Intra-daily variations in volatility and transaction costs in the Credit Default Swap market

Andras Fulop[1]　　Laurence Lescourret[2]

This version: October 2009

[1]ESSEC Business School, Avenue Bernard Hirsh, 95021 Cergy-Pontoise, France. Tel: +(33)(0)134433647. email: fulop@essec.fr. Andras Fulop is research fellow at CREST (Paris).

[2]ESSEC Business School, Avenue Bernard Hirsh, 95021 Cergy-Pontoise, France. Tel: +(33)(0)134433362. email: lescourret@essec.fr. Laurence Lescourret is research fellow at CREST (Paris). We thank Gaëlle Le Fol, Anna Obizhaeva, and participants at conferences (International Risk Forum (IRF) in Paris, the 2008 EFA meeting in Athens, the 4th Annual Central Bank Workshop on the Microstructure of Financial Markets in Hong Kong, the 2008 Credit conference in Venice, the Third Annual Risk Management Conference in Singapore) and seminars (National University of Singapore-Risk Management Institute, Tinbergen Institute Rotterdam, and the National Bank of Hungary) for useful comments. Laurence and Andras gratefully acknowledge Euronext as data sponsor, and Mathieu Alazard from GFI for assistance on data. All remaining errors are our own.

Abstract

This paper describes and analyzes the interdealer broker market for Credit Default Swaps (CDSs). Using data on 18 single-name 5-year CDSs, we find that CDSs exhibit reverse J-shaped patterns for trading and quoting activity in the U.S., and U-shaped patterns in Europe and in Japan. We also find that CDSs trade at most 4.6 times a day, which is comparable to corporate bonds. We develop a methodology based on a state-space model of bid and ask quotes adapted to thinly traded securities. The model deals with (i) price discreteness, (ii) data errors, (iii) heterogeneity of quotes, and (iv) jumps in the efficient CDS spreads. We estimate the model using particle filtering and the Monte Carlo EM algorithm. Our main empirical conclusions are that a) the volatility of the efficient CDS spread is highest at the beginning of the trading day and declines subsquently, b) transaction costs exhibit a reverse J-shaped pattern across all CDSs, c) heterogeneity in dealers' private costs (e.g., urgency to hedge, private needs to trade) is associated with the magnitude of transaction costs.

Keywords : Credit Default Swap, Intra-daily patterns, Stochastic transaction costs, Volatility, Interdealer market.

JEL Classification Numbers: C22, G12

1 Introduction

The market for credit derivatives has experienced a tremendous growth in the last decade becoming one of the leading markets. This market grew to $38 trillion by 2008, from a total notional amount of $600 billion in 1999.[1] Despite the strategic position of the market for credit derivatives, very little is known about the liquidity of credit derivative contracts, and, in particular, about the main regularities observed in the credit default swap (CDS) market, which accounts for around 85% of the credit derivatives market. This paper tries to fill this gap by documenting intra-daily variations in the volatility and transaction costs using a sample of single-name CDSs over a period of 24 months.[2]

An analysis of intraday patterns in trading activity and volatility is useful for understanding the way participants behave in the credit derivatives market, and thus the informational content of CDS prices in different periods of the trading day. Our analysis is based on data provided by GFI, a leading interdealer broker, which provides voice- and electronic brokerage. The dataset contains time-stamped information on intraday trade and quote for corporate CDSs traded on the electronic platform, but does not include information on volume (quoted depth or traded volume). Our sample consists of the 18 most active CDSs, whose underlying reference entity may be North American, European or Japanese.

Across all CDSs, the overnight trading activity is close to insignificant. Trading activity thus mainly occurs during local business hours, with a daily peak in the frequency of quotes and trades early during the morning (8.00). The trading activity then subsides slowly for North American contracts. Trading frequency and quoting frequency exhibit a U-shape

[1]Source: The 2009 ISDA Operations Benchmarking Survey, International Swaps and Derivatives Association, Inc.

[2]A single-name CDS is a derivative contract whose payoff depends on whether there is a default ("credit event") in an underlying obligation in a given time interval. Settlement can take two forms, physical and cash. In the case of physical settlement, the protection buyer can deliver the defaulted security for its par value, while for cash settlement he receives the difference between the par value and the after-default market price. In exchange the protection seller receives an upfront insurance fee and/or a regular insurance payment that lasts to the maturity of the CDS contact or the default of the underlying.

for European and Japanese CDSs. We also observe that the relative quoting frequency is higher than the trading frequency during the morning (filling the order book), and then turns to be lower in the afternoon.

We find that the most traded corporate CDS - Telecom Italia - trades in average 4.6 times a day, and has around 25 bid and ask quotes per day. More generally, U.S. CDSs trade on a par with European contracts (around 3 trades a day), while Japanese CDSs trade less (1.7 trades a day). Trading activity is thus comparable to their counterparts in the corporate bond market. The quoting activity is also low: we observe an average of 15 bid and ask pairs for European CDSs and less for U.S. and Japanese contracts (in average, 5 bid and ask pairs). Limit order books are thus very thin. Accordingly, the cost of immediacy turns out to be significant in this market: the cost of turning around a position is, in average, around 70 cents for an investment with a notional of $100. In comparison, Biais and Declerck (2007) find that the costs of a round trip in the European corporate bond market is 30 cents.

In order to go deeper in the analysis, we develop a methodology based on a state-space model of bid and ask quotes, which is suited to the illiquidity of order books. We use Hasbrouck's model (1999) as a workhorse to build a dynamic model, in which bid and ask quotes are rounding transformation of an implicit efficient CDS premium and stochastic transaction costs, whose mean can be time-varying. The dynamics of the efficient premium allow for deterministic time-varying volatility and jumps. Further, we allow for data errors to deal with outliers. To solve the resulting non-linear filtering task we employ particle filters and we use the Monte Carlo EM algorithm to estimate the model parameters.

In this framework, we document that the (fundamental) volatility of the efficient CDS premium reaches a peak during the preopening period (6.30-8.00) across all CDSs contracts. European and Japanese CDSs exhibit a small increase in volatility during the post-close (15.00-18.00), unlike U.S. CDSs whose fundamental volatility remains low and flattened. We observe a very clear J-shaped pattern for transaction costs across all CDSs, although not all U.S. CDSs have a significant increase in transaction costs at the end of the day. A clear pattern for heterogeneity in dealers' private values emerges to be a small but significant increase during lunch hours and another one during the post-close for European and Japanese CDSs, but not for U.S. CDSs. Overall, the results seem to corroborate the hypothesis of large information asymmetry during the preopening period for all CDS contracts. For European and Japanese CDSs, the increase of transaction costs, along with the one in the dispersion of dealers' private costs tend also to provide some support for the hypothesis of CDS dealers trying to control their inventory position by the end of the trading day. For U.S. CDSs, we fail, however, to find any evidence in favor of U.S. dealers using IDB platforms for inventory motives.

We uncover evidence that mean transaction costs are positively correlated to the fundamental variance and highly positively correlated with heterogeneity in dealers' private costs, in accordance with models of limit order book (e.g., Foucault, 1999 or Roşu, 2009). We also find a high negative correlation between transaction costs and the number of quotes, consistent with larger transaction costs when the book is thinner; the correlation is also highly positive between the number of trades and the number of quotes, confirming that dealers consume liquidity when the book is the least empty.

Does the ongoing crisis affect these intraday variations? We check whether our results are robust to the subprime mortgage crisis which became apparent end of July 2007. We find a drop in the daily number of trades and an increase in fundamental volatility after the credit crisis was revealed. However, intraday patterns, though weaker, seem robust to the onset of the credit crisis.

Our paper makes three distinct contributions. First, to our knowledge, this is the first paper to document intra-daily patterns in the volatility and trading costs of single-name CDSs, whose underlying country of issuance may be European, Japanese or North American. The second contribution is to provide a methodology suited to estimate transaction costs in illiquid markets. In particular, the model deals with jumps in the efficient CDS premium and outliers, which are a prevalent feature of infrequent high frequency data. Third, the paper contributes to the literature on IDB markets. Reiss and Werner (1998), Bjønnes and Rime (2005) or Reiss and Werner (2004) find strong evidence supporting that dealers mainly use the IDB to rebalance their inventory at the end of the day. In contrast, our results show that CDS dealers might trade in the IDB for private information and inventory control motives.

Our paper belongs to the nascent literature dealing with liquidity in the CDS market. Using five different measures of liquidity, Tang and Yan (2007) find that the more liquid a CDS is, the lower its premium gets. Bongaert, de Jong and Driessen (2007) find that expected liquidity generates a liquidity premium for the protection seller. Ashcraft et Santos (2009) define liquidity in the CDS market by the number of daily quotes posted by dealers-contributors. The authors find that firms whose CDSs become more liquid, benefit from a reduction in the cost of their bank loans. Another important role that the CDS market has come to play is in the price discovery process. Blanco, Brennan and Marsh (2005) find a significant lead for CDS premia over credit spreads in the bond market in the price discovery process. Hull, Predescu and White (2005) report evidence on the informational efficiency of the CDS market, by showing that CDS premium changes predict rating changes. Acharya and Johnson (2007) uncover evidence of significant information revelation in the credit default swap (CDS) market, consistent with the occurrence of insider trading. Our paper complements these results by documenting the period of the day during which, most presumably, informed traders play.

This paper is organized as follows. Section 2 presents and analyzes some salient features of the credit default swap market. Section 3 describes the data used in this study and provides preliminary evidence on intradaily patterns in market activity for CDS trading. Section 4 presents the econometric model and discusses the filtering and estimation method. Section 5 reports the model estimates and contains robustness checks. Section 6 concludes the paper.

2 Overview of the CDS market

This section presents the principal trading features of the market for single-name CDSs and also discusses why traders trade CDSs. It aims to give a comprehensive view of the forces that may govern intraday volatility and price dynamics in this market.

2.1 Market organization

Credit derivatives are traded over-the-counter (OTC), which means that there is no physical location and no central organized exchange on which orders are matched. Instead, the CDS market operates 24 hours a day through an electronic network of banks, hedge funds and other institutional investors. In particular, the CDS market is organized as a decentralized dealer market. Trading is conducted in bilateral non-anonymous communication over the telephone: dealers trade directly with clients, and among themselves. In the CDS market, dealers are not required explicit continuous presence (unlike dealers on the NASDAQ). Nor do they face rules limiting the size of the bid-ask spread they choose to post, or limiting changes in their prices (unlike the specialist on the NYSE, or designated market-makers on traditional exchanges). Dealers usually provide quotes only on request, and these quotes are not publicly available. Therefore, there is a low level of pre- and post-trade transparency on the trading process.[3] Assessing the level of implicit transaction costs is thus a sensitive issue, as in any other opaque fragmented OTC market.

Although customer-dealer trading is crucial, trading between dealers represents the most important part with more than 80% of total trading volume across all single-name CDSs.[4] Interdealer trading may occur through direct negotiation between dealers, or indirectly through voice/electronic brokers. This indirect trading through interdealer brokers (IDBs), which involves voice-based systems and/or electronic platforms (e.g., GFI, Creditex, or Tullett Prebon), plays an increasing role. According to the International Swaps and Derivatives Association (ISDA), interdealer brokerage represents 34% of trades in 2004.

Interdealer brokers collect, post dealer limit orders in a computerized network, and

[3]The credit crisis has pushed regulators to propose changes to legislation governing OTC derivatives trading in order to improve transparency. In particular, standardized CDSs will be traded with a centralized counterparty (CCP). Information on transaction prices, traded volumes and exposures will be available, at least to regulators. In 2009, two operators (ICE/TCC and CME - Citadel) started processing and clearing of CDS index transactions and plan to do the same for liquid single-name CDSs in the future. More transparency in the OTC derivative market could thus be enforced in the near future.

[4]Source: The Depository Trust & Clearing Corporation (DTCC), Trade Information Warehouse Data (Section I), October 2009.

execute trades between dealers by matching buyers and sellers. Dealers are also allowed to directly submit limit or market orders into the computerized platform, which then matches orders. Only dealers have access to IDB systems, and they are not required to submit quotes to IDBs. Dealers using IDB platforms place and hit quotes anonymously. Even the size of their limit orders may not be disclosed in some IDBs.[5]

Compared to the direct interdealer market, IDBs are costly since dealers must pay the broker' s fee. IDBs provide, however, dealers that are clients of the broker with platforms which broadcast anonymous limit prices on screens. In particular, dealers have access to information on prices of current quotes, and direction and size of past transactions (information on the order flow). In that respect, IDBs are more transparent, more centralized than the current direct interdealer market, contributing thus to decrease search costs for a matching counterparty.[6] As computerized platforms, they also enhance the speed of execution. Estimating transaction costs in these interdealer electronic platforms may thus provide a lower bound on expected transaction costs borne by final investors.

2.2 Who trades CDSs?

Academics and practitioners point out several reasons for trading CDSs. First, participants transfer credit risks related to clients, counterparty risk, or reduce concentration risk to an industry, although Duffie (2007) or Minton, Stulz and Williamson (2008) report that loan hedging and direct credit risk mitigation may not be the most important reason to enter CDS contracts. Second, CDSs are used for regulatory capital management, where an institution enters a CDS contract to decrease the amount of regulatory capital it needs to hold. Finally, participants also trade for speculation and information motives. Acharya and Johnson (2007) find evidence of insider trading on the CDS market. Fitch (2009) reports that "[...] a very significant portion of the CDS activity of the largest banks

[5] The anonymity of IDB systems is a key success factor of these platforms. Dealers choose to trade anonymously because their identity is an additional piece of information that could be exploited by competitors.

[6] See Günduz, Lüdecke and Uhrig-Homburg (2007) for a comparison between the brokered and the direct interdealer market for single-name CDSs.

was influenced by trading and structuring activities and the intermediation role played by them [...].” In summary, CDS may be traded either for liquidity reasons arising from heterogeneity in traders’ private characteristics (target risk exposure, investment horizon) or to profit from private information about the underlying firm (e.g. free cash flows, credit events). Compared to some OTC markets with macroeconomic drivers (government bonds or currencies), private information is likely to be a more critical motive to trade corporate credit instruments. Since the corporate bond market is illiquid and difficult to leverage and short, the CDS market may be the preferred venue for information-based trading.[7] Accordingly, Blanco, Brennan and Marsh (2005) find that the CDS market leads the corporate bond market in terms of price discovery, consistent with informed traders trading first single-name CDSs.

What are the likely determinants of the bid-ask spreads posted by the CDS market-makers? First, in light of the above discussion, market-makers need to be rewarded for the risk of trading against better informed traders. Second, dealers also face inventory risks when their long positions do not offset perfectly with short positions. The risk to hold suboptimal portfolios is a cost for which dealers have to be compensated. Transaction costs for CDSs should thus reflect cost components identified by the microstructure literature: inventory-holding costs, asymmetric information costs and order-processing costs.[8] In addition, since the CDS market is rather concentrated (the top 5 dealers provide around 50% of the total notional amount bought and sold, according to ECB, 2009), transaction costs could also contained an extra component which would reflect the weakness of the competition among CDS dealers (a rent component).[9]

[7]IMF (2007) reports that haircuts on high-yield bonds went above 10% on the summer 2007.

[8]Order-processing costs might be interpreted as the market-maker’s compensation for presence costs such as labor, communication, transfer taxes, etc..

[9]According to ECB (2009), the top 5 dealers are JP Morgan, Goldman Sachs group, Morgan Stanley, Deutsche Bank, and Barclays group.

