinsame Ausfälle vor. Mit diesem Ansatz werden *risikoneutrale* PDs und *physische* Asset-Return-Korrelationen aus der Höhe der Preise und dem Gleichlauf (Co-movement) von Single-name-CDS-Spreads abgeleitet. Anschließend benutzen wir diese Schätzungen in einer konkreten Anwendung unseres Ansatzes zur Berechnung von Prognosen für Tranchenspreads eines bekannten CDS-Index (Dow Jones CDX North America Investment Grade Index) und vergleichen diese mit empirischen Spreads am CDS-Indexmarkt.

Wir stellen fest, dass die Prognosen für die Spreads nicht in der Lage sind, den entsprechenden empirischen Spreads mit Blick auf die betreffenden Tranchen Rechnung zu tragen. So sind insbesondere die beobachteten Sicherheitsaufschläge für den Fall einer großen (kleinen) Zahl von Ausfällen wesentlich höher (niedriger) als die Prognosen. Weitere Untersuchungen zeigen, dass solche Abweichungen wahrscheinlich durch das Vorhandensein einer Korrelationsrisikoprämie bedingt sind – diese ist definiert als die *risikoneutrale* Asset-Return-Korrelation, die zur Preisgestaltung von CDS-Indextranchen verwendet wird, abzüglich der den Prognosen für die Spreads zugrunde liegenden *physischen* Korrelation. Die Korrelationsrisikoprämie beträgt durchschnittlich 66 % der physischen Korrelation.

Unsere Analyse gibt auch Aufschluss über das gemeinsame intertemporale Verhalten von Korrelationsrisikoprämie und physischer Korrelation. Wenn die zeitgleiche Realisierung physischer Korrelationen steigt, nimmt die

Korrelationsrisikoprämie – im Einklang mit der Markterwartung eines geringeren aufwärts gerichteten Korrelationsrisikos oder einer größeren Korrelationsrisikoneigung der Marktteilnehmer – tendenziell ab. Darüber hinaus ist eine positive Veränderung der Korrelationsrisikoprämie mit einem künftigen Anstieg der physischen Korrelationen verbunden, was darauf hindeutet, dass entweder der Markt das Korrelationsrisiko tendenziell richtig bepreist, oder sich Veränderungen in der Einstellung gegenüber solchen Risiken zu einem späteren Zeitpunkt in den Marktergebnissen niederschlagen.

1 Introduction

Portfolio credit risk has three key components: probability of default (PD), loss given default (LGD) and the probability distribution of joint defaults.¹ The last component, which has received least attention owing to the traditional focus of academic researchers and market practitioners on single-name credit events, has recently gained in importance as a result of the rapid development of innovative products in structured finance. Such products, which allow investors to trade portfolio credit risk, include collateralized debt obligations (CDO), CDOs of CDOs (or CDO²), nth-to-default credit default swaps (CDS) and CDS indices.² The prices of these financial instruments rely heavily on estimated probabilities of joint defaults (see Hull and White, 2004; Gibson, 2004) but here is no consensus on how market participants construct such estimates.

There are two popular approaches to deriving the probability distribution of joint defaults on the basis of firm-level data. One approach estimates this distribution directly from historical data on defaults (Daniels et al., 2005; Demey et al., 2004; Jarrow and van Deventer, 2005; Das et al., 2007). Since defaults are rare events, however, this approach could lead to large estimation errors, especially for portfolios comprising investment-grade entities. The second approach derives the probability of joint defaults indirectly, on the basis of equity market data and firms' balance sheet information. This approach builds on the Merton (1974) framework and exploits (i) the notion that a default occurs when a borrower's assets fall below some threshold and (ii) the insight that, since equity is a call option on the value of the firm, equity prices reflect asset values. This implies that asset return correlations underpin, together with individual PDs, the probability distribution of joint defaults and, from a practical point of view, can be extracted from readily observable equity-return correlations.³ Although it is inherently model dependent, the second

¹The mainstream of the credit risk literature focuses on PD: see Duffie and Singleton (2003) for an overview. The growing literature on LGD includes Altman and Kishore (1996), Jarrow (2001), Covitz and Han (2004), Pan and Singleton (2005) and Carey and Gordy (2006).

²For a general discussion of products used for trading portfolio credit risk, see BCBS (2004).

³The proprietary Global Correlation model by Moody's KMV is one such example. See Das and Ishii

approach possesses an important advantage because it uses risk parameters that can be estimated from large data sets.⁴

In this paper, we follow the spirit of the second approach but propose a method for estimating the probability distribution of joint defaults on the basis of data from the rapidly-growing CDS market. Our estimation method allows us to investigate to what extent the pricing of the credit risk of *individual* companies could help one understand market valuation of *portfolio* credit risk. As a concrete instance of such valuation, we consider five-year spreads on different tranches of a popular CDS index: the Dow Jones CDX North America investment-grade index (CDX.NA.IG.5Y). In order to dissect these tranche spreads, we use data on single-name CDS spreads and market estimates of LGDs, which imply two sets of credit risk parameters that vary over time and across firms: *risk-neutral* PDs (i.e. PDs incorporating a premium for the risk of individual defaults) and *physical* (i.e. actual) asset return correlations.⁵ Together with the data on LGD estimates, these estimated credit risk parameters deliver “predicted” tranche spreads that can be compared directly to their empirical counterparts from the CDS index market.

The two sets of estimated credit risk parameters relate to different aspects of the probability distribution of joint defaults and, as a result, have different impacts on predicted tranche spreads. On the one hand, an increase in PDs entails higher expected default rates, which raises the spreads of all tranches. On the other hand, a change in asset return correlations affects the shape of the probability distribution of joint defaults and, thus, has different pricing impacts across tranches. The reason is that an individual tranche spread relates only to a particular segment of this distribution.

We find that the index and single-name CDS markets employ similar risk-neutral PDs.

(2001) and Crosbie (2005) for details.

⁴Zhou (2001) attempts to relate the two approaches to estimating the probability distribution of joint defaults. Specifically, he studies analytically the link between asset return and default correlations in a first-passage credit risk model.

⁵In comparison to bond spreads, CDS spreads are widely considered as embedding less noisy information about market valuation of default risk. This is because CDS spreads respond more quickly to changes in credit conditions (Blanco et al., 2005; Zhu, 2006) and are affected to a lesser extent by non-credit factors (Longstaff et al., 2005).

There is a close match between the cross-sectional average of our PD estimates and the average PD implied by spreads on the overall (or single-tranche) CDS index. Furthermore, this match is stable over time and underpins the finding that PDs implied by the single-name CDS market explain to a large extent the intertemporal evolution of empirical tranche spreads of the CDS index.

At the same time, predicted spreads fall short of accounting for their empirical counterparts tranche by tranche. For instance, observed spreads imply prices of protection against catastrophe credit events (i.e. against a large number of defaults) that are much higher than those implied by predicted spreads. In addition, observed prices of protection against a small number of defaults are lower than predicted. This finding implies that the *risk-neutral* asset return correlation, which is used for pricing portfolio credit risk, is higher than the physical correlation underpinning predicted spreads.

The wedge between risk-neutral and physical asset return correlations suggests the presence of a correlation risk premium. We measure this premium by subtracting a homogenized version of the physical asset return correlations, which stays constant across pairs of firms and fits predicted tranche spreads as closely as possible on each day, from the corresponding risk-neutral correlation, which is fitted to observed tranche spreads. The time average of the correlation risk premium amounts to 66% of the time average of the homogenized physical correlations (which equals 0.15).

We also study the joint intertemporal behavior of the correlation risk premium and the homogenized physical correlation. First, we find that the premium tends to decrease when the contemporaneous realization of physical correlations increases. Thus, when asset return correlations are higher, the market tends to anticipate lower upward correlation risk or have greater appetite for correlation risk. Second, we find that a positive change in the correlation risk premium is associated with a future increase in physical correlations. This suggests that either the market tends to price correctly correlation risk or changes in the attitude towards such risk feed back into market outcomes down the road.

By analyzing the pricing of portfolio credit risk on the basis of information from the single-name credit market, this paper stands apart from related literature that fits flexible models to observed spreads of CDS index tranches. Longstaff and Rajan (2006), for example, use a multi-factor model of loan losses and derive that three credit risk factors – two of which occur with a low probability but have industry- and economy-wide impacts – are needed in order to explain fully tranche spreads of the CDS index considered here. Pursuing a similar objective, Kalemanova et al. (2007) and Moosbrucker (2006) demonstrate that tranche spreads of CDS indices are consistent with Lévy processes driving default trigger variables. In a different exercise, Hull et al. (2006) construct “implied correlations” – which are similar to the risk-neutral correlations derived here – under various modeling assumptions and explore how such assumptions can help account for empirical tranche spreads. However, in contrast to the analysis here, all these papers do not draw parallels between different credit markets, do not estimate physical asset return correlations from data and do not identify a correlation risk premium.

Such a premium is studied in the context of the equity market by Driessen et al. (2006). Specifically, this paper finds that the risk-neutral correlation implied by equity-index options is on average 63% higher than the physical correlation estimated from the corresponding single-name market. This result is strikingly in line with our findings in the context of credit markets.

The remainder of the paper is organized as follows. Section 2 outlines the structure of the CDS index market and explains how index tranches are priced. Then, Section 3 explains how we construct predicted spreads of index tranches on the basis of data from the single-name CDS market. Section 4 describes the data used in this paper. Section 5 compares predicted with observed spreads in the index market. The messages of this comparison, which unveils inter alia the presence of a correlation risk premium, are discussed in Section 6. The final section concludes.

2 The CDS index market

The market for a CDS index, which allows traders to buy and sell protection against portfolio credit risk, delivers two sets of prices. The first set is a time series of *single-tranche spreads*, which are effectively the prices of protecting the *entire* notional amount of the index against losses caused by defaults of the entities in this index. Under risk neutrality, single-tranche spreads reveal the market's *expectation* of default losses but are insensitive to the market's perception of and attitude towards the probability of default clustering.

These perception and attitude do affect, however, the second set of prices, which comprises several time series of *multi-tranche spreads*. Each time series consists of the effective prices of protection against a particular range (or “tranche”) of credit losses on the notional amount of the index. For example, the tranche relating to the first losses – and, thus, carrying the highest level of credit risk – is known as the equity tranche. If none of the entities in the index defaults, the investor in this tranche (i.e. the protection seller) receives quarterly a fixed premium payment (or “spread”) on the tranche's principal, which is typically 3% of the total notional amount of the index. If defaults occur, this investor stands ready to compensate its counterparty (i.e. the protection buyer) for any credit losses that do not exceed the outstanding principal of the equity tranche. At the same time, this principal and the associated premium payments are reduced for the remainder of the contract's life in order to reflect ongoing credit losses.⁶ Similarly, an investor in the so-called mezzanine tranche is typically responsible only for losses between 3% and 7% of the total notional amount, while investors in the two senior and two super-senior tranches are responsible only for losses between 7% and 10%, 10% and 15%, 15% and 30%, and 30% and 100% of the total notional amount, respectively. Thus, the higher the seniority of the tranche, the less likely it is that the corresponding investor will need to make payments to the protection buyer.

⁶For the CDS index contract we consider below, a default triggers an immediate adjustment to the payments by the protection seller and buyer. In our calculations, however, we impose the simplifying assumption that such adjustments are made quarterly.

A multi-tranche spread is of great use to the analysis in this paper because it pertains to a particular segment of the probability distribution of defaults and, thus, reveals the market's perception of and attitude towards default clustering. To see why, observe that, in a CDS index consisting of 100 equally-weighted entities with LGDs of 50%, the spread of the equity tranche is effectively the price of protection against the first 6 defaults in the underlying portfolio. For a given expectation of default losses, weaker interdependence of defaults across entities raises the probability of there being a few (i.e. up to 6) defaults and, as a result, raises the spread of the equity tranche. Conversely, stronger interdependence of defaults increases the probability of default clustering – e.g. of there being zero or a lot of defaults – which lowers the equity tranche spread. At the same time, greater default clustering raises the spreads of the senior tranches, because (referring back to the stylized example) these spreads are the prices of protection against the 14th to the 20th and the 20th to the 30th defaults, respectively.

