

How liquid is the CDS market?*

Andras Fulop

ESSEC and CREST

Avenue Bernard Hirsch

95021 Cergy-Pontoise, France.

Email: fulop@essec.fr

Laurence Lescourret

ESSEC and CREST

Avenue Bernard Hirsch

95021 Cergy-Pontoise, France.

Email: lescourret@essec.fr

First Draft: October, 2007

This Draft: December, 2007

Abstract

A common belief is to qualify the credit default swap(CDS) market as very liquid. However, looking at intra-daily CDS data on individual firms from a major inter-dealer broker, we find only limited support for this view. In fact, bid-ask spreads and daily number of trades in our CDS data are more comparable to corporate bond markets than to equity markets. To dig deeper in our data set, we estimate a state-space model of CDS bid and ask quotes on our data. Our model allows for price discreteness, data-errors, heterogeneity of the quotes, jumps in the efficient spreads and intra-daily patterns both in the volatility of the efficient CDS premium and proportional transaction costs. We estimate the model using particle filtering and the Monte Carlo EM algorithm. The volatility of the efficient premium and transaction costs exhibit the usual J-shaped intra-daily pattern observed in equity markets. Also, volatility is much lower during overnight periods and transaction costs much higher.

Keywords : Credit Default Swap, Liquidity, Stochastic Transaction Costs, Inter-dealer Market.

*We thank GFI for providing the data. All errors or omissions are ours.

JEL Classification Numbers:

1 Introduction

Credit Default Swaps (CDS) are arguably one of the most successful financial innovations of the last decade. It is a well-established fact that they provide a more up-to-date picture of creditworthiness than corporate bonds (Blanco, Brennan, and Marsh (2005)) or credit ratings (Hull, Predescu, White (2005)). Also they facilitate taking relatively large long and short positions in the credit markets, improving its efficiency. As a result, it is tempting to go one step further and treat CDS premia as a pure measure of credit risk, by-and-large free of the influence of market microstructure phenomena that bedevil the life of any investigator dealing with the corporate bond markets. In this paper we take a microscopic view of the CDS market using high-frequency data from GFI, a major CDS inter-dealer broker (IDB) and try to check whether the latter view is warranted.

Our data consists of three representative US names (Ford Credit, GMAC, and Sears Acceptance) and one European name (France Telecom). The dataset we analyze contains a complete record of bid and ask quotes, and transaction prices in 2004, 2005, and 2006. We do not observe the volume traded, or the depth posted, or the ID of participants. The access in the IDB is limited to dealers, who trade large amounts. We find that the most active CDS name, Ford, trades on average 4 times a day, while the other three CDSs trade less frequently. Since US corporate bonds trade on average 1.9 per day, and European corporate bond trade on average 4 time a day, the trading frequency of reputed liquid CDSs is quite similar to the corporate bond market.¹ Then we examine the tightness of the market and we find that the cost of turning around a position is around 40 basis points 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.

As a further step in our analysis we build a dynamic model of discrete bid and ask quotes, based on Hasbrouck (2003). The bid and ask quotes are rounding transformation of an implicit efficient premium, and stochastic market-making costs, where the mean proportional costs can be time-varying. The dynamics of the efficient premium allows 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

¹Though the average trade size is much higher in the CDS market

we use the Monte Carlo EM algorithm to estimate the model parameters.²

With this modeling framework we can compare intradaily patterns in the trading activity, volatility of the efficient price, and average transaction costs. We find that the number of quotes per unit of time leads the number of trades per unit of time, which is consistent with either a price discovery process and/or price competition among dealers in the IDB system. The analysis reveals that the volatility has a J-shape during business hours (between 7:30 to 3:30 New York GMT), and is very low during overseas trading. Turning to transaction costs, our results show that the mean costs attain their highest level before the jump in trades, and then declines sharply when the trading frequency picks up, which is consistent with price competition among dealers. Then, transaction costs increase 50% during overseas hours when trading is thin. Overall we find much stronger support in the IDB for price discovery than for the presence of dealers with an inventory motive.

Our results point towards a view that while CDS markets may be much deeper than corporate bond markets, transaction costs are at least as high. Also, looking at intradaily patterns, we find some preliminary evidence for the presence of informed traders. All this suggest that CDS markets are relatively efficient informationally, but in a sense the relatively high transaction costs are the price to be paid for this efficiency. This view is consistent with several recent studies using daily data. Blanco, Brennan, and Marsh (2005) find that the CDS market leads the bond market in determining the price of credit risk, suggesting that informed traders trade first in the CDS market. Acharya and Johnson (2006) find prima facie evidence that informed traders play in the CDS market. Besides, Hull, Predescu, White (2005) provide evidence on the informational efficiency of this market, by showing that CDS premium changes predict rating changes. Our study complements this literature by adding a piece of evidence using high-frequency data.

We also add to the nascent literature on IDB markets. Extant empirical literature (Reiss and Werner (1998), Bjonnes and Rime (2005), Reiss and Werner (2004)) find support for dealers who primarily use the IDB to rebalance their inventories at the end of the day. Interestingly, we don't find much evidence for the inventory motive in our data. In contrast, our findings are consistent with either a price discovery process or competition among dealers.

Last, the econometric methodology we work with could be used to measure transaction costs for illiquid assets. This can be useful for investigating the microstructure of illiquid

²The estimation methodology draws heavily on Duan and Fulop (2007a,b)

stocks or derivatives.

2 Overview of the CDS market

2.1 Main characteristics

Credit derivatives are traded over the counter (OTC), which means that there is no physical location and no central organized exchange where orders are matched. Instead, the CDS secondary market operates 24 hours a day through an electronic network of banks, hedge funds and other institutional investors. Quotes are posted by professional intermediaries, as bank dealers. However, dealers are not required explicit continuous presence. 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). Thus the CDS secondary market is organized as a decentralized dealer market. Currently there is a lack of complete post-trade infrastructure for such OTC credit derivatives. In the EU, there is no central counterparty clearing house. In the US, since 2006, there is at least a central warehouse (the Depository Trust and Clearing Corporation), which has created a giant database for credit derivatives.

Dealers trade directly with clients, and among themselves. Trading is conducted in a bilateral non-anonymous communication over the telephone. However, indirect trading via interdealer brokers (IDBs) plays an increasing role. Interdealer brokers collect and post dealer quotes. They also execute trades between dealers by matching buyers and sellers. Only dealers have access to IDB systems, and they are not required to submit quotes to IDBs. Interdealer brokerage involves voice-based systems, and/or electronic platforms, and represents 34% of trades in 2004, according to the International Swaps and Derivatives Association (ISDA) report.

2.2 Trading CDSs in the interdealer market

Trading organization. Interdealer trading in the CDS market can be direct (through phone trading) or brokered (via open limit order books). There are some major well-known differences between these two organizations: (i) the level of fragmentation of both

systems, (ii) the degree of anonymity, (iii) the level of transparency of both organizations.³ All these key dimensions impact the liquidity of the market. The source of our data are the anonymous centralized electronic IDB market.

The direct interdealer market follows the organization of the secondary OTC market of CDSs, being a fragmented dealer market. Trading is initiated by one dealer calling another for a bilateral quote, i.e., the quote is not available to other dealers (opaque system), which distinguishes it from a quote posted to an IDB. This means that trading in that market potentially involves search costs for a matching counterparty.