2.3 The role of the IDB market

The interdealer market is of paramount importance to maintain a minimum level of liquidity in the market for CDSs. As mentioned above, it is very unlikely that the need of credit protection by a dealer's client is perfectly balanced by trades with other clients (in terms of notional amount, or maturity of the swap traded). Dealers thus rely on their ability to trade with other dealers to smooth over any imbalances in their books. That is the reason why inter-dealer trading is so considerable in this market.[10] The IDB market offers in particular a unique laboratory to estimate the quality of the CDS market since this centralized and transparent structure may be considered as free of search costs, and gives firm (and not indicative) quoted prices. In addition, rents extracted by market-makers should be a less sensitive concern in this anonymous inter-dealer system.

Why do dealers trade in particular in the IDB market? First, dealers might use IDBs to unwind inventory position after trading with customers. They might also use IDBs to build position in anticipation of some client trades. Inventory control and risk-sharing could thus be a motive behind interdealer trading, as suggested by Ho and Stoll (1983). Secondly, since IDB platforms guarantee anonymity of the participants and since private information is a motive to trade CDSs, the level of information-based trading could be high in these facilities. Huang (2002) find that anonymous platforms (ECNs) are the favored trading venue of informed trades. However, the few empirical studies that exist on IDBs seem to find more evidence for the inventory motive.[11]

In the London Stock Exchange, Reiss and Werner (1998) show that interdealer markets are mainly used by equity dealers to share inventory risks. Bjønnes and Rime (2005) also find evidence that, in the FX market, IDBs are used by dealers to manage inventory

[10]Because the number of dealers is rather small (no more than 20 dealers worldwide) and because dealers' positions in the inter-dealer market are so large, counterparty risk (i.e. dealers' exposure to the risk that another dealer default) is crucial. In that perspective, implementing CCP organizations is key to enforce the role of the CDS market in the economy.

[11]The most well-known IDBs of dealer markets are GovPX for the US Treasury bond market, CanPX for the Canadian bond market (see, e.g., Boni and Leach, 2004 ; Huang et al., 2002 ; D'Souza et al., 2003), EBS and Reuters 2000-2 for the FX market (see, e.g., Bjønnes and Rime, 2005).

positions. Second, Reiss and Werner (2004) find that IDBs in the LSE attract uninformed order-flow, whereas the direct non-anonymous inter-dealer market attracts informed inter-dealer trades, contrarily to standard information-based microstructure theories. There are two inter-related reasons. First, Reiss and Werner (2004) observe that when adverse selection is perceived to be high, limit orders are cancelled in IDB platforms, and liquidity dries up. Posting a limit order is thus riskier since it may not be executed. Secondly, information leaks out when posting a limit order in the computerized trading platform.

In light of these studies, one expects to observe intraday patterns in the IDB market for CDSs that are consistent with end-of-day inventory control by individual traders and informed trading. First, trading activity and the number of trades should pick at the end of the business day, to reflect the desire to close open positions before the overnight hours. Second, if adverse selection is important in the CDS IDB market, one should also observe an especially high level of transaction costs and volatility at the beginning of the day.

3 Data

3.1 A preliminary look at the data

3.1.1 *The data*

The data are from GFI, a major CDS inter-dealer broker. GFI is ranked No. 1 in credit default swap trading in the last 8 years (Risk Magazine, 2008). According to GFI, it represents 60% of the inter-dealer brokerage activity. GFI data are thus representative of the interdealer market activity.

GFI is a hybrid voice-electronic execution platform for CDSs dedicated only to dealers. Retail and institutional investors (the "buy-side") are not eligible to use the electronic trading platform or voice-brokers directly. The GFI electronic platform — CreditMatch — is a continuous open limit order book. It has a minimum trade size of $1 million. Only GFI brokers can see the identity of limit order traders, and the depth of the market. They also have information on transactions (size, price, direction, and identity) for the current and last trading days. In contrast, clients — bank dealers — can only observe anonymous limit orders, and the last trades of the day (price, volume, and direction, but not the identity except when counterpart of the trade). As in a standard limit order book, dealers may be providers of liquidity (limit order traders). They may also consume liquidity when they hit a quote or place an aggressive quote into the order book.

3.1.2 *The sample*

The dataset we use in this paper runs from March 31, 2006, through March 31, 2008. Our GFI dataset contains trade and quote information on 1,944 worldwide CDS contracts, whose seniority may be 1 to 10 years. The number of CDS contracts whose maturity is 5-year represents 87.6% of the dataset. The regional breakdown is as follows: North American contracts constitute 36.68% of our data, European contracts represent 30.76%, and Asian entities 27.83%, while other worldwide CDS are minority (only 3.60% CDS contracts are Australian, 0.46% are African, and 0.67% are South American). The dataset is also industry clustered. Table 1 reports that the industries that are most active by number of contracts and by percentage of quotes and trades across all CDS. Among the most active for our sample period are industries such as Banks, Financial Services, Electricity, Real Estate or Fixed Line Telecom.

Our sample consists of 18 5-year senior CDS contracts, which is obtained as follows. First, we limit our study to a set of the most active corporate single-name contracts in terms of number of bid and of ask quotes. Second, the regional diversity of the raw dataset is represented. The final sample is made up of 6 U.S. underlying reference entities (American Axle & Manufacturing, Ford Motor Corp, General Motor Corp, Goodyear Tire, Lear and Visteon), 6 European entities (Casino Guich, Deutsche Telecom, ITV Plc, Portugal Telecom Intl, Telecom Italia and Valeo) and 6 Japanese entities (ACOM, Aiful, Orix Corp, Promise Co, Softbank Corp. and Takefuji). There are 8 investment grade firms with rating above and equal to “BBB” and 10 high yield firms with rating lower than “BBB”.

The data sample consists of bid quotes, ask quotes, and trade prices expressed in basis point. Our data is time stamped down to a second. Our local time benchmark is New York GMT for U.S. CDSs, London GMT for European CDSs, and Tokyo GMT for Japanese CDSs. Bid and ask quotes posted on the platform are firm and anonymous. We have no identifiers of the dealers who post limit orders, and no orders identifiers (sequence number), implying that we cannot track deletions or amendments of any recorded orders. Besides, we have no information on the size of routed orders. As a consequence, we cannot observe the displayed "depth" of the book. Moreover, unlike standard data on electronic trading venues (NASTRAQ, TAQ, Euronext data, etc.), we do not have information on the best limits (the maximum bid and the minimum ask that prevail at any moment). We thus cannot trace the inside spread (difference between the best ask and the best bid), or reconstruct the limit order book.

Turning to data on trades, only the time-stamped transaction price is mentioned. There is therefore no available information about traded CDS volume nor about notional.[12] Besides, transaction prices are not linked to underlying orders. The reasons are that (i) there is no order identifiers as previously mentioned, and (ii) trade prices may result from an electronic matching in CreditMatch, or from an OTC transaction negotiated over the phone between the voice broker and the client. Consequently, we cannot trace executed trades back to initial orders. Moreover, signing trades (determining whether it is a buy or a sell trade) is not easy. Less than 15% of our trades could indeed possibly be signed with the Lee and Ready algorithm.[13] In addition, in this subsample, there is still the uncertainty on the source of the matching (voice or electronic). We let this issue for further research, and focus on an estimation of the cost of the committed liquidity based on the displayed bid and ask quotes.

Nevertheless, before dropping trade data from the main analysis, we use the number of trades per unit of time (the trade frequency) to analyze some stylized facts of market

[12]From interviews with traders, it seems that the minimum trade size is €$5 - 10$ million in Europe. In the U.S., CDS are only traded in relatively large slot ("granularity" of CDS contracts). For instance, a trade on a \$ $20 - 30$ million for the most liquid North American CDSs is quite standard. Note also that the most common practice is to trade contracts whose maturity is 5-year.

[13]We take a 30 minutes window, as suggested by Chan et al. (2002) for equity options.

activity. We also examine another proxy of trading activity: the number of price limits per unit of time.[14]

3.1.3 *Introductory stylized facts*

Figure 1 represents a typical week of trading for the 5-year CDS contract on Telecom Italia. Fig. 1 plots bid quotes, ask quotes and trade prices expressed in basis point during 4 trading days (May 2, 3, 4 and 5 2006). Noticeably, for any date, quotes and trades are irregularly spaced during the day and may be heterogeneous even when they are recorded at the same time. The data points are sparsely located. Finally, there is almost no overnight trading activity for the CDS on Telecom Italia during these 4 trading days.

Figure 2a represents the histogram for the trading activity (the average frequency of price limits and of trades per unit of time) of the pooled U.S. CDSs during our sample period, while Figure 2b and 2c respectively depict the trading activity of the pooled European and Japanese CDSs.[15] Several interesting patterns are revealed by these histograms. First, although trading could take place round-the-clock, the trading and quoting activity of all CDSs mainly occurs during business hours (8.00 to 17.30 local G.M.T.) of the country of issuance. The reason may be that news concerning a CDS underlying firm is presumably announced during the local business hours of its headquarters. Second, U.S. CDSs exhibit slightly different trading patterns from European or Japanese CDSs.

The trading and quoting activity of U.S. CDSs is concentrated during morning hours. The daily peak in the number of quotes occurs during the New-York preopening period (6.30 to 8.00), then it falls steadily, ending by 18.00. The daily number of trades shoots up around 10.00 and subsides along the trading day. Unreported results show that there is some weak but significant trading and quoting activity when London opens (3:00 in New York GMT hours) for Ford Motors Corp and General Motor Corp, probably because of

[14] We will use interchangeably price limits or quotes.

[15] The pooled histograms are representative of the market activity on any specific contract. Figures representing the frequency of quotes and of trades for each contract in the sample are, however, available from the authors upon request.

the listing of underlying firms in London. The absence of trading activity at the end of the business day seems to weaken the hypothesis of U.S. dealers using the IDB platform to close or hedge open inventory position before the overnight period.

For European CDSs, the daily number of quotes and trades exhibit a clear U-shape pattern. The number of limit orders and trades placed in the book rises rapidly at the beginning of the trading day (6.30 in London GMT hours) and reaches a peak at 8.00. There is high quoting and trading activity in the afternoon hours (in particular, when New York opens at 15.00 in London GMT hours), while we observe more moderate activity during lunch hours. Japanese CDSs trade mainly during Tokyo business hours, and also exhibit a U-shape pattern for the frequency of price limits and of trades.[16] There is a spike in the number of limit orders placed around 10.00 and a pronounced decrease in quotes and trades between 11.00 and 13.00, which may also be explained by the slowdown of the trading activity during lunch time. In particular, the Tokyo Stock Exchange closes from 11.00 to 12.30. Unlike U.S. or European CDSs, we notice that the trading activity reaches its daily spike around 15.00 while still being important during the morning.

In summary, the trading and quoting U-shape patterns of European and Japanese CDSs resemble those documented in many other OTC markets: see, for instance, the bond market (Fleming, 1997), or the Nasdaq stock market (Chan, Christie and Schultz, 1995, or Barclay and Hendershott, 2003).[17] U.S. CDSs have a similar trading and quoting activity except for the end of the day. To study the intraday patterns of CDSs, we divide the trading day into four periods: the preopen (6.30-8.00), the daytime trading (8.00-15.00), the post-close (15.00-18.00) and the overnight (18.00-6.30) periods.

[16] Trade data for Japanese CDSs turn out to be mis-stamped due to vocal trading prior to 2007. We thus show the histogram for the period from January 1, 2007 to March 31, 2008.

[17] The Nasdaq stock market has been approved to operate as a national securities exchange in January 2007, and was formerly an OTC market.

3.2 Descriptive statistics

Table 2 reports descriptive statistics for the 18 5-year CDS contracts constituting the sample for the period 2006-2008. The most heavily traded CDS contract – Telecom Italia – trades on average 4.6 times a day. U.S. CDSs trade on a par with European CDS contracts (around 3 trades a day), while Japanese contracts trade less (1.7 trades a day). Therefore, the daily number of trades for CDSs is low compared to their counterparts in the equity market, but similar to their counterparts in the corporate bond market. Edwards, Nimalendran and Piwowar (2006) find that the sample of the more active US corporate bonds trades in average 3.7 times per day, whereas Edwards, Harris and Piwowar (2007) find a lower figure (1.9) for their sample which also includes less active bonds. Biais and Declerck (2007) find that Euro-denominated corporate bonds trade on average 4 times per day.[18]

Table 2 shows that even for these reputedly liquid CDS contracts, the number of quotes (bid or ask) is quite low, compared to the equity market. For instance, the most active contract, Telecom Italia, has around 11,700 bid/ask pairs over the period 2006-2008, or, equivalently 25 bid and ask quotes per day. More generally, in our sample, European CDS contracts have more than 5,500 pairs (around 17 bid/ask pairs per day), while U.S. or Japanese entities have a smaller number of quotes: around 2,000 pairs (or less) corresponding to an average of 5 bid/ask pairs per day. According to GFI brokers, U.S. names have a smaller number of quotes compared to European names, partly because U.S. dealers traditionally utilize more voice-brokers than electronic platforms (CreditMatch) to get quotes or to trade. In contrast, the smaller quoting activity of Japanese CDSs is related to their smaller trading activity.

Due to the scarcity of information, it is therefore important to use as much information in the data as possible, in particular, to use the individual bid or ask observations to run the estimation of the econometric model presented below.[19]

[18]Even if the trading frequency is similar between the bond market and the CDS market, the overall trading volume should be higher in the CDS market as the size of individual trade is larger.

[19]In contrast, Huang et al. (2002) and others filter out one-sided quotes when using inter-dealer trading

Table 2 also reports statistics on the cost faced by a trader who wishes to simultaneously sell and buy contracts. This cost reflects the cost of immediacy, and is usually termed the cost of a round-trip. To measure the cost of a round-trip in the CDS market, one cannot simply use the quoted bid-ask spread as is the case in the equity or the bond market. The reason is that entering a long position and unwinding it leaves the investor with a negative cash-flow stream with a frequency and a duration equal to the frequency and maturity of the CDS contract, where the amount of the individual cash-flows is equal to the bid-ask spread. Consequently, to compute the cost of turning around a position, one needs to take the present value of this cash-flow stream. Further, this cash flow stream stops if the underlying defaults, one thus needs to compute a risky annuity to get the PV factor.

Assuming a constant default intensity, this can be obtained in the same way as a risk-free annuity, but using a default-risk adjusted rate to discount the future cash-flows. In particular, if r is the continuously compounded risk-free rate, s is the CDS premium, and R is the assumed recovery rate, then, the risk-adjusted discount rate is approximately $r + \frac{s}{1-R}$. For instance, for Telecom Italia, setting $r = 0.05$, $s = 0.0069$ (the average CDS quote midpoint as a proxy of the CDS premium), $R = 0.4$ as it is usual for CDS pricing, and working with quarterly payments on a 5-year CDS yields a risky annuity factor of 4.3. This in turn implies that the cost of a round trip for an investment in Telecom Italia with a notional of 100 euros is 16.4 cents. Table 2 reports the results for each of the 18 CDSs.