3 Predicted tranche spreads

In order to obtain *predicted* tranche spreads for a CDS index, it suffices to input an estimate of the probability distribution of joint defaults and data on LGDs and risk-free rates in the numerical methodology developed in Gibson (2004). For each particular tranche, this methodology delivers the expected present value of the principal, EP , and the expected present value of contingent payments, EC , made by the protection seller. Denoting the tranche spread by θ , the present value of the expected fee revenue of the protection seller, $\theta \cdot EP$, has to equal EC . The tranche spread is then calculated as:

$$\theta = \frac{EC}{EP}$$

Since this paper will rely on the Gibson (2004) methodology, the heart of the empirical exercise is the construction of the probability distribution of joint defaults. Under typical assumptions on the stochastic distribution of borrowers' asset returns (i.e. Gaussian or

Student-t), the probability distribution of joint defaults has two key components: the PDs of constituent entities and the corresponding asset return correlations. The rest of this section describes how we estimate *risk-neutral* PDs and *physical* asset return correlations from single-name CDS spreads.

3.1 Estimating risk-neutral PDs

In order to uncover risk-neutral PDs from single-name CDS spreads, we follow the framework of Duffie (1999). In a typical single-name CDS contract – written on firm i at date t – the protection buyer agrees to make constant periodic premium payments – determined by the CDS spread $s_{i,t}$ – to the protection seller until the contract matures – at time $t + T$ – or a default occurs, whichever happens first. If a default occurs before $t + T$, the protection seller compensates the protection buyer for the realized credit loss.

Under market clearing, the present value of CDS premium payments (the left-hand side of equation (1)) has to equal the present value of protection payments (the right-hand side):

$$s_{i,t} \int_t^{t+T} e^{-r_\tau \tau} \Gamma_{i,\tau} d\tau = LGD_{i,t} \int_t^{t+T} e^{-r_\tau \tau} q_{i,\tau} d\tau \quad (1)$$

where r_τ stands for the risk-free rate of return, $q_{i,\tau}$ denotes the (annualized) unconditional risk-neutral default intensity of borrower i , $\Gamma_{i,\tau} \equiv 1 - \int_0^\tau q_{i,v} dv$ is the associated risk-neutral survival probability over the following τ years, and $LGD_{i,t} \in [0, 1]$ is the date- t expectation of loss given default.⁷ Under the standard simplifying assumptions that r_τ and $q_{i,\tau}$ are expected to be constant over time, equation (1) implies that the one-year *risk-neutral* PD equals:

$$PD_{i,t}(1) = q_{i,t} = \frac{a_t s_{i,t}}{a_t LGD_{i,t} + b_t s_{i,t}} \quad (2)$$

where $a_t \equiv \int_t^{t+T} e^{-r_t \tau} d\tau$ and $b_t \equiv \int_t^{t+T} \tau e^{-r_t \tau} d\tau$.

We use equation (2) in order to estimate a *daily* time series of borrower-specific risk-neutral one-year PDs on the basis of time series of CDS spreads, expected LGDs and risk-free

⁷Equation (1) incorporates the assumption that LGD is independent of the variable(s) triggering default events.

rates of return.

3.2 Estimating physical asset return correlations

We model the cross-sectional interdependence of default events as driven by the correlation of entity-specific “default trigger” random variables. Each entity-specific default-trigger variable is a *one-dimensional summary* of credit quality and is extracted from the corresponding CDS spread. As such, the default-trigger variables comprise all the information that is deemed relevant, identified and processed by the single-name CDS market. Such information includes: (i) balance sheet and stock market information about the entity (i.e. information that determines credit outlook in traditional structural model à la Merton (1974)) and (ii) information about systemic “frailty” factors (of the type examined recently by Das et al., 2007; Duffie et al., 2007, in their study of default contagion).

In order to be able to draw straightforward parallels with the extant literature on portfolio credit risk (see Hull and White (2004)), we henceforth refer to the default-trigger variable as “the value of the firm’s *assets*” and incorporate it in the following model. We start by assuming that, under the risk-neutral measure, the asset value process of entity i is:

$$\frac{dV_{i,t}}{V_{i,t}} = \mu_i dt + \sigma_i dW_{i,t} \quad (3)$$

where μ_i denotes the drift, σ_i the asset volatility and the shock $W_{i,t}$ is a standard Wiener processes. Then, given a default boundary D_i , we define the distance-to-default variable $DD_{i,t} \equiv \frac{\ln V_{i,t} - \ln D_i}{\sigma_i}$. By Ito’s Lemma, $dDD_{i,t}$ has a drift $\mu_i^* = \frac{\mu_i - \sigma_i^2/2}{\sigma_i}$ and a unit variance. Postulating that entity i defaults τ years into the future if $DD_{i,\tau} < 0$, we obtain that the probability of default equals:

$$PD_{i,t}^M(\tau; DD_{i,t}, \mu_i^*) = \Phi\left(\frac{-DD_{i,t} - \tau\mu_i^*}{\sqrt{\tau}}\right) \quad (4)$$

where $\Phi(\cdot)$ stands for the standard normal CDF. Given two time series of PD estimates,

the corresponding asset return correlation equals:

$$\begin{aligned}
\rho_{ij} &\equiv \text{corr}(\Delta \ln V_{i,t}, \Delta \ln V_{j,t}) \\
&= \text{corr}(\Delta DD_{i,t}, \Delta DD_{j,t}) \\
&= \text{corr}(\Delta \Phi^{-1}(PD_{i,t}^M(1; DD_{i,t}, \mu_i^*)), \Delta \Phi^{-1}(PD_{j,t}^M(1; DD_{j,t}, \mu_j^*))) \\
&\approx \text{corr}(\Delta \Phi^{-1}(q_{i,t}), \Delta \Phi^{-1}(q_{j,t}))
\end{aligned} \tag{5}$$

where the PD horizon τ is set to 1 year and Δ denotes the first difference in discrete time.

A calculation of asset return correlations on the basis of (5) warrants several remarks. First, the model underlying this procedure assumes that the default boundary (D_i), asset volatility (σ_i) and drift (μ_i^*) remain constant over time. In order to relax this assumption, one needs to incorporate additional stochastic processes and estimate their parameters. Given the available data, the errors produced by such an estimation would be large enough to render the exercise useless. Second, the underlying model implies that $DD_{i,t}$ follow unit root processes, which is supported by the data.⁸ Third, since $\Delta \ln(V_{i,t})$ stands for an *actual* asset return, equation (5) delivers a measure of the *physical* asset return correlation. Finally, the fourth line in (5) holds only as an approximation because $q_{i,t}$ is derived in (2) under the assumption of a time-invariant default intensity, whereas $PD_{i,t}^M(\tau; DD_{i,t}, \mu_i^*)$ in equation (4) violates this assumption. Robustness checks, outlined in the following subsection and discussed further in Section 5.3.1, reveal that the approximation in the fourth line of (5) is remarkably good.

Our overall empirical procedure can be summarized as follows. For each day and pair of entities in our sample, we estimate asset return correlations on the basis of the associated time series of PD estimates over the previous six months and equation (5).⁹ Then, for each

⁸More precisely, a battery of Phillips-Perron tests fail to reject the unit-root null for 132 of the 136 distance-to-default time series we construct. In addition, a unit root process provides a reasonable approximation to the dynamics in the remaining 4 series.

⁹We also calculate time series of asset return correlations on the basis of three months of data. These alternative estimates are more volatile, both in the cross section and over time, but do not alter our main conclusions.

day in the sample, we rely on our PD and asset return correlation estimates and a Monte Carlo simulation technique (outlined in Appendix A.1) in order to generate the probability distribution of joint defaults in a given portfolio. Used as an input to the Gibson (2004) methodology (recall Section 3), this probability distribution leads to *predicted* spreads of CDS index tranches.

3.3 Alternative estimators of asset return correlations

The adopted mapping from PDs to asset return correlations, which is specified in equation (5), is essentially a short-cut solution and can be criticized for introducing inconsistency in the empirical procedure. The inconsistency crystallizes in that, as mentioned above, the last equality in (5) holds only as an approximation. Here, we propose two alternative correlation estimators that circumvent this inconsistency.

The first alternative is to estimate asset return correlation on the basis of distance-to-default variables extracted directly from single-name CDS spreads. Specifically, we combine equations (1) and (4) using the fact that $q_{i,t}$ can be rewritten as $\frac{dPD_{i,t}^M}{dt}$. For a given value of the drift term μ_i^* , we then obtain a one-to-one mapping between a time series of CDS spreads and a time series of $DD_{i,t}$. This eliminates the need to resort to the fourth line in (5) and delivers asset return correlation estimated within a coherent framework.

Another alternative is to estimate asset return correlations within a coherent first-passage model, in which the PD over a given horizon equals the probability that the borrower's assets fall below some threshold at any point in time over this horizon. In a first-passage model, the probability that firm i defaults over the next τ years equals:

$$PD_{i,t}^{FP}(\tau; DD_{i,t}, \mu_i^*) = 1 - \Phi\left(\frac{DD_{i,t} + \tau\mu_i^*}{\sqrt{\tau}}\right) + \exp(-2DD_{i,t} \cdot \mu_i^*) \Phi\left(\frac{-DD_{i,t} + \tau\mu_i^*}{\sqrt{\tau}}\right) \quad (6)$$

We combine equation (6) with (1) to derive a mapping between single-name CDS spreads and the distance-to-default variable. This mapping leads to another set of asset return correlations.

In order to carry out either of the two alternative procedures, we experiment with several values of μ_i^* suggested by the literature, keeping each of these values constant in the cross section. We find that the correlation estimates are virtually insensitive to the exact value of μ_i^* and below report results only for $\mu_i^* = 0$.

4 Data

We work with two large data sets, which are described in this section. In addition to these data, we obtain 5-year Treasury rates from Bloomberg, which we use to proxy for the risk-free rate of return.

The first data set is provided by JP Morgan Chase and pertains to 5-year contracts written on the CDS index Dow Jones CDX North America investment-grade index (CDX.NA.IG.5Y). These standardized contracts are highly liquid on the secondary market. We use single-tranche spreads for the “on-the-run” CDX.NA.IG.5Y index, as well as spreads for the equity, mezzanine and two senior tranches of the same index.¹⁰ At each point in time the CDS index consists of 125 entities that represent major industrial sectors and are actively traded in the single-name CDS market as well. All entities have equal shares in the total notional principal of the index. The composition of the index is updated semi-annually – in a new release – in order to reflect events such as defaults, rating changes, and mergers and acquisitions. We consider three releases of the CDX.NA.IG.5Y index, launched respectively on November 13, 2003, March 23, 2004 and September 21, 2004.¹¹ The total number of entities that appear in at least one of these releases is 136.

The second data set pertains to the single-name CDS market and is provided by Markit, which has constructed a network of leading market participants who contribute pricing information across several thousand credits on a daily basis. Markit aggregates the information

¹⁰We abstract from the two super-senior tranches because the spreads on these tranches are likely to be affected substantially by non-credit factors, such as administrative costs and a liquidity premium. Although the analysis of such factors is important, it is beyond the scope of this paper.

¹¹This period is free of episodes of market distress and is, thus, likely to feature relatively stable credit risk parameters. This conjecture is largely in line with our empirical findings.

it receives and releases daily “consensus” CDS spreads and LGD estimates for each credit in its database. In line with the contractual terms of the CDX.NA.IG.5Y index, we use time series of 5-year senior unsecured single-name CDS spreads associated with the no-restructuring clause (see ISDA, 2003) and denominated in US dollars. For the period from April 24, 2003 to September 27, 2005, we download CDS spreads and the associated LGD estimates for all 136 entities that belonged to at least one of the CDS index releases we consider.