On the other hand, while IDBs are not costless (there is the interdealer broker's fee to pay), they provide platforms which broadcast dealer quotes on screens, thus decreasing search costs. Precisely, IDB systems are centralized books for brokers that can better match buyers and sellers with a larger vision of the market and of the order flow. Brokers collect dealers' quote and depth information, and post the data into the system that are provided anonymously to the other dealers. The organization of these electronic platforms are more transparent, more centralized, and should also enhance the speed of execution.

Since trades are executed anonymously in IDBs market, they enable dealers to trade in size. Moreover, the size of limit orders is undisclosed in some IDBs.⁴ Anonymity of quotes and trades is known to lower volatility and enhance liquidity (see Foucault, Moinas, and Theissen, 2007)).

Trading motives. Why do dealers usually trade in anonymous IDBs? As Huang, Cai and Wang (2002) point out, trading in IDBs may be visualized as the second step of a two-stage game. The first stage concerns trades between dealers and clients in the opaque, fragmented quote-driven market. Then, dealers trade in the interdealer market either to manage their inventory positions (i.e. to unwind or to build a position) as suggested by Ho and Stoll (1983), or/and to trade on private information. Given the anonymity of IDBs platforms, the level of information-based trading could be high in these facilities: IDBs could be the favored trading venue of informed interdealer trades (see, for instance, Barclay and Hendershott (2004)).

³Using daily data from Bloomberg and CreditTrade, Gndz, Ldecke, and Uhrig-Homburg (2006) analyze difference in the quoted b/a spread for 200 CDSs between the IDB, CreditTrade, and the direct OTC market. They find that the quoted spread is higher in IDBs.

⁴The anonymity of IDB systems is a key factor of these platforms. Dealers choose to trade anonymously because their identity is an additional information that could be exploited by competitors.

However, the few empirical studies that exist on IDBs seem to find more evidence for the inventory motive.⁵ First, in the London Stock Exchange, Reiss and Werner (1998) show that interdealer markets are mainly used by equity dealers to share inventory risks. Bjonnes and Rime (2005) also find evidence that, in the FX market, IDBs are used by dealers to manage 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. Reiss and Werner observe that when adverse selection is perceived to be high, limit order are cancelled in IDB platforms, and liquidity dries up. Besides when a dealer who possesses a time-sensitive private information posts an informed limit order in the IDB it may not be executed, and information may leak out.

In light of these studies, a priori, one expects to observe intraday patterns in IDBs for CDSs that are consistent with the inventory paradigm and less affected by information asymmetries. 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 indeed not of major importance on the CDS IDB market, one should not 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

We use intra-daily bid and ask quotes from GFI, a major CDS interdealer broker. GFI is a hybrid voice-electronic execution platforms for CDSs. Dealer may be providers or consumers of liquidity. Retail and institutional investors (the "buy-side") are not eligible to use the platform directly. The platform has a minimum trade size of \$1 million. Only the broker can see the depth of the market; however, dealers can observe the last trade (price, volume, and direction). Bid and ask prices posted on the platform are firm. According to GFI, it is the leading broker since it would represent 60% of the interdealer brokerage

⁵The most well-known IDBs of dealer markets are GovPX for the US Treasury bond market, CanPX for the Canadian bon market (see Boni and Leach (2004), Huang et al (2002), D'Souza et al (2003)), EBS and Reuters 2000-2 for the FX market, and, also, SelectNet and SuperSOES for the Nasdaq, or ??? for the LSE.

activity. As a broker, GFI earns commissions directly from trade counterparts.

The dataset that GFI provided us contains CDS quote and trade price for 4 entities (3 from US and 1 from Europe) whose CDS contracts are reputed very liquid. The 4 names are Ford, GMAC, Sears Acceptance and France Telecom, all of them are on a 5-year contracts on senior debt. Ford and GMAC were downgraded from investment to speculative bond on May, 5 2005. France Telecom is graded as investment (see Table 1, Panel A).

Our sample runs from January 2004, through December 2006. The sample consists of bid, ask quotes, and trade prices expressed in basis point. Our data is time stamped down to a minute. Our local time benchmark is New York GMT. There are no identifiers of the dealers who post quotes or who trade. No information about the depth of quotes, nor information about the size of each transaction is provided.⁶

We examine two measures of trading activity: the number of bid or/and ask quotes per unit of time (the quoting frequency), and the number of trades per unit of time (the trade frequency).

As Figure 1 illustrates, quotes and trades are irregularly spaced . Although trading could take place round-the-clock, the trading activity (number of quotes and trades) of the 3 US entities mainly occur during the New York business hours (7:30 am to 5:30 pm), as Figures 2, 3, and 4 illustrate.⁷ However, Ford and GMAC exhibit some weak but significant trading activity when London open at 7:30 am (2:30 am New York), unlike Sears Acceptance. In contrast, the main peak for the trading activity is in London for the European name. Trading of the European CDS France Telecom starts at 7:30 am London. Then it passes to New York at 12:30 am London (7:30 am New York), and continues there until 5:30 pm (see Figure 5).

It is worth noticing, first, that, during New York business hours, we observe a J-shape pattern for the number of trades for the 3 US names. The number of trades does not spike upward at the end of the day, as it is usually the case in other markets (U-shaped pattern for volume). This intraday pattern is quite unexpected, since it seems to go against the hypothesis of IDBs as inventory management tools. Indeed, in case IDBs would have been

⁶From interviews with traders, it seems that the minimum trade size is 5-10 million in Europe and in US: CDS contracts are only traded in relatively large slot ("granularity" of CDS contracts). For instance, a trade on a 20 – 30 million for the most liquid European CDS is quite standard. Note also that the most common practice is to trade on a 5-year contract.

⁷The reasons may be that news concerning these US firms are probably announced during New York trading hours, and holders of these CDS contracts are presumably US banks or hedge funds.

used by dealers to close or hedge open inventory position before the overnight period, we should observe a jump upward at the end of the business day.

Second, we find that for the 4 names, the peak of the number of trades follows the busiest period in terms of number of quotes (lead-lag pattern between quotes and trades). This pattern may reflect either the price competition among dealers in the IDB platform to fill the book before trading happens, or/and a price discovery process similar to what we observe during the preopening periods in regulated equity markets.⁸

3.2 Descriptive statistics

Table 1 (Panel B) reports some descriptive statistics on the trading activity. For the 4 CDS contracts reputed to be very liquid, we compute the average number of trades per day. France Telecom trades on average 1.42 a day. Concerning the 3 US entities, GMAC and Ford trade on average 4.46 a day, and 2.57 a day respectively. Sears Acceptance trades less (0.78 a day). The trading 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 that covers also less active bonds. Biais and Declerck (2007) find that Euro-denominated corporate bonds trade on average 4 times per day.⁹

Table 1 (Panel C) reports also some descriptive statistics on quotes. First, even for these reputed liquid CDS contracts, the number of quotes (bid or ask) is quite low, compared to the equity market. For instance, GMAC and Ford are active CDSs with more than 1400 bid/ask pairs during three years, whereas we observe only 568 bid/ask pairs for Sears Acceptance. Thus, it is important to use as much information in the data as possible, in particular, to use the individual bid or ask observations.¹⁰

Second, the level of CDS spreads seems to be related to the rating of the entity. Average quoted b/a spreads are around 10 bp for CDS of speculative-grade entities (the 3 US names), and lower than 3 bp for the European CDS of the investment-grade bond France

⁸Equity markets have regulated trading hours. There is a clear distinction between the overnight hours, and the trading hours, unlike OTC 24-hours markets

⁹Even if the trading frequency is similar between the bond market and the CDS market, overall trading should be much higher in the CDS as the size of individual trade is much higher.