The results show that, across all CDSs, the average cost of a round-trip is around 58.4 basis points (bp). Three important features of this round trip cost are to be noted. First, European CDSs appear to be less expensive than U.S. or Japanese CDSs. The cost of rounding transactions for European contracts is, in average, 27 bp, while the average cost of a round-trip amounts to 48 bp for U.S. contracts and to 76 bp for Japanese CDSs in the sample. Analyzing a sample of Euro-denominated corporate bonds, Biais and Declerck (2007) find an average quoted spread around 30 cents. Edwards, Harris and Piwowar (2007) find that the effective cost is 20 cents for large trades on US corporate bonds. Second, it is worth noticing that the cost of round-trip trades is inversely related to the

system data.

number of trades as illustrated by Figure 3. This pattern is consistent with previous studies documenting a negative correlation between transaction costs and liquidity proxies such as number of trades (see, among others, Chordia, Roll and Subrahmanyam, 2001). Third, these numbers suggest that (i) bid-ask spreads and transaction costs in the CDS market are not lower than those in the corporate bond market ; (ii) trading frictions are likely to be important in this market and one needs to be careful in interpreting CDS premia as a pure measure of credit risk.

4 A dynamic model of quotes placed in illiquid order books

This section presents a methodology adapted to the rather illiquid and weak activity of the CDS interdealer market. In particular, we link the limit prices to an unobserved efficient price and cast the resulting estimation task as a missing-data problem.[20] Filtering and estimation in the resulting non-linear state-space system is solved using simulation-based methods. This modelling framework allows us to estimate intra-daily patterns of the implicit efficient CDS premium and intra-daily patterns of the stochastic variation of transaction costs.

4.1 Econometric Model

We consider a model conditional on the arrival of bid and ask prices of the CDS premium resulting from posting limit orders in the book. Let $\tau_i, i = 0, ..., T$ denote the joint arrivals of these data points. Let $\mathcal{D}_i$ denote the information set up to τ_i. We either have a bid or an ask observation at τ_i denoted by B_{τ_i}, or A_{τ_i} respectively, or we may have both. Similarly to Hasbrouck (1999) we assume a latent variable model for our data. In particular, we link all observed ask and bid prices to an unobservable efficient log CDS price, M_{τ_i}.

[20]The implicit efficient price is the standard estimate of the full-information price based on all available public information (proxied by the observed order flow).

The unobserved efficient CDS price model. The efficient price innovation allows fat-tails, a prevalent feature of our data. One potential reason for this sudden large moves are, for instance, downgrades. Therefore, the logarithm of the unobserved efficient CDS premium evolves as:

$$m_{\tau_i} = m_{\tau_{i-1}} + \exp(\frac{\sigma^2_{\tau_{i-1}}}{2})\sqrt{\Delta\tau_i}u_i \tag{1}$$

where $m_{\tau_i} = log(M_{\tau_i})$, $\Delta\tau_i = \tau_i - \tau_{i-1}$, and u_i is a mixture of normal distributions $MM(1; \sigma^2_J; \lambda)$, where λ is the probability of the component with variance $1 + \sigma^2_J$. The log fundamental variance, $\sigma^2_{\tau_{i-1}}$, has a deterministic diurnal pattern, modelled using a flexible cubic spline specification. Note that this specification is on purpose more flexible compared to the one used by Hasbrouck (1999) for equities. The reason is that we cannot lean on previous literature on CDSs to guide the choice of the expected shape, as Hasbrouck (1999) does.

The transaction cost model. To develop an economic intuition behind the model, suppose that dealers, who submit limit orders in the computerized platform, are exposed to several costs and risks. These costs are assumed to reflect (i) the possibility of non-execution (Foucault, Kadan and Kandel, 2005; or Roşu, 2009) and (ii) the valuation risk due to public and private information (see Glosten, 1994; Foucault, 1999 or Roşu, 2009).

The mean compensations for liquidity provision on the buy (bid) or on the sell (ask) side is denoted κ_{τ_i}.

$$ln(\kappa_{\tau_i}) = \mu_{\tau_i} + m_{\tau_i} \tag{2}$$

We model the deterministic element of the transaction costs, μ, using a flexible cubic spline specification. We incorporate the level of the efficient premium in our model to get rid of the artificial level effect, in accordance with the literature related to CDS spreads (e.g., Acharya and Johnson, 2007).

We also assume the existence of a dealer-specific independant private value for the asset. In particular, liquidity suppliers may vary by their private needs to hedge and by their level of impatience, thus they may value the speed of execution of their limit orders differently (see, e.g., Foucault, Kadan and Kandel, 2005). To account for this heterogeneity in private costs, we suppose that individual quotes can differ from the average with normal errors (in a proportional sense: the volatility of the private value is proportional to the level of the spread):

$$c_i^A \sim N(0, \exp(\sigma^2_{c,\tau_i}) \times \exp(2m_{\tau_i})) \quad (3)$$

$$c_i^B \sim N(0, \exp(\sigma^2_{c,\tau_i}) \times \exp(2m_{\tau_i})) \quad (4)$$

A flexible cubic spline is used to model the intra-daily seasonal pattern of the heterogeneity in dealers' private costs, σ^2_{c,τ_i}.

In summary, the limit order trader who determines his ask quote is assumed to be driven by a mean cost of liquidity supply, κ, and an idiosyncratic cost reflecting his private value, c.

We also have to deal with outliers. Thus, we allow for data errors with some probability p. Later, in the implementation of the model, we set $p = 1\%$, which can be interpreted as if 1% of our data is disposed of. This yields terms of

$$\varepsilon_i^A \sim q_i^A N(0, \sigma^2_\varepsilon) \quad (5)$$

$$\varepsilon_i^B \sim q_i^B N(0, \sigma^2_\varepsilon) \quad (6)$$

where q_i^A and q_i^B are Bernoulli's with probability p.

Finally, the bid and ask prices set in the limit order book are given by:

$$a_{\tau_i} = M_{\tau_i} + \kappa_{\tau_i} + c_i^A + \varepsilon_i^A \quad (7)$$

$$b_{\tau_i} = M_{\tau_i} - \kappa_{\tau_i} + c_i^B + \varepsilon_i^B \quad (8)$$

Then, if the tick size is K, the discrete bid and ask prices are given by

$$A_{\tau_i} = Ceiling(a_{\tau_i}, K) \quad (9)$$

$$B_{\tau_i} = Floor(b_{\tau_i}, K) \quad (10)$$

where the floor and ceiling functions round down and up respectively to the next multiple of K.

Note that there is no nonnegativity restrictions on the costs for liquidity provision c_i^A and c_i^B. Following the findings of Goettler, Parlour and Rajan (2007), limit buy orders (resp. limit sell orders) may indeed be posted above (resp. below) the efficient price due to large private values of dealers for the asset.

4.2 Filtering and ML estimation

The previous econometric model allows an insight into the dynamics of CDS quoted prices and transaction costs in two ways. First, given the fixed model parameters and the observed data up to τ_i, $\mathcal{D}_i$, we are able to infer the continuous bid quote (b_{τ_i}) and ask quotes (a_{τ_i}), the efficient log CDS premia, m_{τ_i}, the data errors, (q_i^A, ε_i^A) and (q_i^B, ε_i^B), and the jump indicator, N_i. This amounts to decomposing the innovations in the observed noisy discrete bid and ask quotes into fundamental innovations, stochastic transaction costs and observation noise (data errors and price discreteness). Second, by estimating the fixed parameters of the model we are able to estimate the level of and intra-daily patterns of the log fundamental variance, transaction costs, and heterogeneity in dealers' private costs.

4.2.1 *Filtering*

The first task is to filter the unobserved dynamic states given the observed data. Our econometric model can be cast in state-space form, where the transition equation is defined by the movement of the efficient premia in (1) while the observation equations are in (2)-(10). Unfortunately, the Kalman Filter cannot be used since our system is both non-gaussian and non-linear. The transition equation (1) is non-gaussian as we allow for fat-tailed innovations, while the data errors in (5) and (6) make the measurement equations (9) and (10) non-gaussian. Nonlinearity is present in the measurement equations both because the log efficient premia enters equations (7) and (8) after an exponential transformation ($M_{\tau_i} = e^{m_{\tau_i}}$) and as the discretization in (9) and (10) is a highly nonlinear operation.

While Kalman Filtering cannot be applied in our model, we can still derive a theoretical recursion to sequentially update the state of the system using Bayes' rule. For simplicity, assume, first, that at τ_i we have an ask observation.[21] Second, we need to link the

[21]The case of a bid or having both a bid and an ask observation are similar.

filtering distribution of the hidden states at τ_i, $f(a_{\tau_i}, m_{\tau_i}, N_i, q_{i,A}, \varepsilon_i^A, N_i \mid \mathcal{D}_i)$ to the previous filtering distribution, subsumed in $f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1})$. This is achieved if we obtain the joint filtering distribution, $f(a_{\tau_i}, m_{\tau_i}, N_i, q_{i,A}, \varepsilon_i^A, N_i, m_{\tau_{i-1}} \mid \mathcal{D}_i)$. The target distribution can then be obtained by marginalizing out $m_{\tau_{i-1}}$. Using Bayes' rule and the conditional independence of the system, we obtain the following relationship

$$\begin{aligned}
& f(a_{\tau_i}, m_{\tau_i}, N_i, q_{i,A}, \varepsilon_i^A, N_i, m_{\tau_{i-1}} \mid \mathcal{D}_i) \\
\propto\ & 1_{\{a_{\tau_i} \in [A_{\tau_i} - K, A_{\tau_i}]\}} f(a_{\tau_i}, m_{\tau_i}, N_i, q_{i,A}, \varepsilon_i^A, N_i \mid m_{\tau_{i-1}}) f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1}) \\
=\ & 1_{\{a_{\tau_i} \in [A_{\tau_i} - K, A_{\tau_i}]\}} f(a_{\tau_i} \mid m_{\tau_i}, q_i^A, \varepsilon_i^A) f(m_{\tau_i} \mid m_{\tau_{i-1}}, N_i) \\
\times\ & p(N_i) p(q_{i,A}) f(\varepsilon_i^A) f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1}) \qquad (11)
\end{aligned}$$

While this relationship gives a theoretical link between the subsequent filtering distributions, we still need a practical way to implement the recursion. The problem stems, first, from the fact that $f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1})$ does not belong to a known parametric distribution function, second that we need to integrate out $(a_{\tau_i}, N_i, q_{i,A}, \varepsilon_i^A, N_i, m_{\tau_{i-1}})$ to arrive to the next-period filtering distribution, $f(m_{\tau_i} \mid \mathcal{D}_i)$. Monte-Carlo is the method of choice for solving these high-dimensional integrals. In particular, we turn to particle filtering (for references, see, e.g., Doucet et al., 2001) and use a swarm of simulated points to approximate $f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1})$.[22]

Assume that an equal-weighted sample of M points represent $f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1})$ and denote these points by $[m_{\tau_{i-1}}]^{(m)}, m = 1, ..., M$. These points are propagated through (11) using importance sampling. We extend each particle by drawing from a proposal distribution

$$\left([a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}, [\varepsilon_{i,A}]^{(m)}\right) \sim r\left(a_{\tau_i}, m_{\tau_i}, N_i, q_{i,A}, \varepsilon_i^A \mid [m_{\tau_{i-1}}]^{(m)}\right)$$

To compensate for the difference between our proposal distribution and the target distribution, we need to attach importance weights, $w_i^{(m)}$, to the sampled particles. These weights are equal to the likelihood ratio between the target and proposal distributions:

$$w_i^{(m)} = \frac{1_{\{[a_{\tau_i}]^{(m)} \in [A_{\tau_i} - K, A_{\tau_i}]\}} f\left([a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}, [\varepsilon_{i,A}]^{(m)} \mid [m_{\tau_{i-1}}]^{(m)}\right)}{r\left([a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}, [\varepsilon_{i,A}]^{(m)} \mid [m_{\tau_{i-1}}]^{(m)}\right)}$$

[22] See Johannes et al. (2008) for a recent financial application of particle filtering on jump-diffusion models.

If we normalize the weights such that they sum up to 1, we get $\pi_i^{(m)} = \frac{w_i^{(m)}}{\sum_{l=1}^{M} w_i^{(l)}}$. Accordingly, the weighted sample

$$\left(\pi_i^{(m)}; [a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}, [\varepsilon_{i,A}]^{(m)}, [m_{\tau_{i-1}}]^{(m)}\right)$$

provides a correct characterization of the target.

The last step of the particle filter is to resample the particles proportional to the weights $\pi_i^{(m)}$. By doing this, we can get rid of particles with low likelihood and multiply the ones that are more likely. Resampling is crucial to avoid the usual curse of dimensionality in sequential importance sampling and ensures that the performance of our algorithm does not deteriorate with time, even if we keep the number of particles, M, fixed. To reduce Monte-Carlo variance, we use stratified sampling to draw indices ξ_m $(m = 1, ..., M)$ from a multinomial distribution with probabilities $\pi_i^{(m)}$. Hence, the particle set

$$\left([a_{\tau_i}]^{(\xi_m)}, [m_{\tau_i}]^{(\xi_m)}, [N_i]^{(\xi_m)}, [q_{i,A}]^{(\xi_m)}, [\varepsilon_{i,A}]^{(\xi_m)}[m_{\tau_{i-1}}]^{(\xi_m)},\right), m = 1, ..., M$$

provides an equal-weighted representation of the target distribution and we can proceed to the next observation.

Under weak assumptions, the empirical distribution function of the particles converges to the filtering distribution as $M \to \infty$ (see, e.g., Del Moral, 2004 for a textbook treatment) and we can use the output of the algorithm to obtain any expectation over the hidden states.

The choice of the proposal distribution, $r\left(a_{\tau_i}, m_{\tau_i}, N_i, q_{i,A}, \varepsilon_i^A \mid m_{\tau_{i-1}}\right)$ is crucial for the efficiency of the filtering algorithm, especially in models with outliers, like the one we use in this paper. Appendix A describes in detail the proposal distribution we use and the resulting particle filtering algorithm.

4.2.2 ***Monte Carlo EM algorithm***

We now address the issue of computing the maximum likelihood (ML) estimates for the model parameters. The particle filtering algorithm described in the preceding section can generate the log-likelihood function for any fixed parameter values. However, it is ill-suited for finding the ML estimates because the log-likelihood function is inherently irregular with respect to the parameters even with the use of common random numbers. This irregularity arises from the resampling step required for any particle filter. We thus adopt an indirect approach to the ML estimation via the EM algorithm of Dempster et al. (1977). The EM algorithm is an alternative way of obtaining the ML estimates for incomplete data models, where incomplete data refers to the situation in which the model contains some random variables without corresponding observations. The EM algorithm involves two steps - expectation and maximization – and hence its name. The first step consists of writing down the complete-data log-likelihood function. Since it is not observable, one needs to compute its expected value by conditioning on the observed data in conjunction with some assumed parameter values. This completes the expectation step. Secondly, in the maximization step, one finds the new parameter values that maximize the expected complete-data log-likelihood function. The updated parameter values are then used to repeat the E- and M-step until convergence to the ML estimates.

However, the E-step of our model is complex and computed using the the particle filter. We are thus using a version of the Monte Carlo EM (MCEM) algorithm.[23] Casting optimization as an EM algorithm problem effectively circumvents the irregularity induced by the particle filter, because in the M-step the parameters used in the filter are fixed. In order to reduce the attached Monte-Carlo noise, note that we use fixed-lag smoothing following the advice in Olsson et al. (2008). Appendix B describes the complete data likelihood in our model and the MCEM algorithm.

Finally, it is worth noticing that, in our model, asymptotic standard error for the maximum likelihood estimate cannot be computed using the negative Hessian matrix nor using the inner product of the individual scores.[24] We thus apply the alternative estimator proposed by Duan and Fulop (2009), which uses the smoothed individual scores to compute the asymptotic error.

[23]For a general introduction to the MCEM algorithm, see, for instance, Wei and Tanner (1990).