The data on single-name CDS spreads warrant two remarks. First, the spreads in our sample do appear highly responsive to changes in credit conditions. This property surfaces in that market staleness – defined as the realization of the same spread on two consecutive days – characterizes only 13% of our sample. By contrast, it is typical for spreads on the corporate bond market to be stale for weeks. The high quality of the CDS spreads prompts us to use daily data in computing asset return correlations. Second, the LGDs provided by Markit reflect market participants’ consensus view on *expected* losses, and therefore need not match *realized* losses. The reported LGDs exhibit little cross-sectional difference and time variation (see Table 1). In the cross section of 136 time averages of LGDs, the 1st and 99th percentiles equal 60% and 63% respectively. In addition, the time series of cross-sectional averages of LGDs fluctuates within a similarly narrow band. In the light of this and in order to eliminate potential noise in the LGD data, we set LGDs to be the same across entities and smooth the resulting time series via an HP filter with a parameter $\lambda = 64000$.

5 Empirical findings

In this section, we analyze predicted tranche spreads. These spreads are underpinned by two sets of credit risk parameter estimates – risk-neutral PDs and physical asset return correlations – which we report in Section 5.1. In Section 5.2, we compare the predicted tranche spreads with empirical tranche spreads of the CDS index and find substantial differences. In Section 5.3, we argue that these differences are robust to a number of alternative correlation

estimates and model specifications that might have been adopted by the market.

5.1 PD and correlation estimates

Using the methodology described in Sections 3.1 and 3.2, we estimate time series of *risk-neutral* PDs and *physical* asset return correlations, respectively. Table 1 reports summary statistics for the two sets of credit risk parameters. The statistics relate to all 136 entities that belong to at least one of the first three releases of the CDS index CDX.NA.IG.5Y.

In general, both PDs and correlations change substantially over time. For example, the cross-sectional average of PDs peaks in August 2004 at about 1.2% and reaches its lowest levels, roughly 0.75%, towards the end of the sample period. As for asset return correlations, they are on a general upward trend during the sample period, increasing from around 7% in late 2003 to 17.5% in March 2005.

Despite the fact that all sample firms are investment-grade entities, the cross-sectional dispersion in the two sets of credit risk parameters is quite pronounced. Across the 136 entities, the time average of risk-neutral PD estimates has a mean of 94 basis points and a standard deviation of 78 basis points, with the maximum level (425 basis points) being eighteen times higher than the minimum level (23 basis points). Similarly, there is marked heterogeneity across pairwise correlation estimates. Correlations can be as high as 80-90% for firms in the same business area, as is the case of Ford Motors Credit Company and General Motors Acceptance Corporation. At the other extreme, there are negative correlations, such as the one between Intel and Amerada Hess Corporation. In principle, as argued by Hull and White (2004), heterogeneity across PDs and pairwise correlations can have important implications for the probability distribution of joint defaults and, by extension, for the pricing of portfolio credit risk instruments.

5.2 Comparing predicted with observed tranche spreads

In order to construct predicted spreads for the four (i.e. equity, mezzanine and two senior) tranches of the CDS index CDX.NA.IG.5Y, we use our LGD data and the PD and

correlation estimates obtained from the single-name CDS market (recall Sections 3.1 and 3.2 for the estimation procedures). In applying the Gibson (2004) methodology, we assume that all random variables are normal. The resulting predicted tranche spreads – which we dub “baseline” spreads in order to distinguish them from alternatives reported in Section 5.3.2 – can be compared directly with the corresponding spreads observed in the CDS index market. Figure 1 plots the time series of the two sets of tranche spreads and Table 2 reports summary statistics (baseline predicted spreads: first row in each panel; observed spreads: last row in the top panel).

Baseline predicted spreads differ substantially from their counterparts in the data. For the equity tranche, predicted spreads are too high over most of the sample and over-predict observed spreads by 10% on average. The differences are more pronounced in the two senior tranches where predicted spreads are too low over the entire sample period and under-predict observed spreads by 37% for the first and 63% for the second senior tranche. As for the mezzanine tranche, the predicted spreads match well the observed spreads on average (306.3 versus 303.9 basis points), but this result masks large pricing differences on individual days. In fact, the mean absolute percentage error between the predicted and observed spreads for the mezzanine tranche is substantial, averaging 15.5% of observed spreads.

At the same time, predicted tranche spreads exhibit statistically significant, albeit not quite large, explanatory power for the evolution of observed spreads over time. Regression results, reported in Table 3 (Panel A), reveal that changes in predicted tranche spreads account for 45% or more of the variability of changes in observed spreads for the equity, mezzanine and first senior tranches.¹² The goodness-of-fit measure drops to 16% for the second senior tranche.¹³

¹²Tables 3 and 4 report two versions of each regression: one version uses the variables in levels and the other one in first differences. While the coefficient estimates in the “levels” regressions are easier to interpret, the associated significance levels and goodness-of-fit measures are questionable because one cannot reject the unit-root hypothesis for both dependent and explanatory variables. This issue is addressed by the “first difference” regressions.

¹³All regressions are based on weekly averages of the dependent and explanatory variables. This restricts the impact of market microstructure noise, which is likely to surface at high frequencies. Running the regressions with daily variables tends to preserve the sign and significance of the coefficients but, not surprisingly,

5.3 Robustness checks

It is possible that the wedge between observed and baseline predicted spreads is due to the fact that market participants use credit risk parameter estimates or rely on pricing frameworks that are different from the ones we have considered so far. A series of robustness checks reveals that this conjecture is *not* borne out.

5.3.1 Robustness of predicted spreads to alternative correlation estimates

Following the procedures outlined in Section 3.3, we examine two alternative estimates of asset return correlations that allow for circumventing the approximation in the fourth line of (5). It turns out that the sample average of asset return correlations under the first (second) alternative procedure is 12.39% (12.41%), while the corresponding average of the original estimates – which underpin the results reported in Sections 5.1 and 5.2 – equals 12.42%. A comparison across the time series of cross-sectional averages of original and alternative correlations reveals a maximum difference of 0.3 percentage points. In turn, considering the cross section of the time averages of correlation estimates, we find that the maximum difference is roughly 0.8 percentage points. Unreported calculations reveal that such differences have a negligible impact on predicted tranche spreads.

In the rest of the paper, we keep on working with the original estimates of PDs and asset return correlations – given by equation (2) and the fourth line of expression (5), respectively – for the following reasons. First, as stated in the previous paragraph, the approximation in the fourth line of (5) generates negligible errors in our correlation estimates. Second, in adopting this approximation, we work with PD estimates that do not depend on an estimate of the drift parameter μ^* . To see this, compare the adopted equation (2) with the alternative specifications in (4) and (6). This is important because, unlike the correlation estimates, the PD estimates are sensitive to the value of μ^* but the available data do not allow us to pin down this value. Finally, it is reportedly popular market practice to adopt the

lowers the goodness-of-fit measures (i.e. adjusted R^2).

constant-default-intensity assumption of equation (2) in pricing credit derivatives products. If this is true, then choosing a non-flat term structure of default intensities – as implied by the alternative specification in (4) and (6) – would be supported by structural credit risk models but would introduce errors in predicted tranche spreads.

5.3.2 Robustness of predicted spreads to alternative pricing frameworks

In this subsection, we report the results of three robustness checks, which explore the implications of different pricing frameworks.

For the first check, we calculate predicted spreads after removing the heterogeneity in PDs and/or pairwise correlations. This proxies for a scenario in which investors in the CDS index market lack information on firm-specific PDs or are unable to derive asset return correlation for each pair of firms. The results of this exercise, summarized in Table 2 (rows 2-4 in each panel), reveal limited revisions to predicted tranche spreads. In fact, shutting off the cross-sectional difference in correlations raises (lowers) the predicted spreads for the equity (two senior) tranches, leading to an even worse match with observed spreads.

Our second exercise is motivated by market commentary, which refers regularly to common factor models of asset returns.¹⁴ We investigate whether the use of such a model for pricing purposes could have a material impact on predicted tranche spreads. We start by estimating a one-factor structure of the estimated correlation matrix (see Appendix B) and then employ this structure in a Gaussian copula to derive tranche spreads (see Appendix A.2). As reported in Table 2 (row 5 in each panel), adopting a one-factor model has a limited impact on our results and, if anything, renders the difference between the two sets of spreads even larger.¹⁵

The last robustness check examines the pricing implications of alternative assumptions

¹⁴Also see Collin-Dufresne et al. (2003), Das et al. (2007) and Giesecke (2004) for discussions of why the assets of different firms may be driven by common factors.

¹⁵The one-factor approximation matches the mean of the original correlation matrix extremely well – producing an average discrepancy of 12 basis points – but tends to underestimate the dispersion in correlation coefficients. To address this issue, we also estimate a two- and a three-factor correlation structures. These generalizations improve substantially the fit of the original correlation matrix but entail negligible revisions of predicted tranche spreads.

regarding the distribution of asset returns. Researchers (see Hull and White (2004)) and market practitioners have argued for the use of *Student-t* distributions, which perform better than a Gaussian distribution in accounting for the fat tails of asset returns observed in the data. In the light of this, we use a one-factor model of asset returns and assume that the common factor and/or idiosyncratic factors follow a *Student-t* distribution with four degrees of freedom.¹⁶ To derive predicted tranche spreads under this assumption, we employ the so-called *t*-copula. The results are reported in Table 2 (rows 6-7 in each panel). The fat tails implied by a *Student-t* distribution raise the probability of a large number of defaults, which reduces somewhat the gap between predicted and observed spreads for the two senior tranches. However, this is at the cost of larger gaps for the equity and mezzanine tranches.

6 The pricing of portfolio credit risk

This section delves further into the differences between baseline predicted and observed tranche spreads. Specifically, we investigate whether the PDs and asset return correlations we extract from the single-name CDS market, albeit correctly estimated, differ from the credit risk parameters employed by CDS index market.

6.1 Market estimates of PDs

The PDs of individual borrowers are an important factor in the pricing of portfolio credit risk. Namely, an overall rise in PDs translates directly into higher portfolio credit risk, which raises the spreads for all index tranches. Thus, it is natural to ask if the PD estimates implied by the single-name CDS market and underpinning predicted tranche spreads are different from the PD estimates used in the CDS index market.

Even though we are not able to pin down the entity-specific PDs used by market participants in the CDS index market, we do extract the cross-sectional averages of these PDs.

¹⁶On the basis of an ad hoc value for asset return correlations, Hull and White (2004) find that assuming *Student-t* distributions with four degrees of freedom for both the common and idiosyncratic factors helps one account well for the tranche spreads of Dow Jones iTraxx EUR 5y (the European counterpart of the CDX.NA.IG.5Y index).

This is done on the basis of a time series of single-tranche spreads, which reveal expected credit losses (recall Section 2), and data on market expectations of LGDs. The average (risk-neutral) PDs implied by the single-tranche spreads are plotted in Figure 2 alongside average PDs implied by the single-name CDS spreads. The two time series differ on average by only 1.21 basis points, which is roughly 1.4% of the average PD implied by single-name CDS spreads. Importantly, these differences do not have a material impact on predicted tranche spreads (see Table 2, row 8 in each panel) and we can conclude that the PD estimates used in the index market are consistent with the PDs embedded in the single-name CDS market.

Moreover, the close match between the two series of average PDs drives the similarity of the time paths of predicted and observed tranche spreads (see Figure 1). To substantiate this claim, we conduct regression analysis. As Table 3 shows, PDs possess significant explanatory power for the level and changes in both predicted and observed tranche spreads (panels B to D). Strikingly, the goodness-of-fit measures indicate that changes in observed spreads are explained better by changes in PD estimates than by changes in predicted spreads. This is evidence that not PDs but another estimated component of predicted spreads may be weakening their co-movement with observed spreads. This other component is the correlation of asset returns, which we analyze in the following subsection.

6.2 Risk-neutral versus physical asset return correlation

In calculating predicted spreads, we have assumed that the *physical* correlation of asset returns implied by the single-name CDS market is used for pricing in the index market. Such an assumption is likely to be violated in practice. In fact, Driessen et al. (2006) find strong evidence that the *risk-neutral* correlations used by investors in equity-index options are substantially higher than the physical correlations implied by the single-name equity market. In order to examine whether a similar phenomenon exists in the credit market, we compare a homogenized version of our physical correlation estimates with the corresponding

risk-neutral correlations, which we estimate on the basis of observed tranche spreads.