¹⁰In contrast, Huang et al. (2002) and others filter out one-sided quotes when using GovPX.

Telecom.

To measure the cost of a round-trip in the CDS market, one cannot simply use the quoted b/a 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 b/a spread. So, 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, so one needs to compute a risky annuity to get the PV factor. For instance, the cost of a round trip for an investment in GMAC with a notional of \$100 is 37 cents. Analyzing a sample of Euro-denominated corporate bonds, Biais and Declerck (2007) find an average quoted spread around 30 cents. Edwards, Harris and Piwowar find that the effective cost is 20 cents for large trade on US corporate bonds.

These numbers suggest that b/a spreads and transaction costs on the CDS market are not lower than on the corporate bond market. Thus, microstructure phenomena (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 Econometric Model

To be able to dig further in our data set we link the quotes to an unobserved efficient price and cast the resulting estimation task as a missing-data problem. In our data, in any minute, we keep the last recorded bids and asks and label an observation a joint observation if they have the same time stamp. We discard all joint observations where the ask is smaller than or equal to the bid. Then, filtering and estimation in the resulting non-linear state-space system is solved using simulation-based methods. This modeling framework allows us to estimate intra-daily patterns of the efficient CDS premium and intra-daily patterns of transaction costs. Also, as a by-product we can estimate the salient features of the efficient premium evolution and come up with a filtered estimate of the CDS premium at any time point.

We consider a model conditional on the arrival of bids and asks of CDS spreads. Let $\tau_i, i = 0, \dots, T$ 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 ask observation at τ_i denoted by B_{τ_i}, A_{τ_i} respectively,

or we may have both. Similarly to Hasbrouck (2003) we assume a latent variable model for our data. In particular, we link all observables to a theoretical efficient log spread, $m_{\tau_i} (= \log(M_{\tau_i}))$, which follows

$$m_{\tau_i} = \mu_m \Delta\tau_i + (1 - \rho \Delta\tau_i) m_{\tau_{i-1}} + \sigma_{f(\tau_{i-1}, \tau_i)} \sqrt{\Delta\tau_i} (u_i + N_i Z_i) \quad (1)$$

where $\Delta\tau_i = \tau_i - \tau_{i-1}$. Here, $\sigma_{f(\tau_{i-1}, \tau_i)}$ denotes the average intradaily speed of the information arrival to the market between τ_i and τ_{i-1} . u_i is standard normal and the remaining term in the innovation is there to allow for fat-tails and possibly skewness. N_i is a Bernoulli with parameter λ and Z_i is an independent normally distributed random variable with mean μ_J and variance σ_J^2 . μ_m defines a trend in the spread evolution and ρ the speed of mean reversion.

Dealers are assumed to be subject to nonnegative costs of market making, denoted $c_{i,A}$ and $c_{i,B}$, as follows.

$$\begin{aligned} c_{i,A} &\sim N(\mu_{\tau_i}, \sigma_{c,A}^2) \text{ constrained to } c_{i,A} > 0 \\ c_{i,B} &\sim N(\mu_{\tau_i}, \sigma_{c,B}^2) \text{ constrained to } c_{i,B} > 0 \end{aligned}$$

We assume that in the absence of tick restrictions a market maker would bid $B_{\tau_i} = M_{\tau_i} - C_{i,B} = e^{m_{\tau_i} - (c_{i,B} + q_i^A \varepsilon_i^A)}$ and offer(ask) $A_{\tau_i} = M_{\tau_i} + C_{i,A} = e^{m_{\tau_i} + c_{i,A} + q_i^B \varepsilon_i^B}$. The expressions $q_i^A \varepsilon_i^A$ and $q_i^B \varepsilon_i^B$ represent data errors where q_i^A, q_i^B are iid Bernoullis that take the value 1 with probability p and ε_i^A and ε_i^B are iid normal with mean zero and variance σ_ε^2 .

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

$$B_{\tau_i} = \text{Floor} \left(e^{m_{\tau_i} - (c_{i,B} + q_i^B \varepsilon_i^B)}, K \right) \quad (2)$$

$$A_{\tau_i} = \text{Ceiling} \left(e^{m_{\tau_i} + c_{i,A} + q_i^A \varepsilon_i^A}, K \right) \quad (3)$$

where the floor and ceiling functions round down and up respectively to the next multiple of K .¹¹

¹¹Duan and Fulop (2007b) applies a similar model to high-frequency stock market transaction data. However, their focus is on the efficient price process not microstructure phenomena. Accordingly, their efficient-price model is richer, but their microstructure model is simpler.

5 Filtering and ML estimation

5.1 Particle Filter for the model

First, we describe the filtering algorithm for the case when we only have an ask observation. Assume that the time index of the new observation is τ_i . Denote the noisy ask observation before discretization by $a_{\tau_i} = m_{\tau_i} + c_{i,A} + q_i^A \varepsilon_i^A$. Our algorithm is based on the following filtering density/distribution

$$\begin{aligned}
& f(a_{\tau_i}, m_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A, c_i^A \mid \mathcal{D}_i, c_i^A > 0) \\
\propto & f(c_i^A \mid a_{\tau_i}, m_{\tau_i}, q_i^A, c_i^A > 0) p(c_i^A > 0 \mid a_{\tau_i}, m_{\tau_i}, q_i^A) f(m_{\tau_i} \mid a_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A) \\
\times & f(a_{\tau_i} \mid A_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A) f(m_{\tau_{i-1}}, N_i, q_i^A \mid \mathcal{D}_i) \\
\propto & f(c_i^A \mid a_{\tau_i}, m_{\tau_i}, q_i^A, c_i^A > 0) p(c_i^A > 0 \mid a_{\tau_i}, m_{\tau_i}, q_i^A) f(m_{\tau_i} \mid a_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A) f(a_{\tau_i} \mid A_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A) \\
\times & p(A_{\tau_i} \mid m_{\tau_{i-1}}, N_i, q_i^A) p(N_i) p(q_i^A) f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1})
\end{aligned} \tag{4}$$

The last expression in (4) suggests a way to sample from the filtering distribution, given a sample of particles representing $f(m_{\tau_{i-1}} \mid \mathcal{D}_{i-1})$. First, extend the state-space to include jumps in the system by simulating from some importance sampler $g_1(N_i)g_2(q_i^A)$. Then, attach weights $w_1 = \frac{p(A_{\tau_i} \mid m_{\tau_{i-1}}, \Delta N_i, q_i^A) p(N_i) p(q_i^A)}{g_1(N_i) g_2(q_i^A)}$. Then, sample (m_{τ_i}, a_{τ_i}) according to $f(m_{\tau_i} \mid a_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A) f(a_{\tau_i} \mid A_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A)$ and compute the weights $w_2 = p(c_i^A > 0 \mid a_{\tau_i}, m_{\tau_i}, q_i^A)$. Last, resample proportional to the weights $w = w_1 w_2$.