[24]None is directly computable because the individual observed log-likelihood function is highly irregular with respect to θ due to using the particle filter.

5 Empirical Results

In this section we investigate the intra-daily seasonality of the fundamental variance and transaction costs, and analyze the nature of the stochastic variation in transaction costs. We also investigate the intra-daily pattern of the dispersion in dealers' private waiting costs as we believe that it is a relevant component of transaction costs in a rather illiquid limit order book.

5.1 Intra-daily variations

In the efficient price and transaction costs models described previously, the fundamental variance $\exp(\sigma^2_{\tau_i})$, the mean transaction costs $\exp(\mu_{\tau_i})$, and the variance in dealers' waiting costs $\exp(\sigma^2_{c,\tau_i})$ are allowed to vary during the day. For the functions $\sigma^2_{\tau_i}$, μ_{τ_i} and σ^2_{c,τ_i}, we adopt a flexible cubic spline function as previously mentioned, with intra-day timepoints at (6.30; 8.00; 9.30; 12.00; 15.00; 18.00) for U.S. and European CDS contracts and at (8.00; 9.00; 11.00; 12.30; 15.00; 18.00) for Japanese CDSs. We use the histograms described in Section 3.1.3 to guide the choice of the intra-daily partition in the spline. For U.S. and European CDSs, the first and the last timepoints correspond to the preopening (6.30-8.00) and overnight (18.00-6.30) periods. The remaining timepoints reflect the daytime trading, and post-close periods. For Japanese CDSs, the first timepoint is 9.00 since there is not enough data to estimate volatility or transaction costs during the preopening period. The Japanese partition then reflects the Tokyo Exchange rhythm: the Tokyo lunch break occurs between 11.00-12.30, while the closing is at 15.00.

Table 3 Panel A reports estimates and asymptotic standard errors for constant parameters. The standard deviation of data errors in basis points (σ_ϵ) is important for all the contracts, consistent with the presence of outliers. The estimates of the other parameters (λ and σ_J) suggest fat-tailed distributions in the innovation of the efficient CDS premia. This result is consistent with the observation that CDS premia are prone to sudden large changes, presumably due to the arrival of important firm specific news.

Table 3 Panel B, C and D presents intra-daily estimates of the variance of the efficient

CDS premia (σ^2), the mean transaction cost (μ) and the variance of dealers' waiting costs (σ_c^2) for U.S., European and Japanese CDSs. Panel B reports results for the US CDSs and Figure 4 depicts the point estimates. Note that, for U.S. CDSs, our sample period is truncated at August 1, 2007, due to the lack of data thereafter. The results suggest that the intra-daily fundamental variance of U.S. contracts exhibit a flattened reverse J-shape pattern. In particular the fundamental variance spikes upward early during the preopening period (6.30 New York GMT), diminishes over the course of the following hour. It tends to increase very slightly during the post-close (after 15.00). This result is in accordance with what has been documented in virtually all U.S markets (see, e.g., Fleming, 1997, Foster and Viswanathan, 1990 or Chan, Christie and Schultz, 1995). Mean transaction costs exhibit a very clear reverse J-shape pattern with a bump at 18.00, while the largest spike for heterogeneity in dealers' waiting costs occurs also during the preopen and does not seem to increase at the end of the trading day.

Figure 5 depicts the point estimates for the European CDSs that are reported in Panel C. The intra-daily fundamental variance, mean transaction costs and variance of dealers' waiting costs clearly exhibit a reverse J-shaped pattern during the European business hours. Figure 5 reveals the presence of a higher fundamental variance during the preopen (6.30) and a subsequent decline at the beginning of the trading period (8.00). The fundamental variance remains at this low level or even lower during the rest of the day. In particular, there is no clear evidence of an increase in the fundamental variance at the end of the day, unlike findings in other European equity markets (see, e.g., Hillion and Suominen, 2004, or Michayluk and Sanger, 2006). This pattern is, however, also reported in Niemeyer and Sandås (1995). Transaction costs and dispersion in dealers' waiting costs are also largest during the preopening (6.30), and decrease at the beginning of the trading day (8.00). They remain relatively stable during the trading period, and increase during the post-close although they do not reach the level of the morning peak.

The estimates for Japanese CDSs are reported in Panel C, and presented graphically in Figure 6. The fundamental variance exhibit a reverse J-shaped pattern across the 6 Japanese CDS contracts in accordance with regularities documented in the Japanese equity market (see, e.g., Lehmann and Modest, 1994, or Hamao and Hasbrouck, 1995), or in the Tokyo foreign exchange market (Ito, Lyons, and Melvin, 1998). Mean transaction costs are also J-shaped over the trading day with a slight increase during lunch hours, consistent with studies related to the Japanese securities market (see, e.g., Amihud and Mendelson 1991, or Lehmann and Modest, 1994). Note that there is no very clear pattern for heterogeneity in dealers' private value.

Numerous papers try to explain U-shaped and J-shaped patterns in volatility and number of trades documented in securities market. The first hypothesis is informed trading. At the beginning of the trading day, as information is not yet incorporated into prices, volatility is high, and dealers widen their spreads as protection against larger adverse selection risks. In addition, Admati and Pfleiderer (1988) show that traders who have discretion over the timing of their trades choose to concentrate their trades in time so as to minimize adverse selection costs created by informed trading. As a result, trades cluster during certain periods of time, meanwhile bid-ask spreads are reduced.

The second hypothesis is related to inventory control consideration. First, dealers attempt to close exposed positions before the end of the trading day, which leads them to trade heavily before the overnight trading halt, explaining the concentration of trades before the close. Secondly, as argued by Brock and Kleidon (1992), the absence of opportunity to trade during the night brings about the need to reestablish optimal holdings in portfolios as soon as the trading starts.[25]

The J-shaped pattern of the fundamental volatility observed across all contracts tends to support the information asymmetry hypothesis during the preopening period, while the small but significant increase in both transaction costs and dispersion in dealers' waiting costs during the post-close for European and Japanese contract seems to corroborate the hypothesis of inventory control at the end of the day. We might fail, however, to find such inventory motives for U.S. CDSs.

[25] Brock and Kleidon (1992) predict that transaction costs get larger when the number of trades clusters, due to the presence of a monopolist specialist setting wider spreads in periods of less-elastic liquidity demands like those at the open and at the close. IDB have a different market organization, that is probably the reason why our findings are at odd with this prediction.

5.2 Tests for intradaily variations

We need to test if the intra-daily patterns suggested in Figures 4, 5 and 6 are statistically significant. Therefore, we compute the difference between two adjacent intradaily ML estimates for the fundamental variance, variance in dealers' waiting costs and mean transactions costs and test whether this difference is different from zero. Table 3 Panel B, Panel C and Panel D list the successive differences to the previous intradaily timepoints for the 18 North American, European, and Japanese CDS contracts, respectively.

Our tests confirm several broad regularities documented above. First, across most CDS names, the fundamental variance is at its highest level during the preopening period (6.30 local time), and subsequently declines the following period, which results in a significant negative difference at 8.00. The peak of the fundamental variance at 6.30 is consistent with the accumulated overnight information and trading halt effect suggested by Brock and Kleidon (1992). We fail, however, to find a systematic higher volatility at the end of the trading day for U.S. and European CDSs, as would predict Brock and Kleidon's (1992) model. The fundamental variance exhibits, however, a significant increase (positive sign at 18.00) across all Japanese CDSs.

Secondly, the reverse J-shape pattern for mean transaction costs is also confirmed: mean transactions costs are high across all contracts during the preopening period, and decline significantly at 8.00 for all U.S. and European contracts and at 11.00 for Japanese contracts, when the quoting and trading activity is high in the IDB limit order book. The J-shape pattern is flattened for U.S. contracts in the sense that only 3 contracts over 6 exhibit a significant positive difference at 18.00 (compared to 15.00). In contrast, we note the presence of a little but significant hump during lunch hours across most of European and Japanese CDSs, and a clear spike at the end of the trading day: all European and 5 over 6 Japanese CDSs show significant positive difference at 18.00.

Third, tests regarding the intra-daily variations in the variance of dealers' waiting costs reveal that, for most of CDS contracts, it decreases significantly during the morning (8.00 for U.S. CDSs, 9.30 for European CDSs, and 11.00 for Asian). For European and Japanese

CDSs, it increases again significantly during lunch hours (12.00 or 12.30), along with mean transaction costs for those contracts. Finally, the dispersion in dealers' waiting costs might increase again at the end of the trading day for European and Japanese CDSs but not for U.S. contracts.

5.3 Correlation between the filtered estimates and market activity

Although the model does not include any trade variables, it is relevant to examine the correlations between some filtered estimates and trade-related variables described in Section 3.1.3, i.e. the number of quotes and the number of trades. Table 5 presents the correlations between mean transaction costs on the ask-side (a proxy for the best ask), the dispersion in quotes on the ask side (a proxy for the ex ante expected heterogeneity in dealers' private costs), innovation in the efficient premia (a proxy for the fundamental volatility), and two measures of market activity (the number of ask quotes and the number of trades).[26]

First, our results show that the number of trades is highly positively associated with the number of quotes, consistent with dealers consuming liquidity when the order book is the least thin, as reported in Biais, Hillion and Spatt (1995). Second, the number of trades is positively correlated to the fundamental volatility as in Jones, Kaul and Lipson (1994) or Huang and Masulis (2003). Third, the number of trades is negatively associated with mean transaction costs, consistent with dealers concentrating their trades when the cost of trading is lower as predicted by Admati and Pfleiderer (1988).

Results reported in Table 5 also show that mean transactions costs in the order book are positively correlated to the fundamental volatility, which is consistent with higher adverse selection costs7. This relation is also predicted by Foucault (1999). Transaction costs are also highly positively correlated with the heterogeneity in dealers' private value. This result underlines the importance of considering investors' heterogeneity (private trading needs, urgency to hedge, etc.) as an important component of transaction costs, at least,

[26]In unreported results, we find very similar correlations by using the bid side.

in the CDS market. In addition, mean transaction costs are negatively correlated to the number of ask quotes, i.e. transaction costs are larger when the order book is thin or almost empty, consistent with other findings (see, e.g., Biais, Hillion and Spatt, 1995 or Niemeyer and Sandås, 1995).

5.4 What is the impact of the subprime crisis?

This section addresses the concern whether the subprime mortgage crisis which became apparent end of July 2007 could have an impact on the stylized facts reported above.

First, we divide the period into two sub-periods according to whether data are recorded before or after August 1, 2007 (see, e.g., Gorton (2008) for a chronology of the events of the 2007 credit crisis). We drop U.S. and Japanese CDSs because we do not have enough observations in the second sub-period to run clean estimations of the model. We thus use the European CDSs to investigate whether our results are sensitive to the credit crisis. Second, we run the model described in section 4.1 on the two different sub-periods. Third, we compute the differences between the ML estimates we get on both sub-periods and test whether these differences are significant. Finally, we also report differences in market activity.

Table 5 presents the results. They show a significant reduction in the daily number of trades after the crisis, while the daily number of quotes does not seem to be unidirectionally affected. There is an increase in the CDS prices volatility and in the heterogeneity of dealers' private costs. There is also a deterioration in mean transaction costs. In unreported results, we find similar, though weaker, intra-daily patterns after the onset of the crisis, which tends to support that the crisis does not explain or change our results.

6 Conclusion

This paper attempts to provide a comprehensive overview of the intradaily patterns of volatility and transaction costs in the CDS interdealer market. We analyze a sample of 18 single-name 5-year CDS contracts whose underlying reference firm may be North American, European or Japanese during a 24-month period: March 31, 2006 through March 31, 2008. Data are from GFI, one of the leading inter-dealer broker, from which we retrieve the most active contracts in terms of bid and ask quotes.

The trading and quoting activity exhibit U-shaped patterns for the European and Japanese contracts, and a J-shaped for U.S. CDSs. The number of trades does not exceed 4 times a day, the number of simultaneous bid and ask quotes is in average 17 pairs per day for European CDSs, and lower for U.S. or Japanese (5 bid/ask pairs) contracts. We find that the mean cost of a round-trip is 58 bp across all contracts. The market for single-name CDSs is thus rather illiquid.

Due to data scarcity and the absence of any information on the best limits that prevail at any moment in the IDB order book, we propose a state-space model of bid and ask quotes based on Hasbrouck (1999) to estimate the volatility of the efficient CDS premium and stochastic transaction costs. The framework allows for data errors, discreteness and jumps in the efficient price. We estimate the model using particle filtering and the Monte Carlo EM algorithm. First, we find that the volatility of the efficient CDS premium is at its highest during the preopening period (6.30), and subsequently declines. Second, the parameter estimates show a clear J-shaped pattern for mean transaction costs across all contracts. The simultaneous peak in volatility and transaction costs early during the morning is consistent with a price discovery/informed trading hypothesis. In addition, the higher number of trades and larger heterogeneity in dealers' private value (e.g., hedging needs) at the end of the trading day also provide supporting, though weaker, evidence for inventory control by dealers at the end of the trading day. The way swap dealers manage their inventory might be an explanation of this weaker effect. We believe that, due to the high transaction costs and thin trading in the CDS market, a swap dealer aggregates her trades from retail clients and trades the net in a single transaction in the interdealer market to get a position in her book close to neutral at the end of the day. She also might use synthetic instrument such as iTraxx credit indices to offset positions on a basket of credits through a single trade.

We find that mean transaction costs are negatively related to the number of trades, and positively to the fundamental volatility and to the heterogeneity in dealers' private value, as predicted by models of limit order book (e.g., Foucault, 1999). We also find that our results are robust to the onset of the subprime crisis, although we find a slight drop in the number of trades and an increase in the fundamental volatility.

Besides contributing to document trading and liquidity patterns in the market for single-name CDSs, this paper also provides a methodology adapted to very thin limit order books to estimate market quality measures, such as volatility and transaction costs. We think this is a promising avenue to analyze liquidity in other thinly traded financial instruments. In particular, it may be of interest to investigate the case of CDS indices, for which transaction costs might be lower, liquidity higher and could encourage more hedging and inventory management behavior and might be less sensitive to informed trading.

Appendix A: Details of the proposal distribution and of the particle filtering algorithm

We use a proposal distribution with three important features.

1. We sample a_{τ_i} from a uniform distribution with a support $[A_{\tau_i} - K, A_{\tau_i}]$. This ensures that the indicator function $1_{\{a_{\tau_i} \in [A_{\tau_i} - K, A_{\tau_i}]\}}$ is one.

2. We use stratified sampling for the jumps in the efficient CDS premia, N_i and the data errors q_i^A. The probability of sampling from the jumps in the CDS premia is $\widetilde{\lambda}$ instead of λ, while the probability of the data errors is some $\widetilde{p}$ instead of p. This ensures that we sample from these events, even if their probability is low. Denote the corresponding sampling probabilities $g^N(N_i)$ and $g^{q_A}(q_{i,A})$.

3. To generate the proposal for m_{τ_i}, we use a linear approximation to equation (7). In particular consider the following relationship

$$a_{\tau_i} = e^{m_{\tau_{i-1}}} + e^{m_{\tau_{i-1}}}(m_{\tau_i} - m^*) + e^{\mu_{\tau_i} + m_{\tau_{i-1}}} + c_{i,A} + \varepsilon_{i,A},$$

   where $c_i^A \sim N(0, \exp(\sigma^2_{c,\tau_i}) \times \exp(2m_{\tau_{i-1}}))$. This is an approximation of (7) with three modifications: (i) we use a first order Taylor expansion around $m_{\tau_{i-1}}$ to linearize $M_{\tau_i} = e^{m_{\tau_i}}$; (ii) we substitute $m_{\tau_{i-1}}$ in the place of m_{tau_i} in the determination of κ_{τ_i} ; (iii) we substitute $m_{\tau_{i-1}}$ in the place of m_{tau_i} in the determination of the variance of c_i^A. Note that conditional on $(m_{\tau_{i-1}}, N_i, q_{i,A})$, a_{τ_i} and m_{τ_i} are jointly normal. We can thus easily sample from m_{τ_i} given a_{τ_i}. Denote the density of the importance sampler as $g^m(m_{\tau_i} \mid a_{\tau_i}, m_{\tau_{i-1}}, N_i, q_{i,A})$.