6.2.1 Correlation and the average level of tranche spreads

Risk-neutral asset return correlations that are larger than their physical counterparts can help explain the gap between predicted and observed spreads. This is suggested by Figure 3, which illustrates the sensitivity of spreads to changes in (homogenous) correlation coefficients. When the correlation coefficient increases, the default of a particular firm is more likely to be driven by the deterioration of common risk factors and hence is more likely to be accompanied by defaults of other firms. This raises (lowers) the probability of a large (small) number of defaults. As a result, higher asset return correlations lead to higher (lower) spreads for the senior (equity) tranches of a CDS index but have an indeterminate impact on the mezzanine tranche. In the light of Figure 1, this implies that replacing the physical correlations underpinning predicted spreads with higher risk-neutral correlations would help us account to a greater extent for observed spreads.

6.2.2 Correlation and changes in tranche spreads over time

Differences between the physical correlation of asset returns and its risk-neutral counterpart could also account for the intertemporal evolution of the differences between predicted and observed tranche spreads. To see this, recall Figure 3 which implies that, if physical correlations match closely risk-neutral ones, then an increase in their level should lower equity tranche spreads and raise senior tranche spreads. However, regression results reported in Table 3 (panel C) show that physical correlations enter either with the wrong sign (“level” regressions) or are statistically insignificant (“first differences” regressions) as explanatory variables of observed spreads .

Table 4 provides further evidence that physical correlations, which underpin predicted spreads, differ from the risk-neutral correlations behind observed spreads. This table contains results from regressions of the gap between predicted and observed tranche spreads and reveals that physical correlation of asset returns is the most important driver of this

gap. More specifically, the physical correlation enters with a statistically significant positive coefficient the regressions for the two senior tranches. This reveals that increases in the physical correlation, which raise the predicted spread for the senior tranches, are not associated with a concurrent rise in observed spreads for those tranches. A symmetric reasoning applies to the “equity tranche” regressions.¹⁷

6.2.3 Correlation risk premium

This subsection describes our measure of a correlation risk premium in the CDS index market. This premium reflects the market-determined compensation for bearing the risk that physical asset return correlations may *increase* above their current level.

Thus, a change in the correlation risk premium has different impacts on the spreads of different index tranches. An increase in this premium signifies a higher value of protection against a large number of defaults, which inflates the spreads of senior tranches. Symmetrically, an increase in the correlation risk premium is tantamount to a lower value of protection against just a few defaults, which depresses the spread of the equity tranche.

In order to derive the correlation risk premium, we start by extracting a time series of risk-neutral correlations from observed tranche spreads in the index market. On each day, the risk-neutral correlation has two general properties. First, it is constrained to be the same for all pairs of entities and is used to calculate four (i.e. equity, mezzanine and two senior) spreads. Second, the specific value of the risk-neutral correlation is picked so that it minimizes the mean squared *percentage* difference between the “fitted” spreads it implies and the corresponding observed spreads.¹⁸ Importantly, our estimate of the risk-neutral correlations incorporates risk-neutral PDs that vary in the cross section and is, thus, different from the popular “implied” correlations that are underpinned by homogenous

¹⁷Table 4 also reveals that PD coefficients are statistically significant and negative in the regressions of the errors in the predicted spreads for senior tranches. Although PDs have a low explanatory power in these regressions, negative coefficients are puzzling because they seem to suggest that a rise in PDs lowers the tranche spread. Background analysis reveals that this result is driven by a non-linear interaction between firm-specific PDs and pairwise asset return correlations.

¹⁸By minimizing percentage differences, we effectively equalize the units of the pricing errors across different tranches.

PDs.¹⁹

Not surprisingly, fitted spreads perform better than predicted ones in matching observed tranche spreads. This can be visualized by comparing Figures 4 and 1. The overall improvement in the match, which transpires in the two senior tranches, reveals that the risk-neutral correlations we estimate capture to a large extent the market’s valuation of the risk of default clustering. At the same time, the relatively poorer match in the equity and mezzanine tranches²⁰ suggests that asset return correlations miss some information that is necessary for a full account of all tranche spreads of the CDS index.²¹ However, our goal in this paper is not to provide such a full account but to employ data from the single-name credit market for the understanding of the pricing of portfolio credit risk. This is what the estimates of risk-neutral correlations allow us to do.

The second component of the correlation risk premium is a “homogenized” physical correlation. This correlation is defined in the same way as its risk-neutral counterpart but is estimated from predicted spreads. In other words, on each day, the homogenized physical correlation condenses the information contained in the physical correlation matrix that we estimate from single-name CDS spreads. The homogenized physical correlation fits predicted spreads extremely well across all tranches, entailing an average absolute percentage error of less than 2%.²²

We define the correlation risk premium as the risk-neutral minus the homogenized physical correlation. The time series of the premium is plotted in Figure 5, alongside the time

¹⁹For the estimation of risk-neutral correlations, we use the same LGDs and risk-free rates that underpin our estimates of risk-neutral PDs and physical asset return correlations (see Sections 3.1 and 3.2).

²⁰Specifically, the absolute percentage errors between fitted and observed spreads average 13%, 18%, 4% and 7% for the equity, mezzanine and the two senior tranches, respectively. The differences in the match across tranches is a manifestation of the so-called “correlation smile” phenomenon (see Amato and Gyntelberg, 2005, for a review).

²¹To attain a better match, researchers have – explicitly or implicitly – modelled default trigger variables as driven by more general processes than the ones studied in this paper (see Longstaff and Rajan (2006), Kalemánova et al. (2007) and Moosbrucker (2006), for example). These general processes effectively provide a greater number of degrees of freedom.

²²The simple cross-sectional average of physical correlation coefficients is on average 2.9 percentage points lower than the homogenized correlation we work with. Switching between the two physical correlation alternatives affects immaterially the intertemporal properties of the correlation risk premium. These properties are examined in Section 6.2.4.

series of the two correlation estimates. During the sample period, the average of the correlation risk premium equals 10.1%, i.e. 66% of the homogenized physical correlation, which averages 15.3%. This result matches almost exactly the finding of Driessen et al. (2006) that, between 1996 and 2003, the risk-neutral correlation implied by equity-index options is 63% higher than its physical counterpart.

6.2.4 Intertemporal links between the correlation risk premium and physical correlations

Figure 5 (bottom panel) highlights the negative relationship between the correlation risk premium and the physical correlation. For instance, the premium peaks at the end of 2003, exactly when the physical correlation attains its lowest levels. More generally, the correlation risk premium is on a downward path over the entire sample period, while the physical correlation is on an upward path. In the rest of this subsection, we examine more closely the joint behavior of the correlation risk premium and the physical correlation.

The correlation risk premium is constructed in this paper on the basis of one forward-looking and one backward-looking variable. On the one hand, the risk-neutral correlations are extracted directly from market spreads and, thus, reflect investors' forward looking perceptions and attitude toward correlation risk. On the other hand, realized physical correlations are estimated from historical data, which renders them backward looking. Importantly, the difference between these two variables, i.e. the correlation risk premium, has three general components. The first component is due to a discrepancy between the currently realized value of physical correlations and the market's expectation of these correlations over the life of CDS contracts. The second component reflects the risk that future realized correlations are higher than their expected value. Finally, the third component is driven by market participants' appetite for correlation risk.

As a result, it is reasonable to expect an interesting joint behaviour of realized physical correlations and the correlation risk premium. If historical data have a bearing on (i) market perceptions of the expected level and/or volatility of future asset return correlations

or (ii) the prevailing attitudes toward risk, then we should observe contemporaneous comovement between the correlation risk premium and physical correlations. In addition, the correlation risk premium may help predict future changes in realized physical correlations if (i) the market's forward-looking perceptions are accurate or (ii) the market's attitude towards correlation risk feeds back into prices down the road. In order to investigate these hypotheses, we conduct two regression exercises.

In the first exercise, we investigate to what extent changes in the correlation risk premium can be accounted for by concurrent changes in the physical correlation. The investigation is in the form of a simple regression, in which we incorporate two additional explanatory variables: lagged changes in the correlation risk premium and changes in the cross-sectional averages of PDs. As reported in Table 5, changes in the physical correlation is the only significant regressor and accounts for roughly 40% of the time variation in changes of the correlation risk premium. By contrast, changes in PDs – a standard proxy for the credit cycle – exhibit no statistical significance.

Table 5 also reveals that the correlation risk premium tends to decline when the concurrent realization of the physical correlation rises. This finding sheds light on market perceptions of and attitude towards correlation risk. First, the finding suggests that investors may perceive mean reversion in the correlation level. Second, it is possible that investors perceive a negative relationship between the level and the volatility of correlation over time.²³ Thus, when our backward-looking estimate of physical correlations is high, either perception induces investors to demand lower compensation for correlation risk, as they attribute a lower probability to further increases in the correlation level. In addition, a third possibility is that investors with higher tolerance for correlation risk tend to dominate the market when correlation risk is high. The concurrent drop of the correlation risk premium would then be a result of the lower price of risk demanded by such investors.

In the second regression exercise, we examine whether the correlation risk premium

²³Unfortunately, owing to the short sample period, we are not able to test for these intertemporal properties of the physical correlation of asset returns.

helps predict future realizations of the physical correlation of asset returns. The results are reported in Table 6. In contrast to the correlation risk premium, changes in the physical correlations do exhibit serial correlation, which is captured by the significant positive coefficient of the lagged dependent variable in the regressions. More importantly, the correlation risk premium moves systematically in the same direction as future realizations of the physical correlation. This finding can be interpreted as an indication that market expectations tend to be validated *ex post*. Alternatively, it is possible that a rise in the aversion to correlation risk, which inflates the associated premium, feeds back into market prices and raises the correlation of asset returns in future periods.

7 Conclusion

This paper has analyzed the pricing of portfolio credit risk in the CDS index market on the basis of information obtained from the single-name CDS market. This analysis has revealed the existence of a large correlation risk premium, defined as the difference between the *risk-neutral* correlation of asset returns, used for pricing in the index market, and the corresponding *physical* correlation. This premium changes over time, co-varying negatively with current estimates of the physical correlation and positively with future realizations of this correlation. The intertemporal behavior of the correlation risk premium, which would be usefully revisited by future research on the basis of longer data series, reveals information about market perceptions of and attitude towards correlation risk.

References

- Altman, Edward and Vellore Kishore (1996), “Almost everything you want to know about recoveries on default bonds,” *Financial Analysts Journal*, vol. 52, 57–64.
- Amato, Jeffrey and Jacob Gyntelberg (2005), “CDS index tranches and the pricing of credit risk correlations,” *BIS Quarterly Review*, pages 73–87.
- Andersen, Leif and Jakob Sidenius (2005), “Extensions to the Gaussian copula: random recovery and random factor loadings,” *Journal of Credit Risk*, vol. 1, 29–70.