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

- **Step 1:** Enlarge the state-space by the jumps in the system by sampling from N_i and q_i^A . Jumps can be rare events so we use stratified sampling to ensure that our simulated sample always contains jumps by putting an equal probability of $\frac{1}{2}$ on having and not having jumps. I.e. we use the importance sampling distributions

$$g_1(N_i) = \frac{1}{2} 1_{\{N_i=1\}} + \frac{1}{2} 1_{\{N_i=0\}} \tag{5}$$

$$g_2(q_i^A) = \frac{1}{2} 1_{\{q_i^A=0\}} + \frac{1}{2} 1_{\{q_i^A=1\}} \tag{6}$$

To arrive at an empirical representation of $f(m_{\tau_{i-1}}, N_i, q_i^A \mid \mathcal{D}_i)$ attach to each of the particles $(m_{\tau_{i-1}}^{(m)}, N_i^{(m)}, (q_i^A)^{(m)})$ the importance weights

$$w_{i,1}^{(m)} = p(A_{\tau_i} \mid m_{\tau_{i-1}}^{(m)}, N_i^{(m)}, (q_i^A)^{(m)}) p((q_i^A)^{(m)}) [1_{N_i^{(m)}=1} p(N_i = 1) + 1_{N_i^{(m)}=0} p(N_i = 0)]$$

- **Step 2:** Corresponding to each particle $(m_{\tau_{i-1}|i}^{(m)}, N_i^{(m)}, (q_i^A)^{(m)})$, sample from the truncated normal density $f(a_{\tau_i} | A_{\tau_i}, m_{\tau_{i-1}|i}^{(m)}, N_i^{(m)}, (q_i^A)^{(m)})$ to generate the particle $(a_{\tau_i}^{(m)}, m_{\tau_{i-1}|i}^{(m)}, N_i^{(m)}, (q_i^A)^{(m)})$. The corresponding particle set represents $f(a_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A, | \mathcal{D}_i)$.
- **Step 3:** Using conditional normality sample from

$$m_{\tau_i}^{(m)} \sim f(m_{\tau_i} | a_{\tau_i}^{(m)}, m_{\tau_{i-1}|i}^{(m)}, N_i^{(m)}, (q_i^A)^{(m)})$$

Attach weights to the particles as

$$w_{i,2}^{(m)} = p(c_i^A > 0 | a_{\tau_i}^{(m)}, m_{\tau_i}^{(m)}, (q_i^A)^{(m)})$$

Now sample from the truncated normal distribution

$$(c_i^A)^{(m)} \sim f(c_i^A | c_i^A > 0, a_{\tau_i}^{(m)}, m_{\tau_i}^{(m)}, (q_i^A)^{(m)})$$

- **Step 4:** 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 denoted by $(m_{\tau_{i-1}|i}^{(m)}, N_i^{(m)}, (q_i^A)^{(m)})$. This equal-weighted particle set is again an empirical representation of $f(m_{\tau_{i-1}}, N_i, q_i^A, | \mathcal{D}_i)$.

One can then proceed to marginalize m_{τ_i} to represent the filtering distribution $f(m_{\tau_i} | \mathcal{D}_i)$

The likelihood value can be computed as

$$\begin{aligned} P(A_{\tau_i} | \mathcal{D}_{i-1}, c_i^A > 0) &= \frac{P(c_i^A > 0 | \mathcal{D}_{i-1}, A_{\tau_i})P(A_{\tau_i} | \mathcal{D}_{i-1})}{P(c_i^A > 0)} \\ &\approx \frac{\frac{1}{M} \sum_{m=1}^M w_i^{(m)}}{P(c_i^A > 0)} \end{aligned}$$

The necessary quantities to implement the algorithm are described in detail in Appendix A. 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.

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 t_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 $I_i = (m_{\tau_0} + m_{\tau_1} + \dots + m_{\tau_i}) / (i + 1)$. 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 has nothing to do with I_i , and thus the algorithm remains unchanged. However, 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_{τ_i}) , i.e., $f(I_i, m_{\tau_i}, | \mathcal{D}_i)$.

5.2 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. Thus we 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 estimate for the incomplete data model, where incomplete data refers to the situation that the model contains some random variable(s) without corresponding observations. The EM algorithm involves two steps - expectation and maximization – and hence its name. One first writes 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. 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. Interestingly, the EM algorithm will converge to the ML estimate under some regularity conditions.

For our ML estimation, the E-step due to its complexity will have to be computed using the particle filter, which means that we are using the Monte Carlo EM (MCEM) algorithm.¹² Casting optimization as an EM algorithm problem effectively circumvents the irregularity induced by the particle filter, because the E-step ensures that the expected complete-data log-likelihood function is smooth with respect to the model parameters that define the complete-data log-likelihood function. Even though the function is still inherently irregular in relation to the assumed parameter values used in computing the expectation, it becomes immaterial as far as optimization is concerned. In effect, one has decoupled optimization from filtering in each iteration. In general, the complete-data representation of the model is not unique. We choose a specification that allows for relatively easy M-steps and where the expected complete-data loglikelihood can be written as a combination of sufficient statistics of the latent

¹²For a general introduction to the MCEM algorithm, see for instance Wei and Tanner (1990).

data and the model parameters. This latter property is important because it means that we need to run the particle filter only once per iteration of the 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 , the jump size $Z_{i,J}$ and of the log efficient CDS premia, m_{τ_i} . The joint conditional likelihood can be written as

$$f(m_{\tau_i}, Z_{i,J}, N_i | m_{\tau_{i-1}}) = f(m_{\tau_i} | Z_{i,J}, N_i, m_{\tau_{i-1}})f(Z_{i,J} | N_i)f(N_i)$$

The loglikelihood of the number of jumps is

$$L_i^N = 1_{\{N>0\}} \ln(\lambda) + 1_{\{N=0\}} \ln(1 - \lambda) \quad (7)$$

Conditional on the number of jumps, the loglikelihood of the innovations due to the jumps is

$$L_i^{ZJ} = 1_{\{N>0\}} \times \left[-\frac{(Z_{i,J} - \mu_J N_i)^2}{2\sigma_J^2(N_i + 1_{\{N=0\}})} - \frac{\ln(\sigma_J^2)}{2} \right] \quad (8)$$

Then, we can write the loglikelihood of the efficient spread innovation as

$$L_i^m = -\frac{[m_{\tau_i} - (\mu_m \Delta\tau_i + (1 - \rho \Delta\tau_i)m_{\tau_{i-1}} + \sqrt{\Delta\tau_i} \sigma_{f(\tau_{i-1}, \tau_i)} Z_{i,J})]^2}{2\sigma_{u, f(\tau_{i-1}, \tau_i)}^2 \Delta\tau_i} - \frac{\ln(\sigma_{u, f(\tau_{i-1}, \tau_i)}^2)}{2} \quad (9)$$