Assume that we have M particles, $[m_{\tau_{i-1}}]^{(m)}$ representing $f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1})$. Then our particle filter with M particles consists of the following steps:

Step 1: Draw $[a_{\tau_i}]^{(m)}$ from a uniform distribution with a support $[A_{\tau_i} - K, A_{\tau_i}]$.

Step 2: Enlarge the state-space with the jumps in the system by sampling from N_i and $q_{i,A}$ using probabilities $\widetilde{\lambda}$ and $\widetilde{p}$ respectively.

Step 3: Sample from m_{τ_i} using a proposal derived from the linearized approximation of the system, drawing

$$[m_{\tau_i}]^{(m)} \sim g^m(m_{\tau_i} \mid [a_{\tau_i}]^{(m)}, [m_{\tau_{i-1}}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}).$$

Step 4: Sample from ε_i^A given $(a_{\tau_i}, m_{\tau_i}, q_{i,A})$ using joint normality

$$[\varepsilon_i^A]^{(m)} \sim g^{\varepsilon^A}(\varepsilon_i^A \mid [a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [q_{i,A}]^{(m)}).$$

Step 5: Attach importance weights to the particles

$$w_i^{(m)} = \frac{f\left([a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}, [\varepsilon_{i,A}]^{(m)} \mid [m_{\tau_{i-1}}]^{(m)}\right)}{r\left([a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}, [\varepsilon_{i,A}]^{(m)} \mid [m_{\tau_{i-1}}]^{(m)}\right)}.$$

Here, the proposal density takes the form

$$\begin{aligned}
& r\left([a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}, [\varepsilon_{i,A}]^{(m)} \mid [m_{\tau_{i-1}}]^{(m)}\right) \\
= \;& g^{\varepsilon^A}([\varepsilon_i^A]^{(m)} \mid [a_{\tau_i}]^{(m)}, [m_{\tau_i}]^{(m)}, [q_{i,A}]^{(m)}) g^m([m_{\tau_i}]^{(m)} \mid [a_{\tau_i}]^{(m)}, [m_{\tau_{i-1}}]^{(m)}, [N_i]^{(m)}, [q_{i,A}]^{(m)}) \\
\times \;& g^N([N_i]^{(m)}) g^{q_A}([q_{i,A}]^{(m)}).
\end{aligned}$$

Step 6: Resample the particle set according to the probability $\pi_i^{(m)} = \frac{w_i^{(m)}}{\sum_{m=1}^M w_i^{(m)}}$ to yield M equal-weighted particles. This equal-weighted particle set is an empirical representation of $f(m_{\tau_i}, N_i, q_{i,A}, \varepsilon_i^A, m_{\tau_{i-1}} \mid \mathcal{D}_i)$.

Note that the particle filter provides a sample on the entire past of the system up to τ_i. Any quantity of interest based on the past particles can be computed and carried forward alongside with m_{τ_i}. This is true because at any time τ_i, m_{τ_i} is sufficient for moving the algorithm forward. Denote by I_i the quantity whose distribution is of interest; for example, one may be interested in $I_i = (m_{\tau_0} + m_{\tau_1} + \cdots + m_{\tau_i})/(i+1)$. Then, in all of the preceding derivations one can use the vector (I_i, m_{τ_i}) in place of m_{τ_i}. Conditional on m_{τ_i}, the system's forward evolution is independent of I_i, and thus the algorithm remains unchanged. The output of the filter at any time τ_i will be a set of particles representing the joint filtering distribution of $(I_i, , m_{\tau_i})$.

The case when we have a bid observation or both is very similar to the situation described above, so to conserve space we do not spell out the details.

Appendix B: Monte Carlo EM algorithm

First, we need a complete data representation of the fundamental innovations. Here, the complete data space we choose consists of the asset jump indicator N_i and of the log efficient CDS premia, m_{τ_i}. The joint conditional likelihood can be written as

$$f(m_{\tau_i}, N_i \mid m_{\tau_{i-1}}) = f(m_{\tau_i} \mid, N_i, m_{\tau_{i-1}}) f(N_i).$$

The loglikelihood of the number of jumps[27] is

$$L_i^N = 1_{\{N>0\}} \ln(\lambda) + 1_{\{N=0\}} \ln(1-\lambda). \tag{12}$$

Accordingly, we can write the loglikelihood of the efficient premia innovation as

$$\begin{aligned} L_i^m &= 1_{\{N_i=0\}} \times \left[-\frac{\left[m_{\tau_i} - m_{\tau_{i-1}}\right]^2}{2e^{\sigma^2_{\tau_{i-1}}} \Delta\tau_i} - \frac{\ln\left(e^{\sigma^2_{\tau_{i-1}}} \Delta\tau_i\right)}{2} \right] \\ &+ 1_{\{N_i=1\}} \times \left[-\frac{\left[m_{\tau_i} - m_{\tau_{i-1}}\right]^2}{2e^{\sigma^2_{\tau_{i-1}}} \Delta\tau_i \left(1+\sigma_J^2\right)} - \frac{\ln\left(e^{\sigma^2_{\tau_{i-1}}} \Delta\tau_i \left(1+\sigma_J^2\right)\right)}{2} \right]. \end{aligned} \tag{13}$$

Second, the remaining part of the complete data loglikelihood represents the parameters for the transaction costs (μ_{τ_i} and c_i^A) and the data errors ε_i^A. Introducing a scaled version of the data error, $\widetilde{\varepsilon}_i^A = \varepsilon_i^A / \sigma_\varepsilon$, the loglikelihood of the observation is

$$L_i^a = -\frac{\left[a_{\tau_i} - e^{m_{\tau_i}} - e^{\mu_{\tau_i} + m_{\tau_i}} - \sigma_\varepsilon \widetilde{\varepsilon}_i^A\right]^2}{2\sigma_c^2 e^{2m_{\tau_i}}} - \frac{\ln\left(\sigma_c^2\right)}{2}. \tag{14}$$

Note that in this paper we set the probability of data errors, p, exogenously so the probability of $q_{i,A}$ is not related to the parameters we estimate. Further, the loglikelihood of $\widetilde{\varepsilon}_i^A$ does not depend on unknown parameters. We can thus ignore it.

The MCEM algorithm can be summarized as follows: (1) Set some initial parameter values, $\theta^{(0)}$; (2) Repeat the following E- and M-steps until convergence.

E-step: Get the conditional expectation of the sufficient statistics in the complete data loglikelihood function of the previous set of parameters. In particular we need to

[27] Throughout we use the convention $0 \times \ln 0 = 0$

approximate expectations of the form

$$E(X \mid \mathcal{D}_T, \theta^{(k-1)})$$

where X_i is some function of the hidden random states, $\mathcal{D}_T$ is the observed data and θ^{k-1} is the parameter vector at the $(k-1)^{th}$ iteration. The particle filter described in section 4.2 can be used to compute these quantities. We run the filter using the parameters $\theta^{(k-1)}$ to generate the particle set that represents the smoothed distribution for X_i. The m-th particle is denoted by $X_{i|T}^{(m)}$. Accordingly, the expectation can be approximated by the sample average as follows:

$$E(X_i \mid \mathcal{D}_T, \theta^{(k-1)}) \approx \frac{1}{M} \sum_{m=1}^{M} X_{i|T}^{(m)}$$

When the sample size T is large, undesirable Monte-Carlo noise will be introduced by the use of the smoothed distribution. Intuitively, the particle filter always adapts to the newest observation, and thus its representation of the distant past is bound to be poor. Olsson et al. (2008) suggest to use the information only up to $i + L$ when computing any quantity that involves the unobserved state variable at time i. The rationale is the forgetting property expected of the dynamic system; that is, for large enough L, the distribution for the unobserved state variable at time i conditional on the information up to $i + L$ will be almost identical to that conditional on the entire sample. Olsson et al. (2008) thus propose to use fixed-lag smoothing by using information only up to $i + L$. Adopting fixed-lag smoothing leads to our approximation as follows:

$$\begin{aligned} & E(X_i \mid \mathcal{D}_T, \theta^{(k-1)}) \\ \approx \;& E(X_i \mid \mathcal{D}_{(i+L)\wedge T}, \theta^{(k-1)}) \\ \approx \;& \frac{1}{M} \sum_{m=1}^{M} X_{i|(i+L)\wedge T}^{(m)} \end{aligned}$$

M-step: Maximize the conditional expected value of the complete-data log-likelihood function obtained in the E-step. First, denote the i^{th} complete-data loglikelihood by $L_i(\theta)$

$$L_i(\theta) = L_i^N(\theta) + L_i^m(\theta) + 1_{(i \in I_A)} L_i^a(\theta) + 1_{i \in I_B} L_i^b(\theta). \tag{15}$$

Here, I_A (I_B) denotes the set where an ask (bid) is observed. Secondly, the M-step can be written as follows

$$\theta^{(k)} = arg\max_{\theta}\sum_{i=1}^{T} E(L_i(\theta) \mid \mathcal{D}_{i|(i+L)\wedge T}, \theta^{(k-1)}).$$

Our model belongs to the exponential family, meaning that the conditional loglikelihood can be written as

$$E(L_i(\theta) \mid \mathcal{D}_{i|(i+L)\wedge T}, \theta^{(k-1)}) = f\left(\theta, E(s_i(X_i) \mid \mathcal{D}_{i|(i+L)\wedge T}, \theta^{(k-1)})\right),$$

where $s_i(X_i)$ are some sufficient statistics for the parameters. This is computationally important, because we only need to run the filter once per iteration to obtain an estimate of $E(s_i(X_i) \mid \mathcal{D}_{i|(i+L)\wedge T}, \theta^{(k-1)})$.

References

[1] Acharya, Viral V., and Timothy C. Johnson, 2007, Insider trading in credit derivatives, Journal of Financial Economics 84, 110–141.

[2] Admati, Anat R., and Paul Pfleiderer, 1988, A Theory of intraday patterns: volume and price variability, Review of Financial Studies 1, 3–40.

[3] Amihud, Yakov and Haim Mendelson, 1991, Volatility, efficiency, and trading: evidence from the japanese stock market, Journal of Finance 46, 1765–1789.

[4] Ashcraft, Adam B., and Joao A.C. Santos, 2009, Has the CDS market lowered the cost of corporate debt? Journal of Monetary Economics 56, 514-523

[5] Barclay, Michael J., and Terrence Hendershott, 2003, Price discovery and trading after hours, Review of Financial Studies 16, 1041-1073.

[6] Biais, B., and F. Declerck, 2007, Liquidity, competition and price discovery in the European corporate bond market, Mimeo, Toulouse School of Economics.

[7] Biais, Bruno, Pierre Hillion, and Chester Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, Journal of Finance 50, 1655–1689.

[8] Bjønnes, G. H., and D. Rime, 2005, Dealer behavior and trading systems in the Foreign Exchange market, Journal of Financial Economics 75, 571–605.

[9] Blanco, R., S. Brennan, and I. W. Marsh, 2005, An empirical analysis of the dynamic relationship between investment grade bonds and credit default swaps, Journal of Finance 60, 2255–2281.

[10] Bongaert, Dion, Frank de Jong, and Joost Driessen, 2008, Derivative pricing with liquidity risk: Theory and evidence from the Credit Default Swap Market, SSRN.

[11] Boni, Leslie and Chris Leach, 2004, Expandable orders in the brokered interdealer market for U.S. Treasury securities, Journal of Financial Markets 7, 145–185.

[12] Brock, William A., and Allan W. Kleidon, 1992, Periodic market closure and trading volume, Journal of Economic Dynamics and Control 16, 451–489.

[13] Chan, K., Chung, P. and Fong, W.M , 2002, The informational role of stock and option volume, Review of Financial Studies 15, 110–141

[14] Chan, K. C., William G. Christie, and Paul H. Schultz, 1995, Market structure and the intraday pattern of bid-ask spreads for NASDAQ securities, Journal of Business 68, 35–60.

[15] Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2001, Market liquidity and trading activity, Journal of Finance, 56, 501–530.

[16] Del Moral, P., 2004, Feynman-Kac Formulae Genealogical and Interacting Particle Systems with Applications, Springer, New York.

[17] Dempster, A.P., N.M. Laird, and D.B. Rubin, 1977, Maximum likelihood from incomplete data via the EM algorithm, Journal of the American Statistical Association, Series B 39, 1–38.

[18] Doucet, A., de Freitas, N. and Gordon, N., eds (2001), Sequential Monte Carlo Methods in practice, Springer-Verlag.

[19] DSouza, Chris, Charles Gaa, and Jing Yang, 2003, An empirical analysis of liquidity and order flow in the brokered interdealer market for Government of Canada bonds, Bank of Canada Working Paper 2003-28.

[20] Duan, Jin C. and Andras Fulop, 2009, A stable estimator for the information matrix under EM, Forthcoming Statistics and Computing

[21] Duffie, Darrell, 2007, Innovations in credit risk transfer: implications for financial stability, mimeo.

[22] Fleming, Michael J., 1997, The round-the-clock market for U.S.Treasury securities, FRBNY Economic Policy Review, 9–32.

[23] Edwards, Amy K., Larry E. Harris, and M.S. Piwowar, 2007, Corporate bond market transaction costs and transparency, Journal of Finance 62, 1421–1451.

[24] Edwards, Amy K., M. Nimalendran, and M. S. Piwowar, 2006, Corporate bond market transparency: Liquidity concentration, informational efficiency, and competition, working paper, the U.S. Securities and Exchange Commission.

[25] European Central Bank (ECB), 2009, Credit default swaps and counterparty risk, Working paper.

[26] Foster, F. Douglas, and S. Viswanathan, 1990, A theory of interday variations in volumes, variances and trading costs in securities market, Review of Financial Studies 3, 593–624.

[27] Foucault, Thierry, 1999, Order flow composition and trading costs in a dynamic limit order market, Journal of Financial Markets 2, 99–134.

[28] Foucault, Thierry, Ohad Kadan and Eugene Kandel, 2005, Limit order book as a market for liquidity, Review of Financial Studies 18, 1171–1217.

[29] Glosten, Larry R., 1994, Is the electronic open limit order book inevitable? Journal of finance 49, 1127–1161.

[30] Goettler Ronald, Christine A. Parlour, and Uday Rajan, 2007, Informed traders and limit order markets, mimeo.

[31] Hamao, Yasushi, and Joel Hasbrouck, 1995, Securities trading in the absence of dealers: Trades and quotes on the Tokyo stock exchange, Review of Financial Studies, 8, 849–878.

[32] Hasbrouck, Joel, 1999, The dynamics of discrete bid and ask quotes, Journal of Finance 54, 2109-2142.

[33] Hillion, Pierre, and Matti Suominen, 2004, The manipulation of closing prices, Journal of Financial Markets 7, 351–375.

[34] Ho T., and H. Stoll, 1983, The dynamics of dealer markets under competition, Journal of Finance 49, 1639-1664.