- Andersen, Leif, Jakob Sidenius, and Susanta Basu (2003), “All your hedges in one basket,” *Risk*, vol. 16, 67–72.
- BCBS (2004), “Credit Risk Transfer,” *Basel Committee on Banking Supervision Joint Forum*.
- Blanco, Roberto, Simon Brennan, and Ian W. March (2005), “An empirical analysis of the dynamic relationship between investment-grade bonds and credit default swaps,” *Journal of Finance*, vol. 60, 2255–2281.
- Carey, Mark and Michael Gordy (2006), “The bank as grim reaper: debt composition and recoveries on defaulted debt,” *Working Paper*.
- Collin-Dufresne, Pierre, Robert Goldstein, and Jean Helwege (2003), “Is credit event risk priced? Modeling contagion via the updating of beliefs,” *Working Paper*, Carnegie Mellon University.
- Covitz, Dan and Song Han (2004), “An empirical analysis of bond recovery rates: exploring a structural view of default,” *Federal Reserve Board Working Paper*.
- Crosbie, Peter (2005), “Global correlation factor structure: modeling methodology,” *Moody’s KMV Documents*.
- Daniels, Robert, Siem Jan Koopman, and Andrée Lucas (2005), “A non-Gaussian panel time series model for estimating and decomposing default risk,” *Working Paper*.
- Das, Ashish and Shota Ishii (2001), “Methods for calculating asset correlations: a technical note,” *Moody’s KMV Documents*.
- Das, Sanjiv, Darrell Duffie, Nikunj Kapadia, and Leandro Saita (2007), “Common failings: how corporate defaults are correlated,” *Journal of Finance*.
- Demey, Paul, Jean-Frédéric Jouanin, Céline Roget, and Thierry Roncalli (2004), “Maximum likelihood estimate of default correlations,” *Risk*, vol. 17, 104–108.
- Driessen, Joost, Pascal Maenhout, and Grigory Vilkov (2006), “Option-implied correlation and the price of correlation risk,” *Working Paper*.
- Duffie, Darrell (1999), “Credit swap valuation,” *Financial Analysts Journal*, vol. 55, 73–87.
- Duffie, Darrell, Leandro Saita, and Ke Wang (2007), “Multi-period corporate default prediction with stochastic covariates,” *Journal of Financial Economics*, vol. 83, 635–665.
- Duffie, Darrell and Kenneth Singleton (2003), *Credit Risk: Pricing, Measurement, and Management*, Princeton University Press.

- Gibson, Michael (2004), “Understanding the risk of synthetic CDOs,” *Finance and Economics Discussion Series*.
- Giesecke, Ikey (2004), “Correlated default with incomplete information,” *Journal of Banking and Finance*, vol. 28, 1521–1545.
- Hull, John, Mirela Predescu, and Alan White (2006), “The valuation of correlation-dependent credit derivatives using a structural model,” *Working Paper*.
- Hull, John and Alan White (2004), “Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation,” *Journal of Derivatives*, vol. 12, 8–23.
- ISDA (2003), “International Swaps and Derivatives Association (ISDA) credit derivatives definitions,” *Supplements and Commentaries*.
- Jarrow, Robert (2001), “Default parameter estimation using market prices,” *Financial Analysts Journal*, vol. 57, 75–92.
- Jarrow, Robert and Donald van Deventer (2005), “Estimating default correlations using a reduced-form model,” *Risk*, vol. 18, 83–88.
- Kalemanova, Anna, Bernd Schmid, and Ralf Werner (2007), “The normal inverse Gaussian distribution for synthetic CDO pricing,” *Journal of Derivatives*, vol. 14.
- Laurent, Jean-Paul and Jon Gregory (2005), “Basket default swaps, CDOs and factor copulas,” *Journal of Risk*, vol. 7, 103–122.
- Li, David (2000), “On default correlation: a copula approach,” *Journal of Fixed Income*, vol. 9, 43–54.
- Longstaff, Francis, Sanjay Mithal, and Eric Neis (2005), “Corporate yield spreads: default risk or liquidity? New evidence from the credit-default-swap market,” *Journal of Finance*, vol. 60, 2213–2253.
- Longstaff, Francis and Arvind Rajan (2006), “Empirical analysis of the pricing of collateralized debt obligations,” *Working Paper*.
- Merton, Robert (1974), “On the pricing of corporate debt: the risk structure of interest rates,” *Journal of Finance*, vol. 29, 449–470.
- Moosbrucker, Thomas (2006), “Explaining the correlation smile using variance Gamma distributions,” *Journal of Fixed Income*, vol. 16, 71–86.
- Pan, Jun and Kenneth Singleton (2005), “Default and recovery implicit in the term structure of sovereign CDS spreads,” *Working Paper*.

Zhou, Chunsheng (2001), “An analysis of default correlation and multiple defaults,” *Review of Financial Studies*, vol. 14, 555–576.

Zhu, Haibin (2006), “An empirical comparison of credit spreads between the bond market and the credit default swap market,” *Journal of Financial Services Research*, vol. 29, 211–235.

Appendix

A Estimating the probability distribution of defaults

This appendix outlines two methods for estimating the probability distribution of joint defaults in a given portfolio, when PDs of individual entities and the asset return correlation across entities are known. The first method relies on Monte Carlo simulations and does not impose any restriction on the structure of the correlation matrix. The second, copula, method requires the correlation structure to be driven by a finite number of common factors.

A.1 Monte-Carlo simulation

Under this method, the probability distribution of defaults in a portfolio of N entities is derived as follows.

1. Generate N random draws x_0 from independent standard normal distributions.
2. Calculate $x = R'x_0$, where R denotes the Cholesky factor of the estimated asset return correlation matrix for the N entities.
3. Denoting the i -th member of x by x_i ($i = 1, \dots, N$) and the associated PD by $PD_{i,t}$, entity i is said to default if and only if $x_i < \Phi^{-1}(PD_{i,t})$.
4. Repeat steps 2 to 4 a large number of times to estimate the probability of $n \in \{0, \dots, N\}$ defaults.

A.2 Copula

The copula method, developed by Li (2000), Laurent and Gregory (2005) and Andersen and Sidenius (2005), relies on a common factor structure of asset returns. For simplicity, we describe the copula method assuming there is only one common factor but the method can be generalized to multiple common factors. Without loss of generality, all factors are assumed to have zero mean and unit standard deviation.

Denoting the common and idiosyncratic factors, the loading coefficient on the common factor and the unconditional PD by M_t , $Z_{i,t}$, α_i and $PD_{i,t}$, we calculate the probability distribution of joint defaults in three steps. In the first step, we calculate the conditional default probability for entity i on date t : $PD(i|M_t)$. Postulating that a default is triggered when the asset value $V_{i,t} = \alpha_i M_t + \sqrt{1 - \alpha_i^2} \cdot Z_{i,t}$ falls below some threshold, we obtain:

$$PD(i|M_t) = G\left(\frac{F^{-1}(PD_{i,t}) - \alpha_i M_t}{\sqrt{1 - \alpha_i^2}}\right) \quad (7)$$

where G and F are the cumulative distributional functions of $Z_{i,t}$ and $V_{i,t}$, respectively. In general, there need not be analytical expressions for these distributions.

The second step is to calculate the conditional probability of an arbitrary number of defaults. Suppose we know the probability of $k \in \{0, 1, \dots, K\}$ defaults in a portfolio of K entities: $p^K(k|M_t)$. Then, adding one more entity leads to the following update of the distribution of defaults:

$$\begin{aligned} p^{K+1}(0|M_t) &= p^K(0|M_t)(1 - PD(K+1|M_t)) \\ p^{K+1}(k|M_t) &= p^K(k|M_t)(1 - PD(K+1|M_t)) \\ &\quad + p^K(k-1|M_t)PD(K+1|M_t) \quad \text{for } k = 1, \dots, K \\ p^{K+1}(K+1|M_t) &= p^K(K|M_t)PD(K+1|M_t) \end{aligned}$$

This recursion starts with the initial condition $p^0(0|M_t) = 1$.

The final step is to calculate the unconditional probability of k defaults in a portfolio of N entities:

$$p^N(k, t) = \int_{-\infty}^{\infty} p^N(k|M_t)\varphi(M_t)dM_t$$

where φ is the probability density function of M_t .

B Estimating a common-factor structure of asset return correlations

This appendix describes how we fit a common-factor structure to a given correlation matrix. The given correlation matrix has entries ρ_{ij} , where i and $j \in \{1, \dots, N\}$ and N is the size of the cross section. This matrix is to be approximated under the assumption that asset returns are underpinned by F common factors $M_t = [M_{1,t}, \dots, M_{F,t}]'$ and N idiosyncratic factors $Z_{i,t}$:

$$\Delta \ln(V_{i,t}) = A_i M_t + \sqrt{1 - A_i' A_i} \cdot Z_{i,t}$$

where $A_i \equiv [\alpha_{i,1}, \dots, \alpha_{i,f}, \dots, \alpha_{i,F}]$ is the vector of common factor loadings, $\alpha_{i,f} \in [-1, 1]$ and $\sum_{f=1}^F \alpha_{i,f}^2 \leq 1$. Without loss of generality, all common and idiosyncratic factors are assumed to be mutually independent and to have zero means and unit standard deviations.

We estimate the loading coefficients $\alpha_{i,f}$ ($i = 1, \dots, N$, $f = 1, \dots, F$) by minimizing the mean squared difference between the factor-driven correlation and the target correlation:

$$\min_{A_1 \dots A_N} \sum_{i=2}^N \sum_{j < i}^N (\rho_{ij} - A_i A_j')^2$$

Andersen et al. (2003) proposes an efficient algorithm to solve this optimization problem. Importantly, besides the “zero mean-unit variance” normalization, this estimation method

imposes no restriction on the distribution of the common and idiosyncratic factors.

Table 1: Risk parameter estimates

	<i>mean</i>	<i>std. dev.</i>	<i>min</i>	<i>5%</i>	<i>25%</i>	<i>50%</i>	<i>75%</i>	<i>95%</i>	<i>max</i>
LGDs (%)									
<i>Daily averages</i>	61.6	0.9	60.3	60.3	60.5	61.7	62.3	62.7	63.6
<i>Averages over time</i>	61.6	0.7	59.0	60.5	61.1	61.5	61.9	62.7	63.7
PDs (basis points)									
<i>Daily averages</i>	93.9	11.8	75.0	77.4	83.0	93.5	105.0	111.4	116.6
<i>Averages over time</i>	93.9	78.2	23.3	34.3	53.0	68.5	92.2	268.9	425.0
Pairwise correlations (%)									
<i>Daily averages</i>	12.0	2.9	6.6	7.0	10.5	12.6	13.7	16.7	17.5
<i>Averages over time</i>	12.0	7.5	-14.8	0.9	7.0	11.5	16.4	24.5	89.1

Note: The summary statistics refer to the sample period between November 13, 2003 and March 18, 2005 and reflect all 136 entities that belong to any of the first three releases of the CDS index CDX.NA.IG.5Y. In each panel, the first row is based on a time series of *daily* cross-sectional averages, whereas the second row is based on a cross section of time averages.

Table 2: Predicted versus observed tranche spreads

A. Predicted tranche spreads (averages), in basis points				
	0-3%	3-7%	7-10%	10-15%
Baseline	1849.7	306.3	68.5	15.9
no dispersion in PDs	1849.3	315.3	74.5	18.7
no dispersion in correlations	1960.7	290.3	52.0	10.0
no dispersion in PDs and correlations	1929.5	307.4	62.0	13.4
One-factor correlation structure	1915.3	289.7	62.2	14.5
<i>t</i> -copula: (4,4)	2087.0	210.5	46.2	20.2
<i>t</i> -copula: (4, ∞)	1899.7	244.8	60.1	24.5
Adjust the level of PDs	1887.8	316.3	71.9	16.9
<i>Memo:</i>				
<i>Observed tranche spreads</i>	<i>1705.4</i>	<i>303.9</i>	<i>111.1</i>	<i>45.5</i>

B. Mean absolute error, in basis points				
	0-3%	3-7%	7-10%	10-15%
Baseline	178.4	45.7	42.7	29.6
no dispersion in PDs	177.3	54.7	36.9	26.8
no dispersion in correlations	256.2	52.1	59.1	35.5
no dispersion in PDs and correlations	231.0	56.2	49.1	32.1
One-factor correlation structure	222.1	46.2	48.9	31.1
<i>t</i> -copula: (4,4)	371.8	95.2	65.3	25.5
<i>t</i> -copula: (4, ∞)	222.5	62.3	51.4	21.6
Adjust the level of PDs	171.1	53.1	39.1	28.1

C. Mean absolute percentage error, in per cent				
	0-3%	3-7%	7-10%	10-15%
Baseline	10.2	15.5	37.5	62.8
no dispersion in PDs	10.2	18.9	31.7	55.2
no dispersion in correlations	14.5	16.5	52.9	76.5
no dispersion in PDs and correlations	13.1	18.7	43.1	67.7
One-factor correlation structure	12.6	14.8	43.2	66.1
<i>t</i> -copula: (4,4)	21.0	30.8	57.8	51.4
<i>t</i> -copula: (4, ∞)	12.7	18.9	44.0	41.9
Adjust the level of PDs	9.8	20.1	32.8	58.9

Note: The reported averages are based on daily predicted tranche spreads calculated between November 21, 2003 and March 18, 2005 (369 business days). Column headings specify the particular tranche of the CDS index. Row headings indicate alternative assumptions behind predicted spreads. The “baseline” results incorporate firm-specific PDs and pairwise asset return correlations estimated from the single-name CDS market and assume that asset returns are Gaussian. The pricing of index tranches in the baseline case adopts the Monte Carlo simulation technique (see Appendix A.1). The other results reflect variations on the baseline scenario and are obtained by: (1) removing the dispersion in PDs on each day; (2) removing the dispersion in correlation coefficients on each day; (3) removing the dispersion in both PDs and correlation coefficients on each day; (4) adopting a single-factor correlation structure that best fits the original correlation matrix (see Appendix B) and relying on a Gaussian copula (Appendix A.2); (5) adopting the same single-factor correlation structure as in (4) but relying on “*t*-copula” (the two numbers in parentheses refer to the degrees of freedom of the common and idiosyncratic factors); and (6) adjusting the level of individual PDs so that the cross-sectional average PD equals the PD implied by the single-tranche index spread on each day.