We also need the complete data representation of the microstructure noises and data errors. Here for simplicity we present formulas when we only have an ask observation.¹³ We include the microstructure errors $c_{i,A}$, the data error indicators $q_{i,A}$ and the data errors $\varepsilon_{i,A}$. Here, some more care is needed to deal with the conditioning on the event $\{c_{i,A} > 0\}$. To obtain a simple complete data loglikelihood, we also include in the complete data space draws of $c_{i,A}$ that were thrown away because of not satisfying the condition. Let k_i denote the number of such draws. Then the complete data loglikelihood is

$$\begin{aligned} L_i^{MI,A} &= -\frac{(c_{i,A} - \mu_{\tau_i})^2}{2\sigma_{c,A}^2} - \frac{\ln(\sigma_{c,A}^2)}{2} + q_{i,A} \left[-\frac{\varepsilon_{i,A}^2}{2\sigma_\varepsilon^2} - \frac{\ln(\sigma_\varepsilon^2)}{2} \right] \\ &+ \sum_{j=0}^{k_i} \left(-\frac{(c_{i,j,A} - \mu_{\tau_i})^2}{2\sigma_{c,A}^2} - \frac{\ln(\sigma_{c,A}^2)}{2} \right) \end{aligned}$$

Then 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.

¹³The cases with only a bid or both observations are almost identical, so to conserve space we do not report them here

- **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 approximate quantities of the form

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

where X_i is some function of the hidden random states, \mathcal{D}_T is the observed data and $\theta^{(k-1)}$ is the parameter vector at the $(k-1)^{th}$ iteration. The particle filter described in section 5.1 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 $X_{i|T}^{(m)}$. Thus, the expectation can be approximated by the sample average as follows:

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

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. Cappe and Moulines (2005) 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. Cappe and Moulines (2005) 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:

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

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

$$L_i(\theta) = L_i^N(\theta) + L_i^{Z_J}(\theta) + L_i^m(\theta) + 1_{i \in I_A} L_i^{MI,A}(\theta) + 1_{i \in I_B} L_i^{MI,B}(\theta) \quad (10)$$

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

$$\theta^{(k)} = \arg \max_{\theta} \sum_{i=1}^T E(L_i(\theta) \mid \mathcal{D}_T, \theta^{(k-1)})$$

A detailed description of the M-step is in Appendix B.

The usual way to compute the asymptotic standard error for the maximum likelihood estimate is to use the negative Hessian matrix or the inner product of the individual scores. But in our case, either one is not directly computable because the individual observed log-likelihood function is highly irregular with respect to θ due to using the particle filter. An alternative estimator proposed by Duan and Fulop (2007a) can be applied to our setting, however, which uses the smoothed individual scores to compute the asymptotic error.

6 Empirical Results

In this section we investigate intraday patterns of volatility and transaction costs. We allow both the diffusion volatility and the proportional costs to vary during the day. For the 3 US entities, the parameter values are estimated during 5 time periods in New-York GMT: 5:30-7:30 (1), 7:30-9:30 (2), 9:30-14:30 (3), 14:30-16:30 (4) and the overnight period 16:30-05:30 (5). For the European name, the parameter values are estimated during 6 time periods (NY GMT): 5:30-7:30 (1), 7:30-9:30 (2), 9:30-14:30 (3), 14:30-16:30 (4) and the overnight period 16:30-05:30 (5). Table 2 shows the parameter estimates and the asymptotic standard errors.

For the 3 US entities, the volatility parameter exhibits the usual J-shaped pattern during the business hours (5:30 am-5:30 pm) (see Hasbrouck (1999) and others). The volatility is at the lowest level during off hours trading, when almost no trade or quote happens. This pattern may signal the price discovery process and/or price competition among traders. However, it is not likely that portfolio re-balancing is the main motive of the dealers.

The intraday pattern of the mean cost of the US entities, μ , is noteworthy. The mean cost attains its highest level prior to the New York business hours (5:30-7:30), when there are few trades, but when the quoting activity is high in the IDB. Then μ

declines sharply from 7:30 to 9:30, when the number of trades picks up. During the late afternoon, μ increases slightly for Ford and GMAC, and it picks up (increase of 50%) during the off-hours period (17:30 - 5:30), when trading is thin. Note that we do not have enough data to estimate the mean costs of Sears during the overnight period. Thus, we observe that the number of trades and the mean costs are inversely related in our data, which is consistent with the theory of liquidity externalities.

7 Conclusion

In this paper, we use a new high-frequency data set for 4 credit default swaps, provided by GFI, one of the main interdealer brokers.

First, we find that bid-ask spreads and roundtrip costs in our data are not substantially lower than their counterparts in the corporate bond markets. This suggests that just like corporate bonds, CDS prices are not likely to be free from microstructure phenomena. So one should be careful when interpreting CDS premia as a pure measure of credit risk.

Second, as we have intra-day observations, we are able to estimate an econometric model allowing for intradaily patterns in transaction costs and volatility. In the model, bid and ask quotes arise from an implicit efficient price and stochastic market-making costs. Further, our 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. The parameter estimates show a J-shape pattern for the volatility of the efficient price, and for transaction costs during New York business hours for the 3 US names. Patterns are different for the European name whose trading activity and volatility pick up during London business hours. For all the four names, we find that volatility is low and transaction costs are higher when trading is thinner (off hours trading).

These results suggest that this IDB does not seem to be used so much by dealers to manage their inventory position. We do not observe in our data any patterns usually consistent with the inventory paradigm, (for instance, we don't find a peak in the number of trades, or volatility at the end of the day). In contrast, the J-shape pattern of volatility and transaction costs may signal the presence of the price discovery process. If this is indeed the case, we have found an example of an IDB

where anonymity promotes informed trades. However, further research is needed on this issue, especially to disentangle the asymmetric information story from price competition among dealers.

Appendix A:

The necessary quantities for executing the particle filtering algorithm when the observation is an ask are described below. The conditional distribution of $f(a_{\tau_i} | m_{\tau_{i-1}}, \Delta N_i, q_i^A)$ is normal with mean and variance

$$\begin{aligned} E[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A] &= \mu_m \Delta \tau_i + (1 - \rho \Delta \tau_i) m_{\tau_{i-1}} + \sigma_{f(\tau_{i-1}, \tau_i)} \sqrt{\Delta \tau_i} N_i \mu_J + \mu_{\tau_i} \\ \text{Var}[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A] &= \sigma_{f(\tau_{i-1}, \tau_i)}^2 \Delta \tau_i (1 + N_i \sigma_J^2) + \sigma_{c,A}^2 + q_i^A \sigma_\varepsilon^2 \end{aligned}$$

Corresponding to A_{τ_i} it must be that $a_{\tau_i} \in (\ln(A_{\tau_i} - K), \ln A_{\tau_i}]$. Thus we can compute

$$\begin{aligned} & p(A_{\tau_i} | m_{\tau_{i-1}}, N_i, (q_i^A)) \tag{11} \\ &= \Phi \left(\frac{\ln A_{\tau_i} - E[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A]}{\sqrt{\text{Var}[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A]}} \right) - \Phi \left(\frac{\ln(A_{\tau_i} - K) - E[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A]}{\sqrt{\text{Var}[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A]}} \right) \end{aligned}$$