[35] Huang, Roger D., 2002, The quality of ECN and Nasdaq market-maker quotes, Journal of Finance 57, 1285-1319.

[36] Huang R., J. Cai, and X. Wang, 2002, Information-based trading in Treasury Note interdealer broker market, Journal of Financial Intermediation 11, 269-296.

[37] Huang, Roger D., and Ronald W. Masulis, 2003, Trading activity and stock price volatility: evidence from the London Stock Exchange, Journal of Empirical Finance 10, 249–269.

[38] Hull, John, Mirela Predescu, and Alan White, 2004, The relationship between credit default swap spreads, bond yields, and credit rating announcements, Journal of Banking and Finance 28, 2789-2811.

[39] Ito, Takatoshi, Richard K. Lyons, and Michael T. Melvin, 1998, Is there any private information in the FX market? The Tokyo experiment, Journal of Finance, 53, 1111–1130.

[40] Johannes M, Polson N, Stroud J., 2008, Learning about jumps and stochastic volatility: Filtering stochastic differential equations with jumps. Forthcoming Review of Financial Studies

[41] Jones, Charles, Gautam Kaul, and Marc Lipson, 1994, Transactions, volume and volatility, 1994, Review of Financial Studies 7, 631-651.

[42] Lehmann, Bruce, and David M. Modest, 1994, Trading and liquidity on the Tokyo stock exchange: a bird's eye view, Journal of Finance 49, 951–984.

[43] Michayluk, David, and Gary C. Sanger, 2006, Day-end effect on the Paris Bourse, Journal of Financial Research, 29, 131–146.

[44] Minton, Bernadette, René M. Stulz, and Rohan Williamson, How muvh do banks use credit derivatives to hedge loans, Fisher College of Business WP 2008-03-001.

[45] Niemeyer Jonas, and Patrik Sandås, 1995, An empirical analysis of the trading structure at the Stockholm Stock Exchange, Working Paper No. 44, Stockholm School of Economics.

[46] Olsson, J., Cappe, R., Douc, R. and Moulines E., 2008, Sequential Monte Carlo smoothing with application to parameter estimation in non-linear state space models, Bernoulli 14, 155–179.

[47] Reiss P. C., and I. Werner, 1998, Does risk sharing motivate interdealer trading? Journal of Finance 53, 1657-1703.

[48] Reiss P. C., and I. Werner, 2004, Anonymity, adverse selection and the sorting of interdealer trades, Review of Financial Studies 18, 599-636

[49] Roşu, Ionid, 2009, A dynamic model of the limit order book, Review of Financial Studies 22, 4601–4641.

[50] Tang, Dragon Yongjun and Hong Yan, 2007, Liquidity and Credit Default Swap spreads, SSRN.

[51] Wei, G.C.G. and M.A. Tanner, 1990, A Monte Carlo Implementation of the EM Algorithm and the Poor Man's Data Augmentation Algorithms, Journal of the American Statistical Association 85, 699–704.

Table 1

Most representative industries

This table reports the top 20 industries in the raw CDS dataset for the period March 31, 2007 to March 31, 2008, by number of CDS contracts and by percentage of quotes and trades (data points). The dataset is from GFI, one of the leader inter-dealer CDS broker, and consists of a total of 1,944 CDS contracts.

20 most active CDS industries		
Industry	By number of CDS contracts	By % of (quotes and trade) data
Banks	294	15.12%
Financial Services	112	5.76%
Electricity	90	4.63%
Sovereign	80	4.12%
Real Estate	74	3.81%
Major Integrated oil/gas company	65	3.34%
Fixed Line Telecom	61	3.14%
Municipal	57	2.93%
Chemicals Specialty	48	2.47%
Insurance Full Line	45	2.31%
Wireless Telecom	39	2.01%
Auto Manufacturers	33	1.70%
Food Products	30	1.54%
Electrical Equipment	28	1.44%
Multi line Retailers (non Food)	28	1.44%
Television	28	1.44%
Pharmaceutical	27	1.39%
Consumer Electronics	25	1.29%
Railways	24	1.23%
Publishing	24	1.23%

Table 2

Summary statistics

This table presents summary statistics for the CDS data used in this study. We use data from GFI, one of the main inter-dealer CDS broker, from which we retrieve time-stamped (down to the second) information on bid quotes, ask quotes, and transaction prices during the period from March 31, 2006 to March 31, 2008. We use a sample of 18 5-year CDSs: 6 U.S. contracts (American Axle & Manufacturing, Ford Motor Corp, General Motor Corp, Goodyear Tire, Lear and Visteon), 6 European CDSs (Casino Guich, Deutsche Telecom, ITV Plc, Portugal Telecom Intl Finance, Telecom Italia and Valeo), and 6 Japanese CDSs (ACOM, Aiful, Orix Corp, Promise Co, Softbank Corp. and Takefuji). These contracts have been selected on the highest number of data points (number of quotes). The table presents variables for each contract. Credit rating is retrieved from Reuters based on ratings of senior unsecured long term issuer as of 01/01/2007 (by S&P when available, or Moody's). Number of trades is the total number of transaction prices we observed during the sample period for the CDS contract. Daily average number of trades is the number of trades divided by the number of days during which at least one trade occurs. Minimum tick is the minimum difference between two consecutive prices (transaction price or quotes). Number of bid (ask) quotes is the total number of bid (ask) quotes observed during the period. Daily average number of bid (ask) quotes is the total number of bid (ask) quotes divided by the number days during which at least one bid (ask) quotes is registered. Number of bid/ask pairs is the total number of "simultaneous" bid and ask quotes we observe in the limit order book. Daily average number of bid and ask pairs is the number of bid/ask pairs divided by the number of days during which at least one pair is registered. Average midpoint is (ask + bid)/2, calculated when we observe a pair of quotes, and averaged over the sample period. Average bid/ask spread is a "quoted" spread (ask - bid) calculated when we observe a pair of quotes, and averaged over the sample period. Risky annuity factor is the cost of the 5-year risky annuity for the 5-year CDS contract using a 5% interest rate and a 40% recovery value. Average cost of a round-trip is the present value of the stream of individual cash-flows (equal to the bid-ask spread), whose frequency is the frequency of the CDS, and whose duration is the lower of the maturity of the CDS and the default time of the underlying reference: it is calculated by multiplying the risky annuity factor and the average bid/ask spread. Average relative cost of a round-trip is the average cost of a round-trip divided by the average midpoint. Panel A presents the summary statistics for the 6 U.S. CDSs. Panel B presents the summary statistics for the 6 European CDSs. Panel C presents the summary statistics for the 6 Japanese CDSs.

Panel A. U.S. CDSs

	American Axle & Manufacturing	**Ford Motor Corp**	**General Motors Corp**	**Goodyear Tire**	**Lear**	**Visteon**	*Mean*
Credit Ratings	BB	B	B	B+	B+	B+	
Number of trades	258	857	961	145	260	217	*449,7*
Daily average number of trades	0,8	2,6	3,2	0,5	0,9	0,8	*1,5*
Minimum tick	1	0,5	1	1	1	2	
Number of bid quotes	1 220	2 858	2 824	867	1 163	985	*1 652,8*
Daily average number of bid quotes	3,9	8,6	9,3	3,2	4,0	3,7	*5,5*
Number of ask quotes	1 124	2 805	2 578	763	1 102	918	*1 548,3*
Daily average number of ask quotes	3,6	8,5	8,5	2,8	3,7	3,5	*5,1*
Number of bid and ask pairs	974	2 151	2 063	674	939	786	*1 264,5*
Daily average number of bid and ask pairs	3,1	6,5	6,8	2,5	3,2	3,0	*4,2*
Average midpoint	372	612	488	321	445	530	*461,3*
Average bid-ask spread	20,9	15,3	13,4	19,6	23,1	28,0	*20,1*
Risky annuity factor	3,8	3,4	3,6	3,9	3,7	3,5	
Average cost of a round-trip	78,8	52,6	48,3	75,6	84,7	99,4	*73,2*
Average relative cost of a round-trip	0,21	0,09	0,10	0,24	0,19	0,19	*0,2*

Panel B. European CDSs

	Casino Guich	**Deutsche Telecom**	**ITV Plc**	**Portugal Telecom Intl Finance**	**Telecom Italia**	**Valeo**	*Mean*
Credit Ratings	BBB-	A-	BBB-	BBB-	BBB+	BBB(*)	
Number of trades	547	1 135	650	889	1 767	850	*973.0*
Daily average number of trades	1.1	2.3	1.3	1.8	3.5	1.9	*2.0*
Minimum tick	0.5	0.5	0.5	0.5	0.5	0.5	*0.5*
Number of bid quotes	6 806	10 491	9 001	10 485	13 403	7 394	*9 596.7*
Daily average number of bid quotes	14.0	21.0	18.7	21.5	26.9	16.4	*19.7*
Number of ask quotes	6 206	9 816	8 319	9 955	12 824	7 018	*9 023.0*
Daily average number of ask quotes	12.8	19.7	17.3	20.4	25.7	15.5	*18.6*
Number of bid and ask pairs	5 514	8 808	7 530	8 923	11 663	6 233	*8 111.8*
Daily average number of bid and ask pairs	11.4	17.7	15.6	18.3	23.4	13.8	*16.7*
Average midpoint	82	48	111	107	69	94	*85.2*
Average bid-ask spread	5.8	3.4	9.7	7.4	3.8	7.5	*6.3*
Risky annuity factor	4.2	4.3	4.2	4.2	4.3	4.2	*4.2*
Average cost of a round-trip	24.8	14.7	40.7	31.0	16.4	31.5	*26.5*
Average relative cost of a round-trip	0.30	0.30	0.37	0.29	0.24	0.33	*0.3*

Panel C. Japanese CDSs

	ACOM	**Aiful**	**Orix Corp**	**Promise Co**	**Softbank Corp.**	**Takefuji**	*Mean*
Credit Ratings	BBB+	BBB+	A-	BBB+	BB-	BBB	
Number of trades	143	380	182	136	79	210	*188.3*
Daily average number of trades	0.3	0.9	0.4	0.3	0.3	0.5	*0.5*
Minimum tick	0.5	0.5	0.25	0.25	0.5	0.25	*0.4*
Number of bid quotes	2 645	4 726	2 530	2 567	2 240	3 455	*3027.2*
Daily average number of bid quotes	6.3	10.6	6.1	6.4	7.2	8.6	*7.5*
Number of ask quotes	2 343	4 417	2 114	2 257	1 949	2 948	*2671.3*
Daily average number of ask quotes	5.6	9.9	5.1	5.6	6.3	7.3	*6.6*
Number of bid and ask pairs	1 899	3 682	1 719	1 808	1 676	2 506	*2215.0*
Daily average number of bid and ask pairs	4.5	8.3	4.1	4.5	5.4	6.2	*5.5*
Average midpoint	82	149	75	79	324	115	*137.3*
Average bid-ask spread	13.2	17.0	11.3	12.7	39.6	17.0	*18.5*
Risky annuity factor	4.2	4.1	4.3	4.3	3.8	4.2	*4.2*
Average cost of a round-trip	56.0	70.1	48.1	54.0	152.3	71.2	*75.3*
Average relative cost of a round-trip	0.68	0.47	0.64	0.68	0.47	0.62	*0.6*

Table 3

Estimation results for the efficient CDS price and transaction costs models

This table presents the ML estimates of the state-space model (1) and (9)-(10) based on bid and ask quotes of 18 5-year CDS contracts for March 31. 2006 through March 31. 2008. Panel A presents the ML estimates for the constant parameters (σ_C, λ, σ_J) for all CDS contracts. Panel B (resp. Panel C and Panel D) contains time-varying ML estimates for the fundamental variance (σ^2). the mean transaction costs (μ) and the quote dispersion (σ_c^2) for each intraday timepoint for the U.S. (resp. European and Asian) CDS contracts. Panel B reports also the difference between two adjacent ML estimates for each timepoint, and the number of contract for which the difference is significantly positive (at 1%) and significantly negative.

Panel A. Estimates for constant parameters (cont.)

	pooled U.S. CDSs	pooled European CDSs	pooled Japanese CDSs
σ_C	39.43	18.82	47.38
λ	0.06	0.05	0.06
σ_J	23.42	28.16	31.10

Panel B. Estimates for varying parameters - U.S. CDSs

	Time					
	6:30	8:00	9:30	12:00	15:00	18:00
Fundamental variance	0,070	0,023	0,030	0,021	0,030	0,032
Pooled difference to previous timepoint		-0,047	0,007	-0,009	0,009	0,001
Nb of contracts with a positive diff		*0*	*2*	*0*	*3*	*0*
Nb of contracts with a negative diff		*3*	*0*	*2*	*0*	*0*
Quote dispersion	0,019	0,010	0,010	0,010	0,011	0,010
Pooled difference to previous timepoint		-0,009	0,000	0,000	0,001	-0,001
Nb of contracts with a positive diff		*0*	*2*	*1*	*2*	*1*
Nb of contracts with a negative diff		*5*	*2*	*1*	*1*	*1*
Mean transaction costs	0,033	0,020	0,019	0,017	0,017	0,019
Pooled difference to previous timepoint		-0,013	-0,002	-0,002	0,000	0,002
Nb of contracts with a positive diff		*0*	*0*	*1*	*1*	*3*
Nb of contracts with a negative diff		*6*	*4*	*5*	*1*	*1*

Panel C. Estimates for varying parameters - European CDSs

	Time					
	6:30	8:00	9:30	12:00	15:00	18:00
Fundamental variance	0,140	0,040	0,037	0,031	0,030	0,025
Pooled difference to previous timepoint		-0,100	-0,002	-0,006	-0,001	-0,005
Nb of contracts with a positive diff		*0*	*1*	*0*	*2*	*0*
Nb of contracts with a negative diff		*4*	*3*	*2*	*2*	*2*
Quote dispersion	0,028	0,019	0,018	0,020	0,018	0,020
Pooled difference to previous timepoint		-0,009	-0,001	0,002	-0,001	0,001
Nb of contracts with a positive diff		*1*	*0*	*5*	*0*	*4*
Nb of contracts with a negative diff		*3*	*5*	*1*	*4*	*1*
Mean transaction costs	0,060	0,029	0,025	0,027	0,024	0,028
Pooled difference to previous timepoint		-0,031	-0,004	0,002	-0,003	0,004
Nb of contracts with a positive diff		*0*	*0*	*5*	*0*	*6*
Nb of contracts with a negative diff		*6*	*6*	*0*	*6*	*0*

Panel D. Estimates for varying parameters - Japanese CDSs

	Time					
	8:00	9:00	11:00	12:30	15:00	18:00
Fundamental variance	0,217	0,020	0,027	0,030	0,022	0,048
Pooled difference to previous timepoint		-0,197	0,007	0,002	-0,008	0,026
Nb of contracts with a positive diff		*0*	*2*	*2*	*0*	*6*
Nb of contracts with a negative diff		*3*	*1*	*1*	*3*	*0*
Quote dispersion	0,036	0,038	0,031	0,033	0,031	0,035
Pooled difference to previous timepoint		0,002	-0,007	0,002	-0,002	0,003
Nb of contracts with a positive diff		*1*	*0*	*4*	*1*	*3*
Nb of contracts with a negative diff		*0*	*5*	*1*	*3*	*0*
Mean transaction costs	0,086	0,080	0,055	0,060	0,049	0,054
Pooled difference to previous timepoint		-0,006	-0,024	0,005	-0,012	0,005
Nb of contracts with a positive diff		*0*	*0*	*4*	*0*	*5*
Nb of contracts with a negative diff		*1*	*6*	*1*	*6*	*0*

Table 4

Correlations between some smoothed quantities and market activity

The table presents the cross-sectional medians of firm-specific correlations based on the model and parameter estimates given in Tables 2 and 3. These correlations are computed for each CDS contract using weekly non-overlapping data. Correlations are defined as follows: mean transaction costs on the ask side are calculated as the weekly mean of the smoothed mean transaction costs on the ask-side. Quote dispersion is the weekly standard deviation of the smoothed mean transaction costs on the ask side. Fundamental variance is the weakly mean of the smoothed squared innovations in the efficient premia. N(ask) is the weekly number of asks. N(trades) is the weekly number of trades.