Table 3: **Explaining the time variation in tranche spreads**

	in levels				in first differences			
	0-3%	3-7%	7-10%	10-15%	0-3%	3-7%	7-10%	10-15%
A. Dependent variable: observed spreads, in basis points (bps)								
Predicted spreads (bps)	0.54 (15.6)	0.88 (11.3)	0.87 (8.7)	0.85 (4.2)	0.55 (7.7)	0.78 (9.8)	0.81 (7.6)	0.64 (3.8)
<i>Adjusted R</i> ²	0.78	0.64	0.52	0.19	0.46	0.58	0.45	0.16
B. Dependent variable: observed spreads (bps)								
PD (bps)	15.34 (23.9)	5.39 (10.9)	2.16 (13.8)	0.89 (8.1)	15.74 (11.3)	4.99 (9.0)	2.26 (9.3)	0.86 (6.6)
<i>Adjusted R</i> ²	0.89	0.63	0.73	0.48	0.65	0.53	0.55	0.37
C. Dependent variable: observed spreads (bps)								
PD (bps)	15.40 (23.5)	5.05 (11.3)	2.04 (15.2)	0.79 (8.8)	15.51 (10.9)	4.99 (8.7)	2.25 (9.0)	0.86 (6.4)
Correlation (%)	1.66 (0.5)	-9.26 (4.5)	-3.28 (5.3)	-2.52 (6.0)	4.63 (0.7)	0.05 (0.0)	0.25 (0.2)	-0.03 (0.1)
<i>Adjusted R</i> ²	0.89	0.71	0.81	0.66	0.65	0.53	0.55	0.37
D. Dependent variable: predicted spreads (bps)								
PD (bps)	24.49 (64.8)	6.26 (77.9)	2.01 (36.3)	0.58 (23.7)	24.23 (48.7)	6.22 (44.3)	1.68 (21.2)	0.39 (10.2)
Correlation (%)	-27.94 (16.8)	3.00 (8.5)	4.15 (17.0)	1.77 (16.4)	-45.55 (20.0)	4.85 (7.6)	6.56 (18.0)	3.03 (17.5)
<i>Adjusted R</i> ²	0.99	0.99	0.95	0.91	0.97	0.97	0.93	0.88

Note: All variables are used as weekly averages and run from November 21, 2003 to March 18, 2005. The *PD* and correlation variables in the regressions are time series of cross-sectional averages of risk-neutral PDs and physical asset return correlations estimated from the single-name CDS market. For the results in the left panel all variables are used in levels; for the right panel, all variables are expressed in first differences. Column headings specify the particular tranche of the CDS index. *t*-statistics are reported in parentheses and significant coefficients (at the 95% confidence level) are in bold. All regressions include a constant term, which is omitted from the table.

Table 4: Explaining errors in predicted tranche spreads

	in levels				in first differences			
	0-3%	3-7%	7-10%	10-15%	0-3%	3-7%	7-10%	10-15%
A. dependent variable: prediction error (bps)								
Correlation (%)	-38.30 (5.8)	11.39 (6.3)	7.40 (18.0)	4.41 (12.7)	-41.46 (4.8)	6.04 (2.5)	5.75 (5.5)	2.59 (4.0)
<i>Adjusted R</i> ²	0.32	0.35	0.82	0.70	0.24	0.07	0.30	0.18
B. dependent variable: prediction error (bps)								
PD (bps)	10.31 (8.5)	0.75 (1.6)	-0.30 (1.4)	-0.37 (2.8)	6.28 (3.1)	1.46 (2.8)	-0.26 (0.9)	-0.32 (2.1)
<i>Adjusted R</i> ²	0.50	0.02	0.01	0.09	0.11	0.08	0.00	0.05
C. dependent variable: prediction error (bps)								
Dependent variable (-1)	0.96 (28.3)	0.97 (31.9)	0.97 (27.4)	0.97 (31.2)	0.36 (3.2)	0.20 (1.7)	0.07 (0.6)	-0.03 (0.2)
<i>Adjusted R</i> ²	0.92	0.94	0.92	0.93	0.12	0.03	0.00	0.00
D. dependent variable: prediction error (bps)								
Dependent variable (-1)	0.84 (14.4)	0.89 (25.2)	0.69 (11.7)	0.82 (16.4)	0.10 (1.0)	0.15 (1.3)	0.01 (0.1)	-0.08 (0.8)
PD (bps)	2.07 (3.1)	0.27 (2.2)	-0.01 (0.2)	-0.03 (0.9)	7.63 (4.3)	1.00 (1.9)	-0.61 (2.6)	-0.51 (3.6)
Correlation (%)	-2.41 (0.8)	2.39 (3.5)	2.59 (5.4)	0.94 (3.7)	-48.17 (6.2)	4.68 (2.0)	6.28 (6.0)	3.05 (4.9)
<i>Adjusted R</i> ²	0.93	0.95	0.94	0.94	0.46	0.12	0.34	0.30

Note: The dependent variable is the weekly average of prediction errors, which are defined as predicted tranche spreads (based on risk-neutral PDs and physical correlation estimates implied by the single-name CDS market) minus observed spreads in the CDS index market. Explanatory variables, specified in row headings, include the first lag of the dependent variable and the weekly average of physical correlations and of PDs. For the results in the left panel all variables are used in levels; for the right panel, all variables are expressed in first differences. All variables run from November 11, 2003 to March 18, 2005. Column headings specify the particular tranche of the CDS index. *t*-statistics are reported in parentheses and significant coefficients (at the 95% confidence level) are in bold. All regressions include a constant term, which is omitted from the table.

Table 5: **Explaining the correlation risk premium**

	dependent variable: Δ correlation risk premium (%)			
Lagged dependent variable			0.09 (0.8)	0.05 (0.5)
Δ homogenized physical correlation (%)		-1.04 (6.8)		-1.01 (6.4)
Δ PD (bps)	-0.07 (1.4)			-0.04 (0.8)
<i>Adjusted R</i> ²	0.02	0.40	0.01	0.41

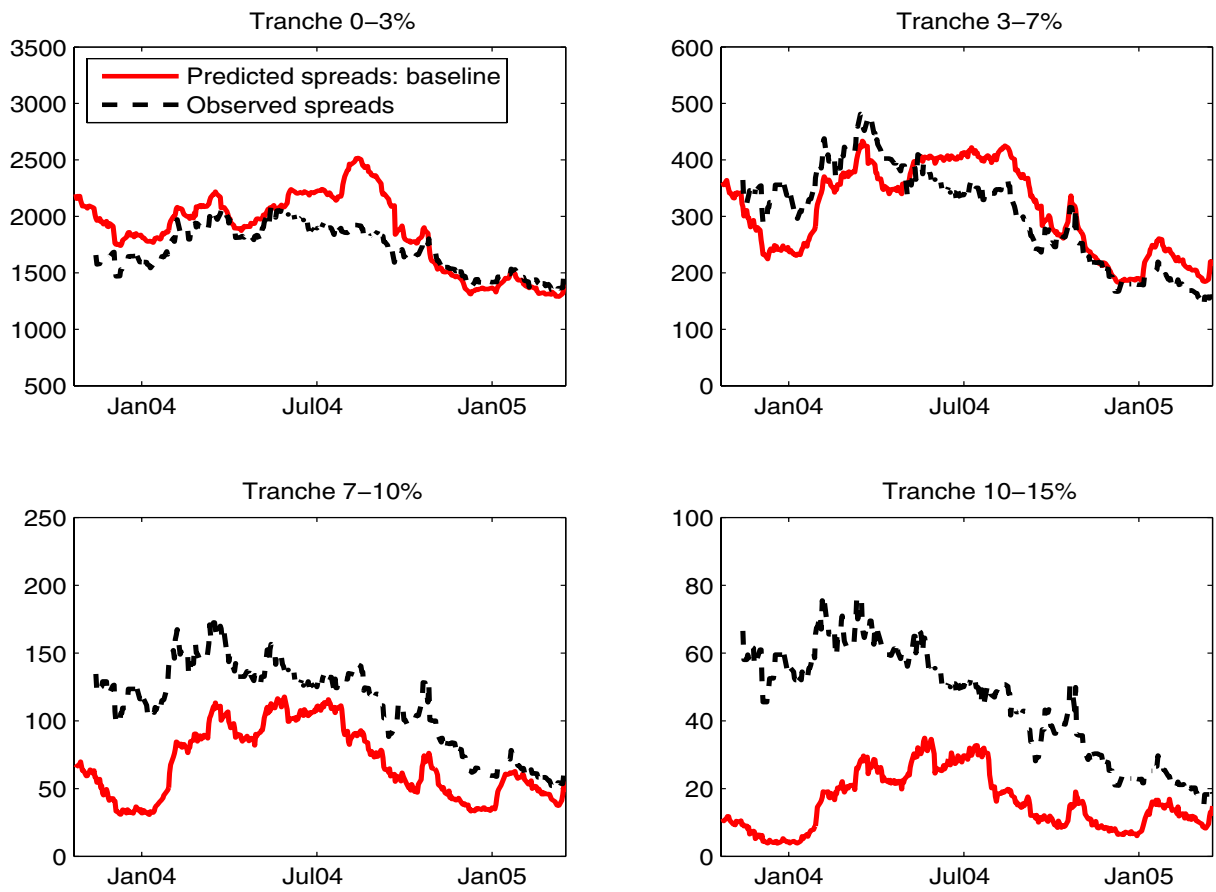
Note: The dependent variable is the change in the weekly average of the correlation risk premium, which equals a risk-neutral correlation minus a homogenized physical correlation. On each day, the risk-neutral and the homogenized physical correlations are constrained to be the same across all pair of firms and across the four CDS tranches. Given this constraint, the risk-neutral and physical correlations are fitted to observed and predicted tranche spreads, respectively, and minimize mean squared percentage errors. Explanatory variables include the first lag of the dependent variable, changes in the weekly average of homogenized physical correlations and changes in the weekly average of PDs. There is no overlap between the sets of observations underlying any two weekly averages. The sample period is from November 21, 2003 to March 18, 2005. *t*-statistics are reported in parentheses and significant coefficients (at the 95% confidence level) are in bold. All regressions include a constant term, which is omitted from the table.