The conditional distribution $f(m_{\tau_i} | a_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A)$ is again normal with mean and variance

$$\begin{aligned} E[m_{\tau_i} | a_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A] &= E[m_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A] \\ &+ \frac{\text{Cov}[m_{\tau_i}, a_{\tau_i} | m_{\tau_{i-1}}, \Delta N_i, q_i^A]}{\text{Var}[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A]} \times (a_{\tau_i} - E[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A]) \\ \text{Var}[m_{\tau_i} | a_{\tau_i}, m_{\tau_{i-1}}, N_i, q_i^A] &= \text{Var}[m_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A] - \frac{\text{Cov}[m_{\tau_i}, a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A]^2}{\text{Var}[a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A]} \end{aligned}$$

where

$$\begin{aligned} E[m_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A] &= \mu_m \Delta \tau_i + (1 - \rho \Delta \tau_i) m_{\tau_{i-1}} + \sigma_{f(\tau_{i-1}, \tau_i)} \sqrt{\Delta \tau_i} N_i \mu_J \\ \text{Var}[m_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A] &= \sigma_{f(\tau_{i-1}, \tau_i)}^2 \Delta \tau_i (1 + N_i \sigma_J^2) \\ \text{Cov}[m_{\tau_i}, a_{\tau_i} | m_{\tau_{i-1}}, N_i, q_i^A] &= \sigma_{f(\tau_{i-1}, \tau_i)}^2 \Delta \tau_i (1 + N_i \sigma_J^2) \end{aligned}$$

Similarly, $f(c_i^A | a_{\tau_i}, m_{\tau_i}, q_i^A)$ is normal with mean and variance

$$\begin{aligned} E[c_i^A | a_{\tau_i}, m_{\tau_i}, q_i^A] &= \mu_{\tau_i} + \frac{\sigma_{c,A}^2}{\sigma_{c,A}^2 + q_i^A \sigma_\varepsilon^2} \times (a_{\tau_i} - m_{\tau_i} - \mu_{\tau_i}) \\ \text{Var}[c_i^A | a_{\tau_i}, m_{\tau_i}, q_i^A] &= \sigma_{c,A}^2 - \frac{\sigma_{c,A}^4}{\sigma_{c,A}^2 + q_i^A \sigma_\varepsilon^2} \end{aligned}$$

Appendix B: M-step of the MCEM algorithm

This appendix describes the parameter updates in the M-step of our MCEM algorithm. First, let us introduce some notation. Let $l_{X_i}^{k-1}$ denote the conditional expectation of a variable X_i given the observed data, \mathcal{D}_T and the parameter vector at the $(k-1)^{th}$ iteration, θ^{k-1}

$$l_{X_i}^{k-1} = E(X_i | \mathcal{D}_T, \theta^{(k-1)})$$

As discussed in the main text, the particle filter described in section 5.1 can be used to compute this quantity jointly with fixed-lag smoothing.

Similarly denote by $l_{X_i}^{k-1, c_A < 0}$ the expectation of X_i conditionally on θ^{k-1} and the event $\{c_A < 0\}$.

$$l_{X_i}^{k-1, c_A < 0} = E(X_i | \theta^{(k-1)}, c_A < 0)$$

These quantities can be analytically computed in our model.

Further, introduce the following notation

$$\begin{aligned} X_{i,1} &= 1_{\{N_i > 0\}} \frac{Z_{i,J} N_i}{N_i + 1_{\{N_i = 0\}}} \\ X_{i,2} &= 1_{\{N_i > 0\}} \frac{(N_i)^2}{N_i + 1_{\{N_i = 0\}}} \\ X_{i,3} &= 1_{\{N_i > 0\}} \frac{Z_{i,J}^2}{N_i + 1_{\{N_i = 0\}}} \\ A_i &= \begin{pmatrix} 1 \\ m_{\tau_i} \\ m_{\tau_{i-1}} \\ Z_{i,J} \end{pmatrix} \begin{pmatrix} 1 & m_{\tau_i} & m_{\tau_{i-1}} & Z_{i,J} \end{pmatrix} \end{aligned}$$

Then, the individual loglikelihood function can be written as

$$E(L_i^N(\theta) | \theta^{k-1}) = l_{N_i}^{k-1} \ln(\lambda) + (1 - l_{N_i}^{k-1}) \ln(1 - \lambda)$$

$$E(L_i^{Z_J}(\theta) | \theta^{k-1}) = -\frac{l_{X_{i,3}}^{k-1} - 2\mu_J l_{X_{i,1}}^{k-1} + \mu_J^2 l_{X_{i,2}}^{k-1}}{2\sigma_J^2} - \frac{\ln(\sigma_J^2)}{2} \times l_{1_{\{N_i > 0\}}}^{k-1}$$

$$E(L_i^m(\theta) | \theta^{k-1}) = -\frac{g_i(\mu_m, \rho) l_{A_i}^{k-1} g_i(\mu_m, \rho)}{2\sigma_{u,f(\tau_{i-1}, \tau_i)}^2 \Delta\tau_i} - \frac{\ln(\sigma_{u,f(\tau_{i-1}, \tau_i)}^2)}{2}$$

where

$$g_i(\mu_m, \rho) = \begin{pmatrix} -\mu_m \Delta\tau_i \\ 1 \\ -(1 - \rho \Delta\tau_i) \\ -\sqrt{\Delta\tau_i} \sigma_{f(\tau_{i-1}, \tau_i)} \end{pmatrix}$$

$$\begin{aligned} E(L_i^{MI,A} | \theta) &= -\frac{l_{c_{i,A}}^{k-1} - 2\mu_{\tau_i} l_{c_{i,A}}^{k-1} + \mu_{\tau_i}^2}{2\sigma_{c,A}^2} - \frac{\ln(\sigma_{c,A}^2)}{2} - \frac{l_{q_{i,A}\varepsilon_{i,A}^2}^{k-1}}{2\sigma_\varepsilon^2} - \frac{\ln(\sigma_\varepsilon^2)}{2} l_{q_{i,A}}^{k-1} \\ &+ l_{k_{i,A}}^{k-1} \times \left[-\frac{l_{c_{i,A}}^{k-1, c_A < 0} - 2\mu_{\tau_i} l_{c_{i,A}}^{k-1, c_A < 0} + \mu_{\tau_i}^2}{2\sigma_{c,A}^2} - \frac{\ln(\sigma_{c,A}^2)}{2} \right] \end{aligned}$$

$$\begin{aligned} E(L_i^{MI,B} | \theta) &= -\frac{l_{c_{i,B}}^{k-1} - 2\mu_{\tau_i} l_{c_{i,B}}^{k-1} + \mu_{\tau_i}^2}{2\sigma_{c,B}^2} - \frac{\ln(\sigma_{c,B}^2)}{2} - \frac{l_{q_{i,B}\varepsilon_{i,B}^2}^{k-1}}{2\sigma_\varepsilon^2} - \frac{\ln(\sigma_\varepsilon^2)}{2} l_{q_{i,B}}^{k-1} \\ &+ l_{k_{i,B}}^{k-1} \times \left[-\frac{l_{c_{i,B}}^{k-1, c_B < 0} - 2\mu_{\tau_i} l_{c_{i,B}}^{k-1, c_B < 0} + \mu_{\tau_i}^2}{2\sigma_{c,B}^2} - \frac{\ln(\sigma_{c,B}^2)}{2} \right] \end{aligned}$$