	Mean transaction costs	Quote dispersion	Fundamental variance	N(ask)	N(trades)
Mean transaction costs	1.00				
Quote dispersion	0.39	1.00			
Fundamental variance	0.32	0.31	1.00		
N(ask)	-0.07	0.27	0.16	1.00	
N(trades)	-0.24	0.03	0.08	0.50	1.00

Table 5

Impact of the subprime crisis on market activity variables and parameter estimates

The table presents the cross-sectional medians of firm-specific correlations based on the model and parameter estimates given in Table 2 and 3. These correlations are computed for each CDS contract using weekly non-overlapping data. defined as follows

Variables	Telecom Italia		Portugal Telecom Intl Finance		Deutsche Telecom		ITV Plc		Valeo		Casino Guich	
Change in the daily number of trades	-1,91	***	-1,16	***	-0,21		-0,67	***	-1,05	***	-0,14	
Change in daily number of quotes	-4,00		-15,41	***	9,67	***	-4,25	*	-2,30		10,74	***
Change in the average fundamental variance	0,010	***	0,003		0,023	***	- 0,036		0,102	***	0,013	*
Change in the average heterogeneity of dealers private costs	0,007	***	0,004	***	0,006	***	0,012	***	0,010	***	0,010	***
Change in mean transaction costs	0,010	***	0,005	***	0,007	***	0,019	***	0,017	***	0,019	***

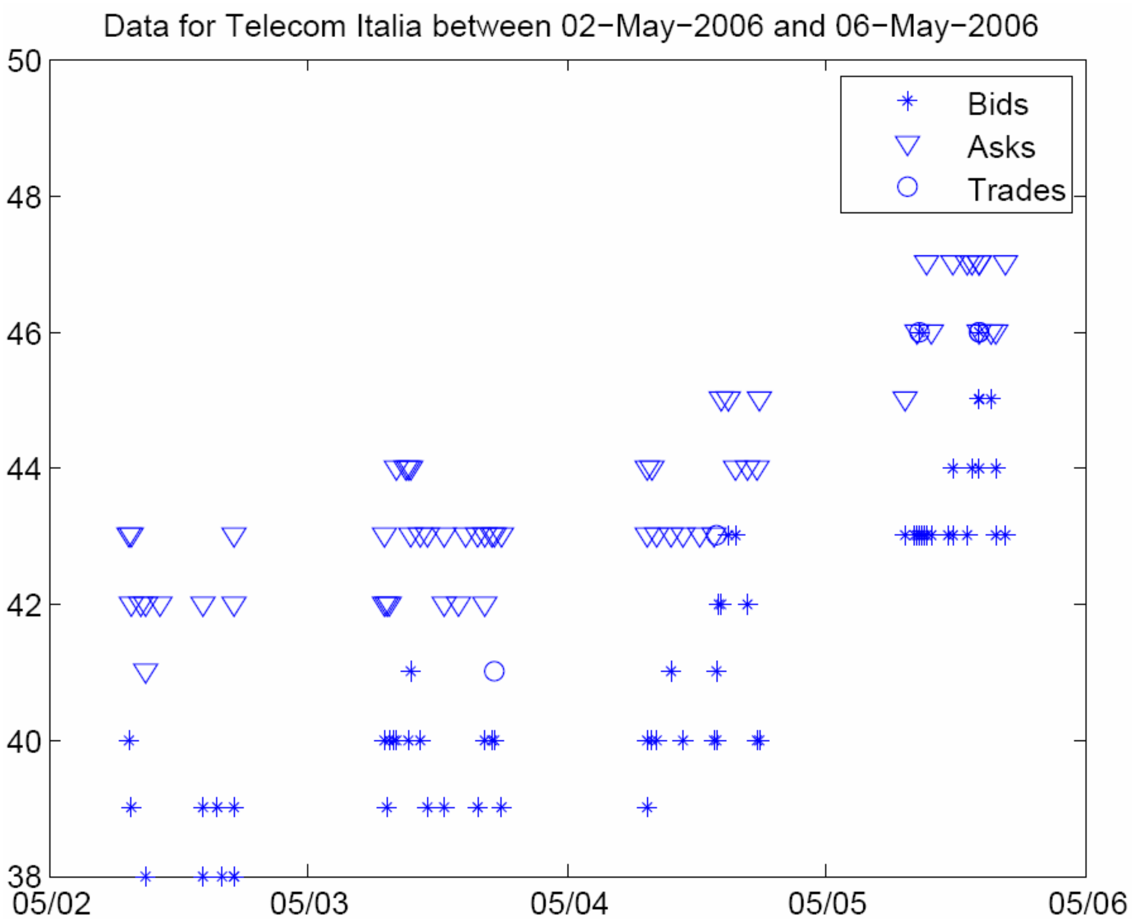

Figure 1: Bid quotes. ask quotes and trade prices expressed in basis point during 4 trading days (May 2. 3. 4 and 5 2006) for the 5-year CDS contract on Telecom Italia

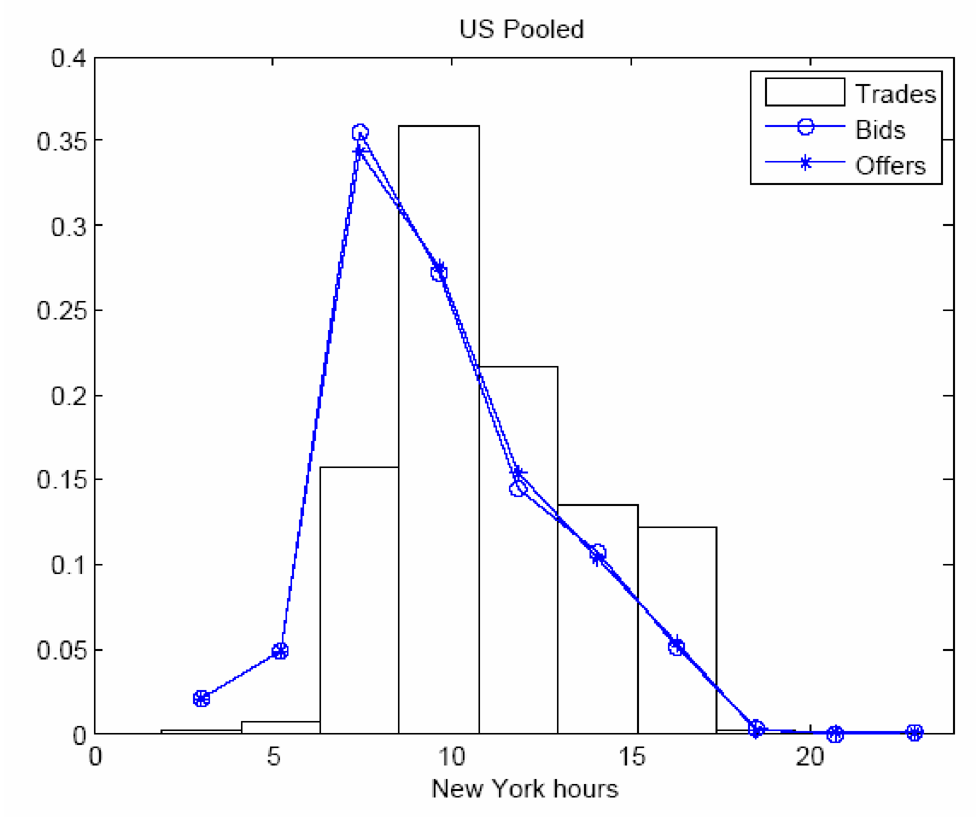

Figure 2a: Pooled frequency of trades and frequency of quotes for the 6 U.S. 5-year CDSs

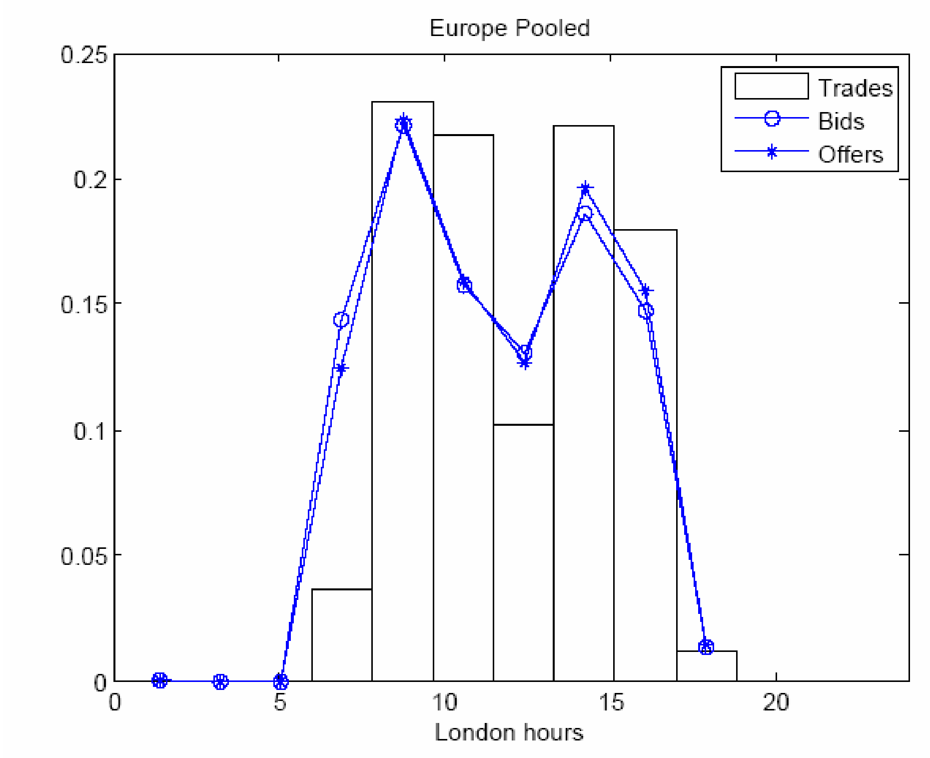

Figure 2b: Pooled intradaily frequency of trades and quotes for the 6 European 5-year CDSs

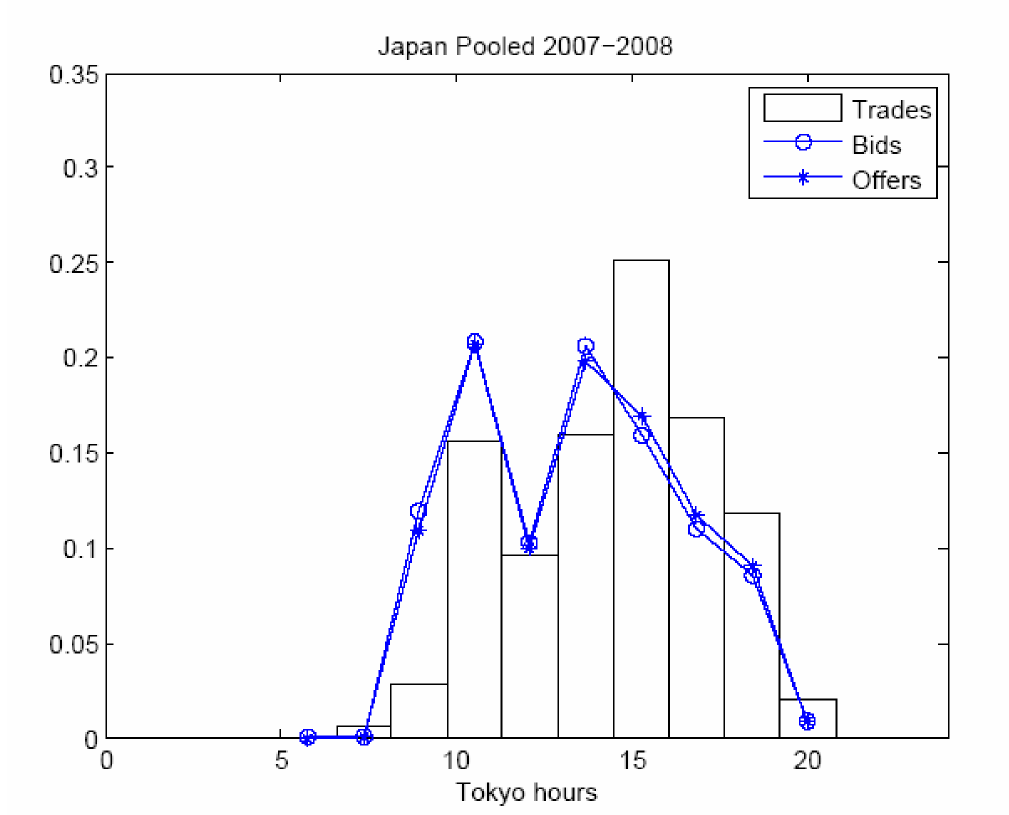

Figure 2c: Pooled intradaily frequency of trades and quotes for the 6 Asian 5-year CDSs

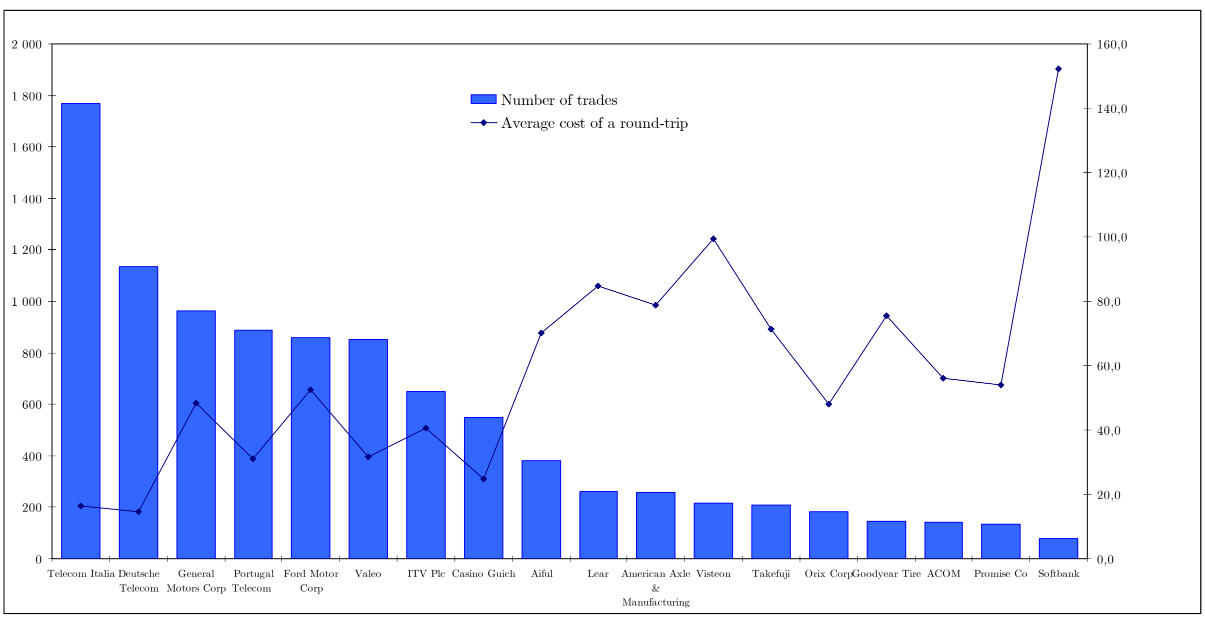

Figure 3: Trading activity and average cost of a round-trip for the CDSs sample

The figure presents the total number of trades and the average cost of a round-trip in basis point for each of the 18 5-year CDS in the sample over the sample period (March 31. 2006 – March 31. 2008)