Table 6: **Explaining realized correlations of asset returns**

	dependent variable: Δ homogenized physical correlation (%)			
Lagged dependent variable			0.27 (2.3)	0.24 (2.1)
Δ Risk-neutral correlation (%)	0.28 (2.2)			0.25 (2.0)
<i>Adjusted R</i> ²	0.05	0.06		0.10

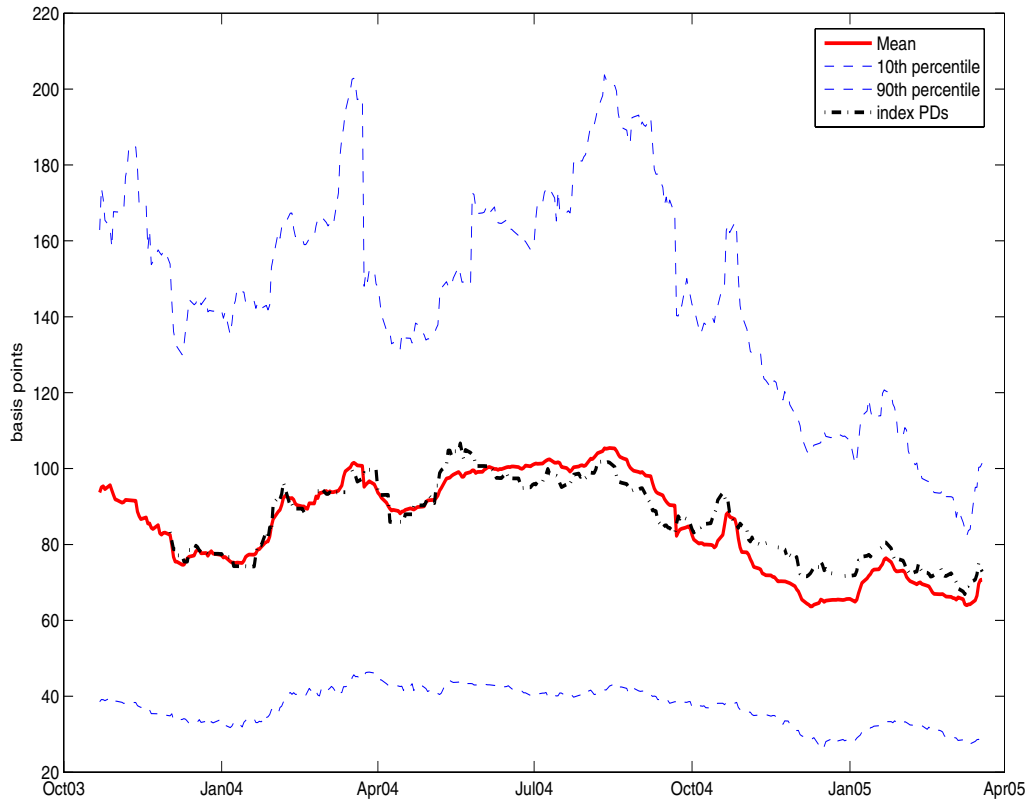
Note: The dependent variable is the first difference of the weekly averages of homogenized physical correlations (see Table 5). Each daily estimate of the homogenized physical correlations is based on single-name CDS spreads observed over the previous 6 months. Explanatory variables include the first lag of the dependent variable and first differences (lagged by six months) of the weekly averages of the risk-neutral correlations implied by observed tranche spreads. There is no overlap between the sets of observations underlying any two weekly averages. The sample period is from November 21, 2003 to March 18, 2005. *t*-statistics are reported in parentheses and significant coefficients (at the 95% confidence level) are in bold. All regressions include a constant term, which is omitted from the table.

Figure 1: Predicted versus observed spreads of CDS index tranches



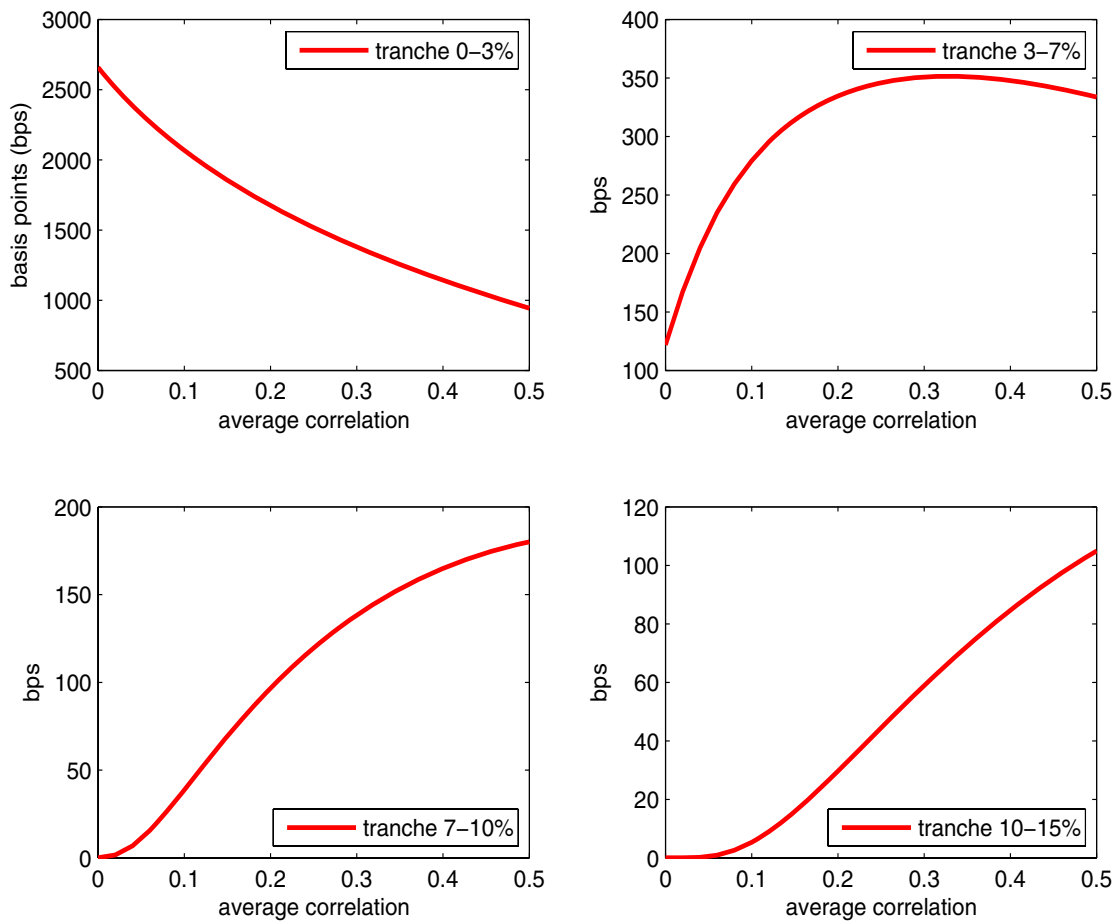
Note: The observed tranche spreads in the CDS index market are provided by JP Morgan Chase. Predicted tranche spreads are based on firm-specific PDs and physical asset return correlations implied by the single-name CDS market as well as on a particular pricing algorithm. This algorithm incorporates Monte Carlo simulations and assumes normally distributed asset returns (see Appendix A.1).

Figure 2: Probabilities of default



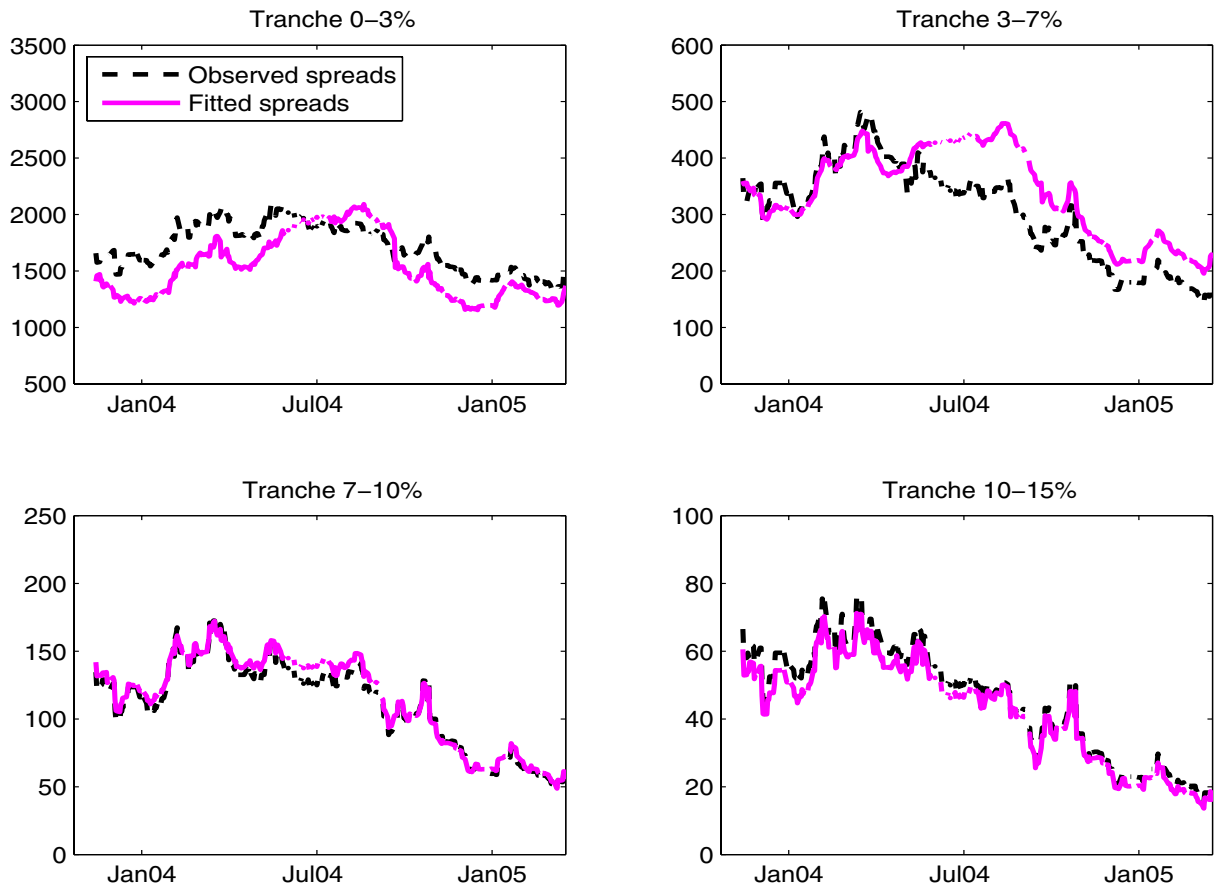
Note: The mean and the two percentiles are based on daily cross sections of PDs, which are estimated from the single-name CDS spreads of the 125 entities in the “on-the-run” release of the CDS index CDX.NA.IG.5Y. The series “index PDs” consists of the average PDs (across the same 125 entities) implied by the single-tranche spreads on the CDS index.