Now we want to come up with new parameter iterates that increase the value of the complete data loglikelihood. Instead of a full-fledged maximization, we use iterations that increase on the objective value. In our case this will be much simpler and still ensure convergence of the algorithm.

$$\begin{aligned} \lambda^{(k)} &= \frac{\sum_{i=1}^T l_{N_i}^{k-1}}{T} \\ (\mu_m, \rho)^{(k)} &= \arg \max_{\mu_m, \rho} \sum_{i=1}^T -\frac{g_i(\mu_m, \rho) l_{A_i}^{k-1} g_i(\mu_m, \rho)}{2(\sigma_{u,f(\tau_{i-1}, \tau_i)}^{(k-1)})^2 \Delta\tau_i} \\ \theta_{\sigma_u}^{(k)} &= \arg \max_{\theta_{\sigma_u}} \sum_{i=1}^T -\frac{g_i(\mu_m^{(k)}, \rho^{(k)}) l_{A_i}^{k-1} g_i(\mu_m^{(k)}, \rho^{(k)})}{2\sigma_{u,f(\tau_{i-1}, \tau_i)}^2 \Delta\tau_i} - \frac{\ln(\sigma_{u,f(\tau_{i-1}, \tau_i)}^2 \Delta\tau_i)}{2} \\ \mu_J^{(k)} &= \frac{\sum_{i=1}^T l_{X_{i,1}}^{k-1}}{\sum_{i=1}^n l_{X_{i,2}}^{k-1}} \\ (\sigma_J^2)^{(k)} &= \frac{\sum_{i=1}^T [l_{X_{i,3}}^{k-1} - 2\mu_J^{(k)} l_{X_{i,1}}^{k-1} + (\mu_J^{(k)})^2 l_{X_{i,2}}^{k-1}]}{\sum_{i=1}^n l_{1_{\{N_i > 0\}}}^{k-1}} \end{aligned}$$

$$\begin{aligned} \theta_{\mu}^{(k)} &= \arg \max_{\theta_{\mu}} \sum_{i \in I_A} -\frac{-2l_{c_{i,A}}^{k-1} \mu_{\tau_i} + \mu_{\tau_i}^2}{2(\sigma_{c,A}^{(k-1)})^2} + l_{k_{i,A}}^{k-1} \times \left[-\frac{-2l_{c_{i,A}}^{k-1, c_A < 0} \mu_{\tau_i} + \mu_{\tau_i}^2}{2(\sigma_{c,A}^{(k-1)})^2} \right] \\ &+ \sum_{i \in I_B} -\frac{-2l_{c_{i,B}}^{k-1} \mu_{\tau_i} + \mu_{\tau_i}^2}{2(\sigma_{c,B}^{(k-1)})^2} + l_{k_{i,B}}^{k-1} \times \left[-\frac{-2l_{c_{i,B}}^{k-1, c_B < 0} \mu_{\tau_i} + \mu_{\tau_i}^2}{2(\sigma_{c,B}^{(k-1)})^2} \right] \end{aligned}$$

$$\sigma_{c,A}^{(k)} = \sqrt{\frac{\sum_{i \in I_A} \left[l_{c_{i,A}}^{k-1} - 2l_{c_{i,A}}^{k-1} \mu_{\tau_i}^{(k)} + (\mu_{\tau_i}^{(k)})^2 + l_{k_{i,A}}^{k-1} \left[l_{c_{i,A}}^{k-1, c_A < 0} - 2l_{c_{i,A}}^{k-1, c_A < 0} \mu_{\tau_i}^{(k)} + (\mu_{\tau_i}^{(k)})^2 \right] \right]}{T_A + \sum_{i \in I_A} l_{k_{i,A}}^{k-1}}}$$

$$\sigma_{c,B}^{(k)} = \sqrt{\frac{\sum_{i \in I_B} \left[l_{c_{i,B}}^{k-1} - 2l_{c_{i,B}}^{k-1} \mu_{\tau_i}^{(k)} + (\mu_{\tau_i}^{(k)})^2 + l_{k_{i,B}}^{k-1} \left[l_{c_{i,B}}^{k-1, c_B < 0} - 2l_{c_{i,B}}^{k-1, c_B < 0} \mu_{\tau_i}^{(k)} + (\mu_{\tau_i}^{(k)})^2 \right] \right]}{T_B + \sum_{i \in I_B} l_{k_{i,B}}^{k-1}}}$$

$$\sigma_{\varepsilon}^{(k)} = \sqrt{\frac{\sum_{i \in I_B} l_{q_{i,B} \varepsilon_{i,B}^2}^{k-1} + \sum_{i \in I_A} l_{q_{i,A} \varepsilon_{i,A}^2}^{k-1}}{\sum_{i \in I_B} l_{q_{i,B}}^{k-1} + \sum_{i \in I_A} l_{q_{i,A}}^{k-1}}}$$

References

- Acharya, V. V., and T. C. Johnson, 2007, Insider trading in credit derivatives, *Journal of Financial Economics* 84(1), 110-141.
- Barclay, M. J., and T. Hendershott, 2004, Liquidity Externalities and Adverse Selection: Evidence from Trading after Hours, *Journal of Finance* 59 (2), 681-710.
- Biais, B., and F. Declerck, 2007, Liquidity, Competition and Price Discovery in the European Corporate Bond Market, Mimeo, Toulouse School of Economics.
- Bjornes G. H., and D. Rime, 2005, Dealer Behavior and Trading Systems in the Foreign Exchange Market, *Journal of Financial Economics* 75(3), 571-605.
- 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.

- Cappe, O., and E. Moulines, 2005, On the Use of Particle Filtering for Maximum Likelihood Parameter Estimation, European Signal Processing Conference (EUSIPCO), Antalya, Turkey.
- 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.
- Duan, J.C. and A. Fulop, 2007a, A Stable Estimator for the Information Matrix under EM, University of Toronto working paper.
- Duan, J.C. and A. Fulop, 2007b, How Frequently Does the Stock Price Jump? An Analysis of High-Frequency Data with Microstructure Noises, University of Toronto working paper.
- Fleming, Michael J., 1997, The round-the-clock market for U.S. Treasury securities, FRBNY Economic Policy Review, 9-32.
- Foucault Th., S. Moinas, and E. Theissen, 2007, Does anonymity matter in electronic limit order markets ?, *Review of Financial Studies* 20(5), 1707-1747.
- Hasbrouck J., 2000, The Dynamics of Discrete Bid and Ask Quotes, *Journal of Finance* 54(6), 2109-2142.
- Hasbrouck J., 2003, Liquidity in the Futures Pits: Inferring Market Dynamics from Incomplete Data, *forthcoming Journal of Financial and Quantitative Analysis*.
- Ho T., and H. Stoll, 1983, The dynamics of Dealer markets under competition, *Journal of Finance* 49, 1639-1664.
- Huang R., J. Cai, and X. Wang, 2002, Information-Based Trading in Treasury Note Interdealer Broker Market, *Journal of Financial Intermediation* 11, 269-296.
- 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.
- Reiss P. C., and I. Werner, 1998, Does Risk Sharing Motivate Interdealer Trading? *Journal of Finance* 53(5), 1657-1703.