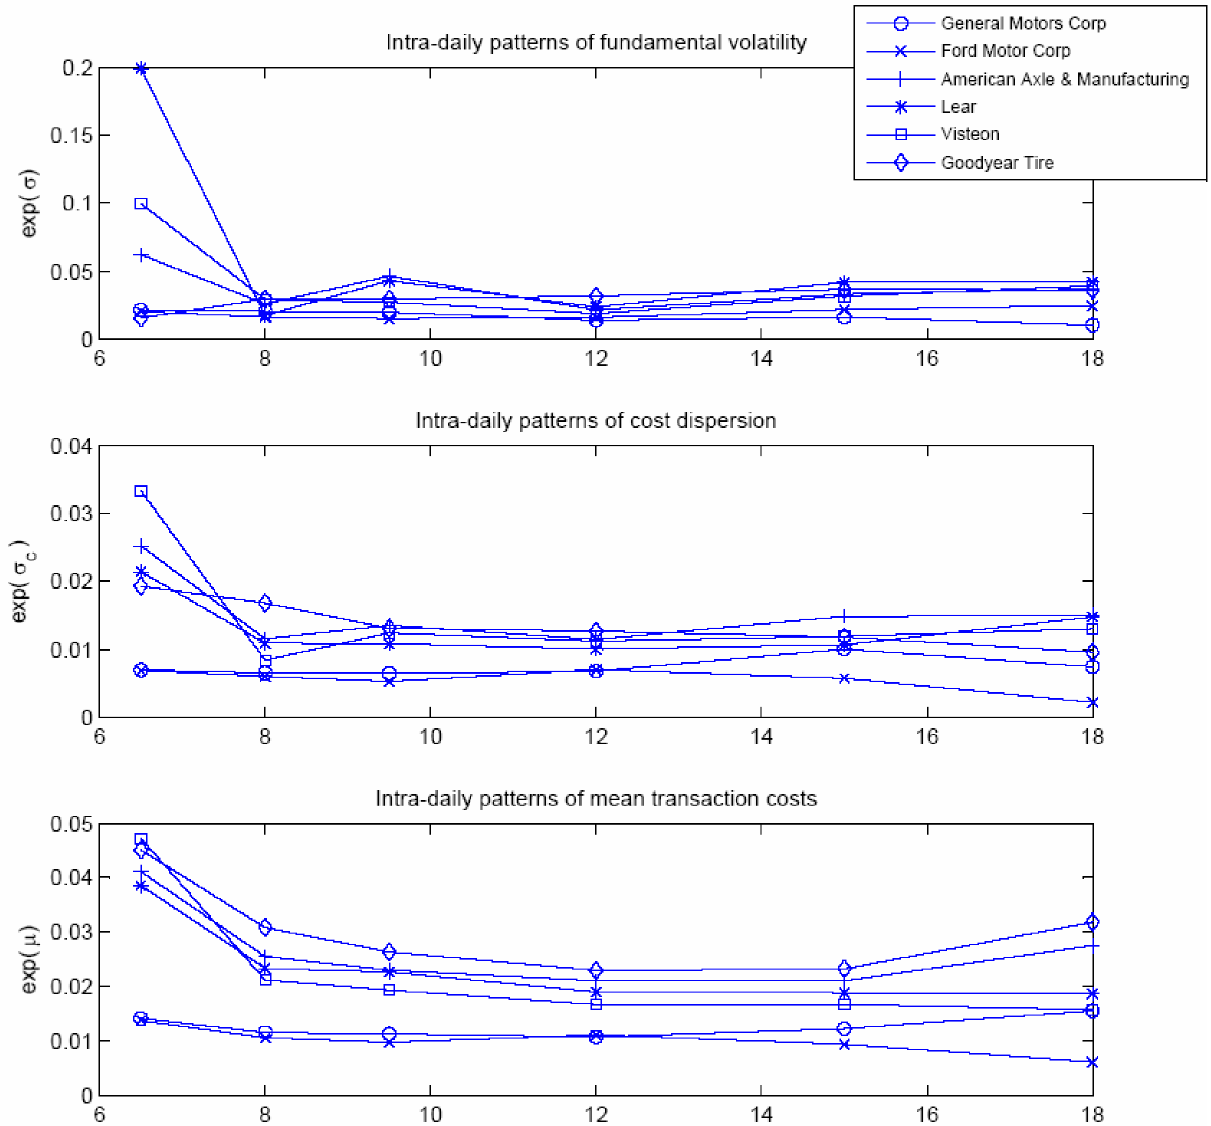

Figure 4: Intra-daily patterns for the 6 U.S. 5-year CDS contracts

From top to bottom. the graphs show the point estimates of the fundamental volatility. the dispersion of dealers' waiting costs and mean transaction costs from 6.30 to 18.00. measured in percentage.

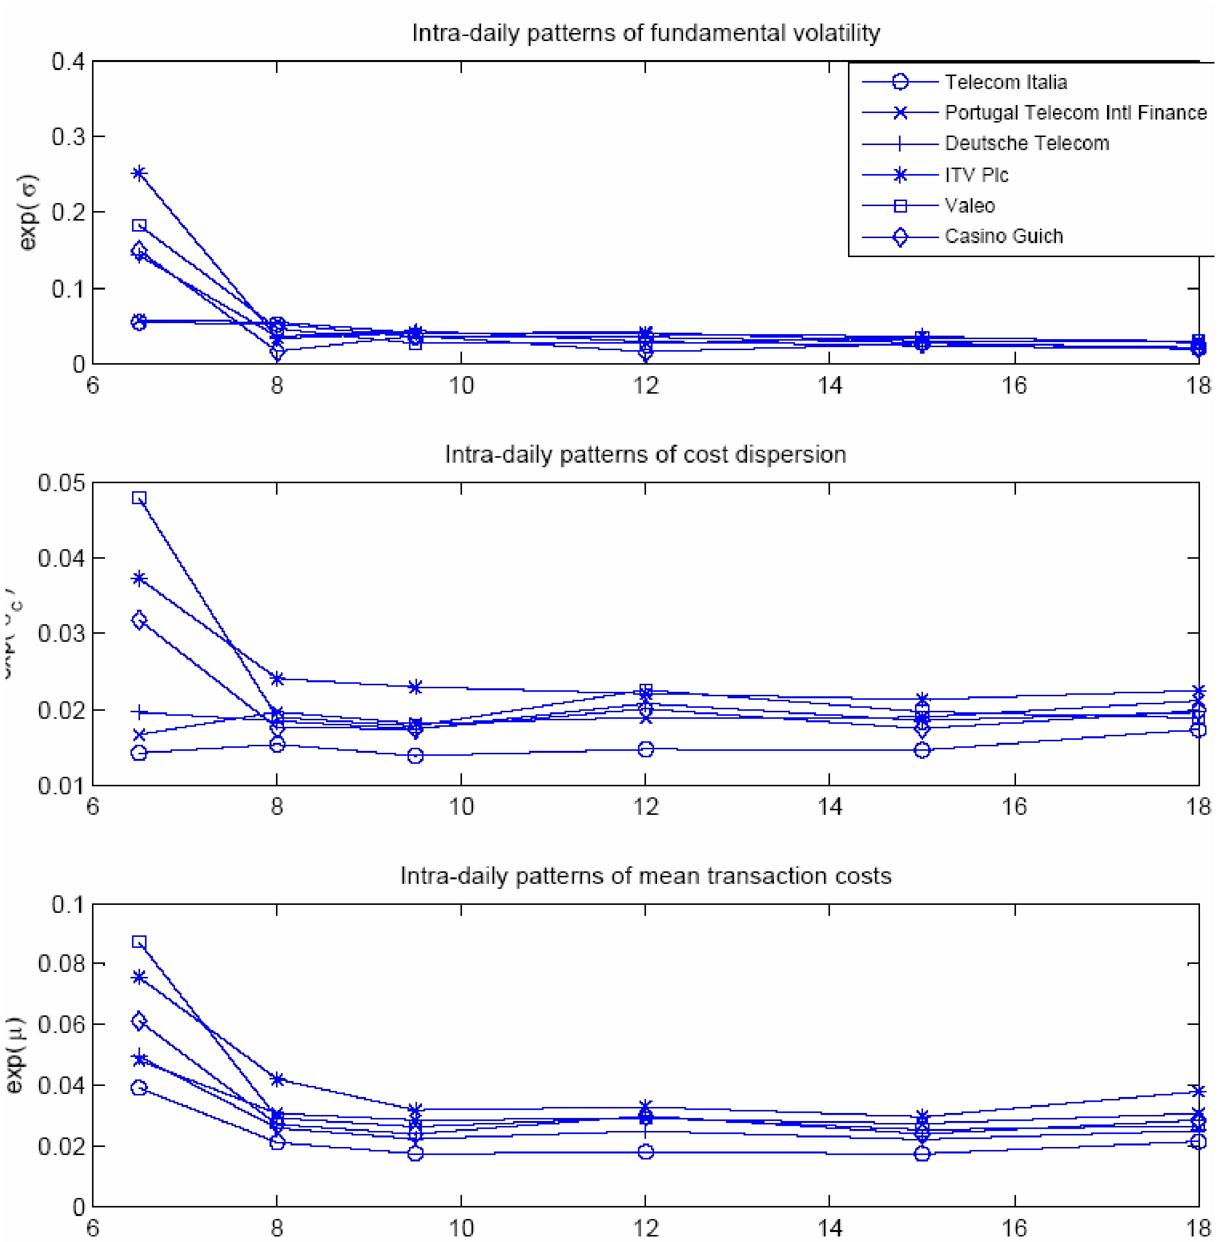

Figure 5: Intra-daily patterns for the 6 European 5-year CDS contracts

From top to bottom. the graphs show the point estimates of the fundamental volatility. the dispersion of dealers' waiting costs and mean transaction costs from 6.30 to 18.00. measured in percentage.

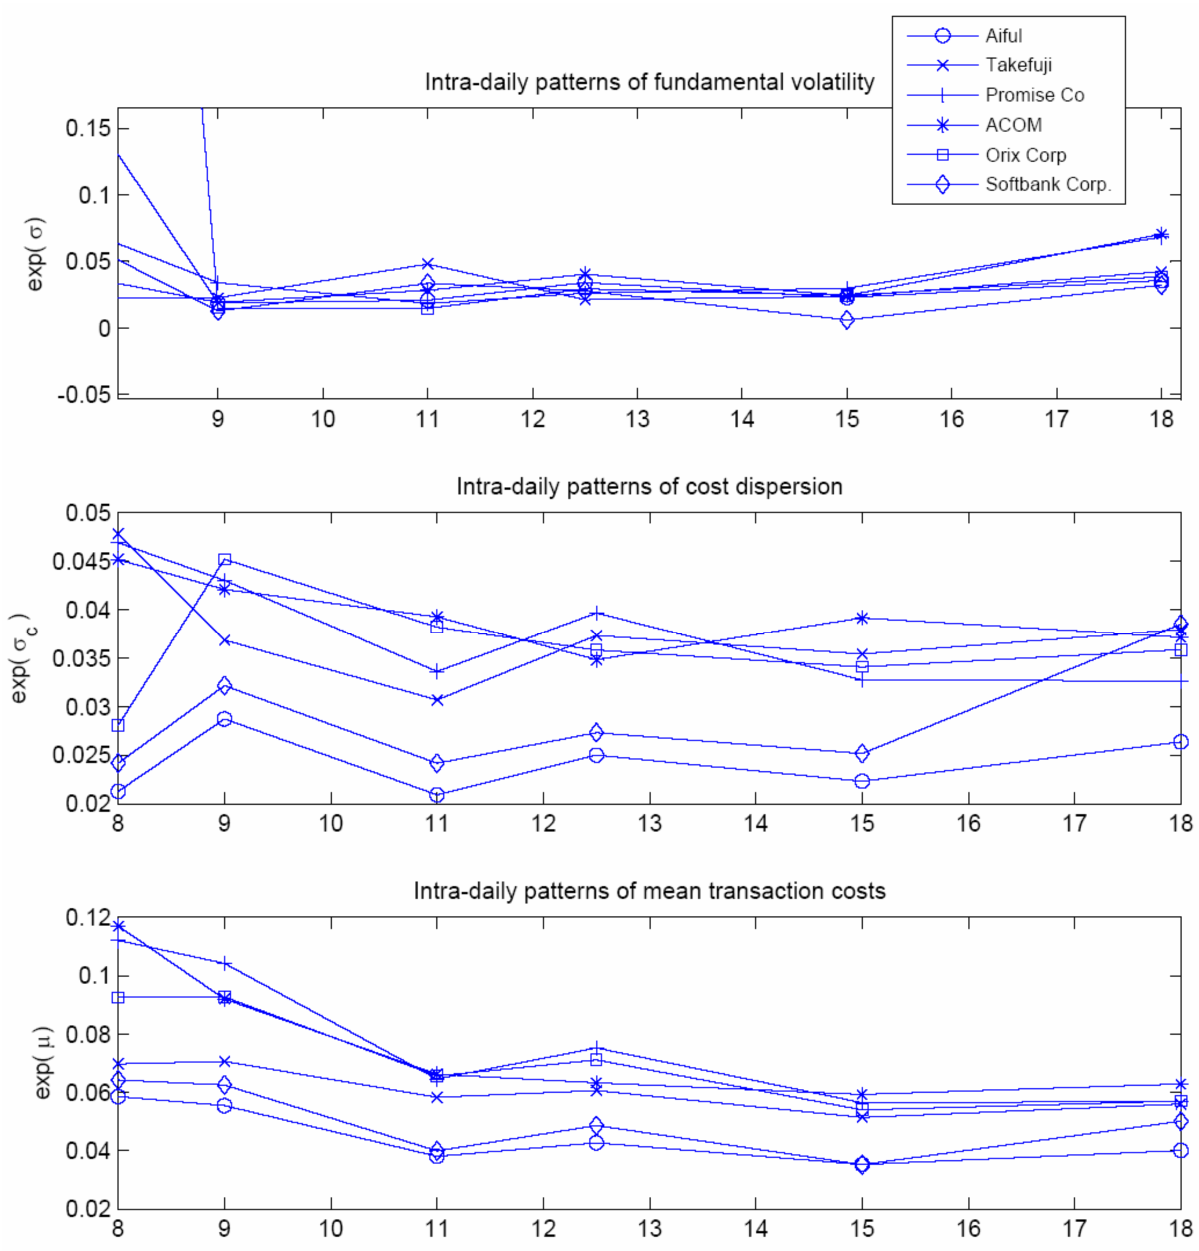

Figure 6: Intra-daily patterns for the 6 Japanese 5-year CDS contracts

From top to bottom. the graphs depict the point estimates of the fundamental volatility. the dispersion of dealers' waiting costs and mean transaction costs from 8.00 to 18.00. measured in percentage.