Figure 3: The sensitivity of tranche spreads to average correlations



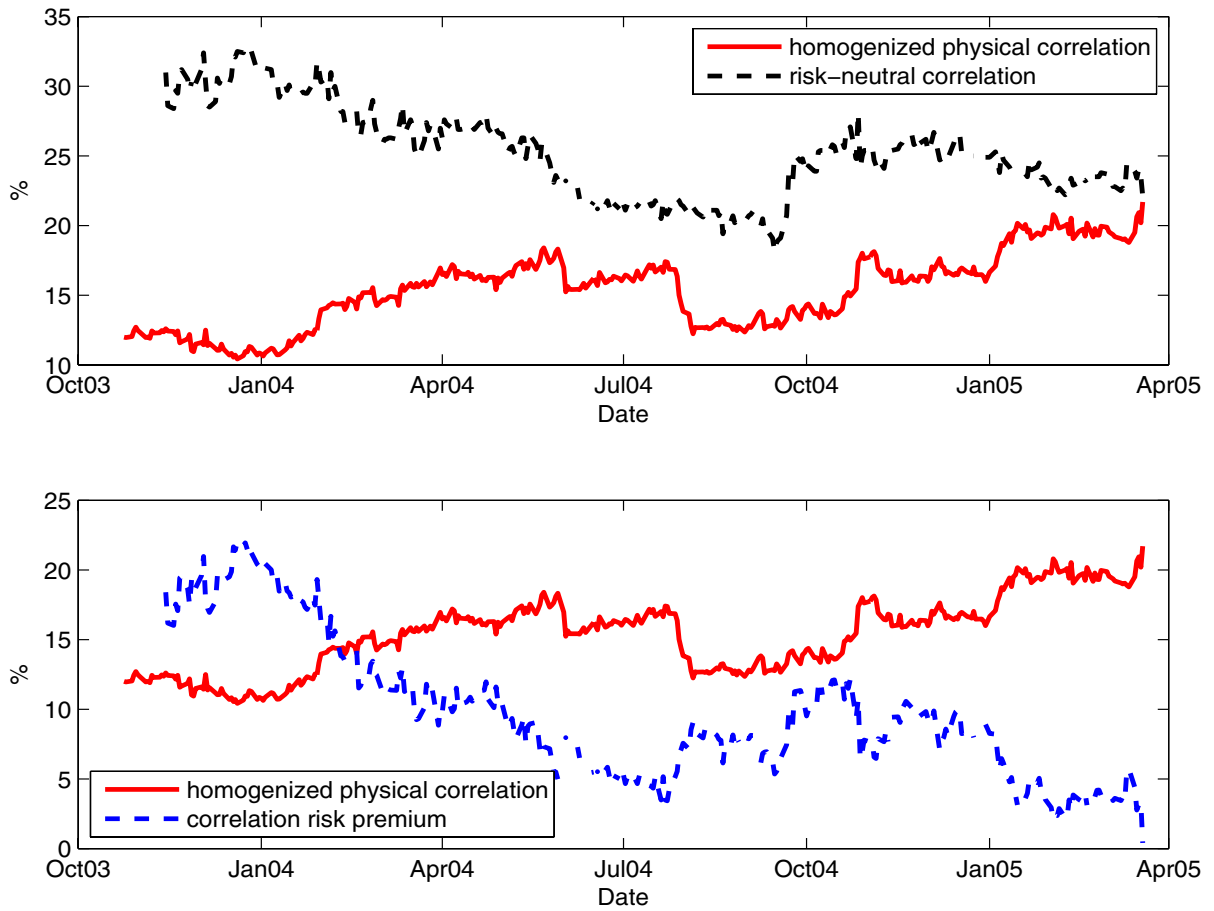
Note: This illustrative example uses the cross section of time averages of PDs, and the average LGD and risk-free rate in our sample. All pairwise correlations are held fixed in the cross section.

Figure 4: Matching observed tranche spreads



Note: Fitted spreads are based on a daily estimate of the risk-neutral correlation, which is the same across all pairs of firms and across tranches. On each day, this risk-neutral correlation minimizes the mean squared percentage error between observed and fitted spreads across the four tranches.

Figure 5: Correlation risk premium



Note: On each day, both the risk-neutral and the homogenized physical correlations are assumed to be constant across pairs of firms and the four (i.e. equity, mezzanine and two senior) tranches. The risk-neutral correlation is fitted to observed tranche spreads, while the physical correlation is fitted to predicted tranche spreads (see Tables 5 and 6 for further detail). The correlation risk premium is defined as the risk-neutral correlation minus the physical correlation.

The following Discussion Papers have been published since 2007:

Series 1: Economic Studies

01	2007	The effect of FDI on job separation	Sascha O. Becker Marc-Andreas Müндler
02	2007	Threshold dynamics of short-term interest rates: empirical evidence and implications for the term structure	Theofanis Archontakis Wolfgang Lemke
03	2007	Price setting in the euro area: some stylised facts from individual producer price data	Dias, Dossche, Gautier Hernando, Sabbatini Stahl, Vermeulen
04	2007	Unemployment and employment protection in a unionized economy with search frictions	Nikolai Stähler
05	2007	End-user order flow and exchange rate dynamics	S. Reitz, M. A. Schmidt M. P. Taylor
06	2007	Money-based interest rate rules: lessons from German data	C. Gerberding F. Seitz, A. Worms
07	2007	Moral hazard and bail-out in fiscal federations: evidence for the German Länder	Kirsten H. Heppke-Falk Guntram B. Wolff
08	2007	An assessment of the trends in international price competitiveness among EMU countries	Christoph Fischer
09	2007	Reconsidering the role of monetary indicators for euro area inflation from a Bayesian perspective using group inclusion probabilities	Michael Scharnagl Christian Schumacher
10	2007	A note on the coefficient of determination in regression models with infinite-variance variables	Jeong-Ryeol Kurz-Kim Mico Loretan

11	2007	Exchange rate dynamics in a target zone - a heterogeneous expectations approach	Christian Bauer Paul De Grauwe, Stefan Reitz
12	2007	Money and housing - evidence for the euro area and the US	Claus Greiber Ralph Setzer
13	2007	An affine macro-finance term structure model for the euro area	Wolfgang Lemke
14	2007	Does anticipation of government spending matter? Evidence from an expectation augmented VAR	Jörn Tenhofen Guntram B. Wolff
15	2007	On-the-job search and the cyclical dynamics of the labor market	Michael Krause Thomas Lubik
16	2007	Heterogeneous expectations, learning and European inflation dynamics	Anke Weber
17	2007	Does intra-firm bargaining matter for business cycle dynamics?	Michael Krause Thomas Lubik
18	2007	Uncertainty about perceived inflation target and monetary policy	Kosuke Aoki Takeshi Kimura
19	2007	The rationality and reliability of expectations reported by British households: micro evidence from the British household panel survey	James Mitchell Martin Weale
20	2007	Money in monetary policy design under uncertainty: the Two-Pillar Phillips Curve versus ECB-style cross-checking	Günter W. Beck Volker Wieland
21	2007	Corporate marginal tax rate, tax loss carryforwards and investment functions – empirical analysis using a large German panel data set	Fred Ramb

22	2007	Volatile multinationals? Evidence from the labor demand of German firms	Claudia M. Buch Alexander Lipponer
23	2007	International investment positions and exchange rate dynamics: a dynamic panel analysis	Michael Binder Christian J. Offermanns
24	2007	Testing for contemporary fiscal policy discretion with real time data	Ulf von Kalckreuth Guntram B. Wolff
25	2007	Quantifying risk and uncertainty in macroeconomic forecasts	Malte Knüppel Karl-Heinz Tödter
26	2007	Taxing deficits to restrain government spending and foster capital accumulation	Nikolai Stähler
27	2007	Spill-over effects of monetary policy – a progress report on interest rate convergence in Europe	Michael Flad
28	2007	The timing and magnitude of exchange rate overshooting	Hoffmann Sondergaard, Westelius
29	2007	The timeless perspective vs. discretion: theory and monetary policy implications for an open economy	Alfred V. Guender
30	2007	International cooperation on innovation: empirical evidence for German and Portuguese firms	Pedro Faria Tobias Schmidt
31	2007	Simple interest rate rules with a role for money	M. Scharnagl C. Gerberding, F. Seitz
32	2007	Does Benford's law hold in economic research and forecasting?	Stefan Günnel Karl-Heinz Tödter
33	2007	The welfare effects of inflation: a cost-benefit perspective	Karl-Heinz Tödter Bernhard Manzke

34	2007	Factor-MIDAS for now- and forecasting with ragged-edge data: a model comparison for German GDP	Massimiliano Marcellino Christian Schumacher
35	2007	Monetary policy and core inflation	Michele Lenza
01	2008	Can capacity constraints explain asymmetries of the business cycle?	Malte Knüppel
02	2008	Communication, decision-making and the optimal degree of transparency of monetary policy committees	Anke Weber
03	2008	The impact of thin-capitalization rules on multinationals' financing and investment decisions	Buettner, Overesch Schreiber, Wamser
04	2008	Comparing the DSGE model with the factor model: an out-of-sample forecasting experiment	Mu-Chun Wang
05	2008	Financial markets and the current account – emerging Europe versus emerging Asia	Sabine Herrmann Adalbert Winkler
06	2008	The German sub-national government bond market: evolution, yields and liquidity	Alexander Schulz Guntram B. Wolff
07	2008	Integration of financial markets and national price levels: the role of exchange rate volatility	Mathias Hoffmann Peter Tillmann
08	2008	Business cycle evidence on firm entry	Vivien Lewis

Series 2: Banking and Financial Studies

01	2007	Granularity adjustment for Basel II	Michael B. Gordy Eva Lütkebohmert
02	2007	Efficient, profitable and safe banking: an oxymoron? Evidence from a panel VAR approach	Michael Koetter Daniel Porath
03	2007	Slippery slopes of stress: ordered failure events in German banking	Thomas Kick Michael Koetter
04	2007	Open-end real estate funds in Germany – genesis and crisis	C. E. Bannier F. Fecht, M. Tyrell
05	2007	Diversification and the banks’ risk-return-characteristics – evidence from loan portfolios of German banks	A. Behr, A. Kamp C. Memmel, A. Pfingsten
06	2007	How do banks adjust their capital ratios? Evidence from Germany	Christoph Memmel Peter Raupach
07	2007	Modelling dynamic portfolio risk using risk drivers of elliptical processes	Rafael Schmidt Christian Schmieder
08	2007	Time-varying contributions by the corporate bond and CDS markets to credit risk price discovery	Niko Dötz
09	2007	Banking consolidation and small business finance – empirical evidence for Germany	K. Marsch, C. Schmieder K. Forster-van Aerssen
10	2007	The quality of banking and regional growth	Hasan, Koetter, Wedow
11	2007	Welfare effects of financial integration	Fecht, Grüner, Hartmann
12	2007	The marketability of bank assets and managerial rents: implications for financial stability	Falko Fecht Wolf Wagner

13	2007	Asset correlations and credit portfolio risk – an empirical analysis	K. Düllmann, M. Scheicher C. Schmieder
14	2007	Relationship lending – empirical evidence for Germany	C. Memmel C. Schmieder, I. Stein
15	2007	Creditor concentration: an empirical investigation	S. Ongena, G. Tümer-Alkan N. von Westernhagen
16	2007	Endogenous credit derivatives and bank behaviour	Thilo Pausch
17	2007	Profitability of Western European banking systems: panel evidence on structural and cyclical determinants	Rainer Beckmann
18	2007	Estimating probabilities of default with support vector machines	W. K. Härdle R. A. Moro, D. Schäfer
01	2008	Analyzing the interest rate risk of banks using time series of accounting-based data: evidence from Germany	O. Entrop, C. Memmel M. Wilkens, A. Zeisler
02	2008	Bank mergers and the dynamics of deposit interest rates	Ben R. Craig Valeriya Dinger
03	2008	Monetary policy and bank distress: an integrated micro-macro approach	F. de Graeve T. Kick, M. Koetter
04	2008	Estimating asset correlations from stock prices or default rates – which method is superior?	K. Düllmann J. Küll, M. Kunisch
05	2008	Rollover risk in commercial paper markets and firms' debt maturity choice	Felix Thierfelder
06	2008	The success of bank mergers revisited – an assessment based on a matching strategy	Andreas Behr Frank Heid

07	2008	Which interest rate scenario is the worst one for a bank? Evidence from a tracking bank approach for German savings and cooperative banks	Christoph Memmel
08	2008	Market conditions, default risk and credit spreads	Dragon Yongjun Tang Hong Yan
09	2008	The pricing of correlated default risk: evidence from the credit derivatives market	Nikola Tarashev Haibin Zhu

Visiting researcher at the Deutsche Bundesbank

The Deutsche Bundesbank in Frankfurt is looking for a visiting researcher. Among others under certain conditions visiting researchers have access to a wide range of data in the Bundesbank. They include micro data on firms and banks not available in the public. Visitors should prepare a research project during their stay at the Bundesbank. Candidates must hold a Ph D and be engaged in the field of either macroeconomics and monetary economics, financial markets or international economics. Proposed research projects should be from these fields. The visiting term will be from 3 to 6 months. Salary is commensurate with experience.

Applicants are requested to send a CV, copies of recent papers, letters of reference and a proposal for a research project to:

Deutsche Bundesbank
Personalabteilung
Wilhelm-Epstein-Str. 14

60431 Frankfurt
GERMANY