- Reiss P. C., and I. Werner, 2004, Anonymity, Adverse Selection and the Sorting of Interdealer Trades, *Review of Financial Studies* 18(2), 599-636
- 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
Characteristics of the Entities

Panel A
Characteristics

	Entity			
	Ford	GMAC	Sears Acceptance	France Telecom
Currency	USD	USD	USD	EUR
Country	United States	United States	United States	France
Sector	Automobile	Automobile	Retail	Telecommunication

Panel B
Credit Rating

In bold, from investment to speculative grade.

Rating type: LT Issuer

Agency: Standard & Poors

Entity	Date	Rating	Credit Event
Sears	01-Jan-04	A-	
	29-Mar-05	BB+	Rating Downgraded
Ford	01-Jan-04	BBB-	
	05-May-05	BB+	Rating Downgraded
	01-May-06	BB-	Rating Downgraded
	28-Jun-06	B+	Rating Downgraded
	19-Sep-06	B	Rating Downgraded
GMAC	01-Jan-04	BBB	
	14-Oct-04	BBB-	Rating Downgraded
	05-May-05	BB	Rating Downgraded
	28-Nov-06	BB+	Rating Upgraded
France Telecom	01-Jan-04	BBB-	
	18-Feb-04	BBB+	Rating Upgraded
	02-Oct-05	A-	Rating Upgraded

Table 2
Characteristics of the bid, ask prices, and trade prices

Panel A
 Descriptive statistics from the raw dataset

Data is between 2004-2006

	Entity			
	Ford	GMAC	Sears Acceptance	France Telecom
<i>Daily Average # bid</i>	5,81	6,65	2,49	4,32
<i>N bid</i>	4265	4936	1356	3078
<i>Average bid</i>	332	344	106,23	41,27
<i>Daily Average # ask</i>	5,12	6,18	2,07	4,24
<i>N ask</i>	3759	4589	1127	3017
<i>Average ask</i>	341	352	114,65	42,78
<i>Daily Average # trade</i>	2,57	4,46	0,78	1,42
<i>N trade</i>	1891	3313	425	997
<i>Average trade price</i>	367	393	113,46	42,38

Panel B
 Results from the dataset used for estimation

	Entity			
	Ford	France Telecom	GMAC	Sears Acceptance
<i>N bid</i>	2611	1640	2991	714
<i>Average bid</i>	335	40,49	351,45	107,34
<i>N ask</i>	2149	1583	2681	509
<i>Average ask</i>	343	40,95	359,25	115,59
<i>N bid/ask pair</i>	1420	1276	1656	568
<i>Average midpoint(*)</i>	331	43,61	331,185	109,42
<i>Average b/a spread(*)</i>	10,08	2,87	9,7	9,26
<i>Risky Annuity Factor(**)</i>	3,84	4,32	3,84	4,20
<i>Average cost of a round-trip</i>	38,68	12,39	37,22	38,89
<i>Minimum tick, K</i>	1	0,5	0,5	1

Data is between 2004-2006

At most 1 bid and ask observation has been kept for an identical time stamp (this is per minute)

Only observations where the ask is higher than the bid has been kept

(*) These statistics have been computed using the observation with both a bid and an ask

(**) We computed the cost of the 5 year risky annuity for the 5-year

CDS using a 5% interest rate 40% recovery value

Table 3
Panel A: US names

Time	Ford		GMAC		Sears Acceptance	
	σ_u	μ	σ_u	μ	σ_u	μ
05:30-07:30	0,6085 <i>0,0338</i>	0,0054 <i>0,0010</i>	0,6954 <i>0,0572</i>	0,0069 <i>0,0010</i>	0,1974 <i>0,1379</i>	<i>0,0282</i>
07:30-09:30	0,7093 <i>0,0315</i>	-0,0102 <i>0,0016</i>	0,8356 <i>0,0322</i>	-0,0066 <i>0,0015</i>	0,6991 <i>0,0918</i>	0,0035 <i>0,0084</i>
09:30-14:30	0,6028 <i>0,0260</i>	-0,0068 <i>0,0015</i>	0,7555 <i>0,0254</i>	-0,0079 <i>0,0014</i>	0,7135 <i>0,0724</i>	0,0034 <i>-0,0080</i>
14:30-16:30	0,6158 <i>0,0426</i>	-0,0044 <i>0,0017</i>	0,8669 <i>0,0569</i>	-0,0045 <i>0,0015</i>	0,6299 <i>0,1972</i>	0,0045 <i>-0,0126</i>
16:30-05:30	0,1958 <i>0,0132</i>	-0,0037 <i>0,0020</i>	0,1735 <i>0,0232</i>	0,0032 <i>0,0016</i>	0,2017 <i>0,0326</i>	0,0070 <i>-0,0694</i>
Parameters						0,0586
$\sigma_{c,A}$	0,0166 <i>0,0005</i>		0,0163 <i>0,0005</i>		0,0380 <i>0,0019</i>	
$\sigma_{c,B}$	0,0159 <i>0,0005</i>		0,0174 <i>0,0005</i>		0,0403 <i>0,0016</i>	
σ_E	0,0816 <i>0,0034</i>		0,0844 <i>0,0031</i>		0,2071 <i>0,0251</i>	
λ	0,1578 <i>0,0108</i>		0,1483 <i>0,0080</i>		0,1222 <i>0,0185</i>	
μ_J	0,4142 <i>0,2386</i>		0,1123 <i>0,2627</i>		1,8867 <i>0,9282</i>	
σ_J	5,8427 <i>0,2084</i>		6,2286 <i>0,2014</i>		9,3800 <i>0,9730</i>	

Panel B: European name

Time	France Telecom	
	σ_u	μ
00:30-02:30	0,2526 <i>0,0353</i>	0,0267 <i>0,0014</i>
02:30-04:30	0,5606 <i>0,0512</i>	0,0124 <i>0,0013</i>
04:30-07:30	0,4262 <i>0,0462</i>	0,0018 <i>0,0018</i>
07:30-10:30	0,3564 <i>0,0380</i>	0,0030 <i>0,0018</i>
10:30-14:30	0,3202 <i>0,0363</i>	-0,0009 <i>0,0023</i>
14:30-00:30	0,0306 <i>0,0523</i>	0,0095 <i>0,0049</i>
Parameters	#N/A	#N/A
$\sigma_{c,A}$	0,0254 <i>0,0008</i>	
$\sigma_{c,B}$	0,0253 <i>0,0008</i>	
σ_E	0,0598 <i>0,0053</i>	
λ	0,2879 <i>0,0243</i>	
μ_J	-0,0484 <i>0,1921</i>	
σ_J	5,0601 <i>0,3840</i>	

Figure 1:
Quote and Trade data for GMAC, for the period 2005 June 1–2005 June 5




Figure 2:
Intradaily frequency of trades and quotes, Ford 2004–2006




Figure 3:
Intradaily frequency of trades and quotes, GMAC 2004–2006




Figure 4:
Intradaily frequency of trades and quotes, Sears 2004–2006




Figure 5:

Intradaily frequency of trades and quotes, France Telecom 2004–2006




Figure 6:
GMAC CDS spreads 2005 May

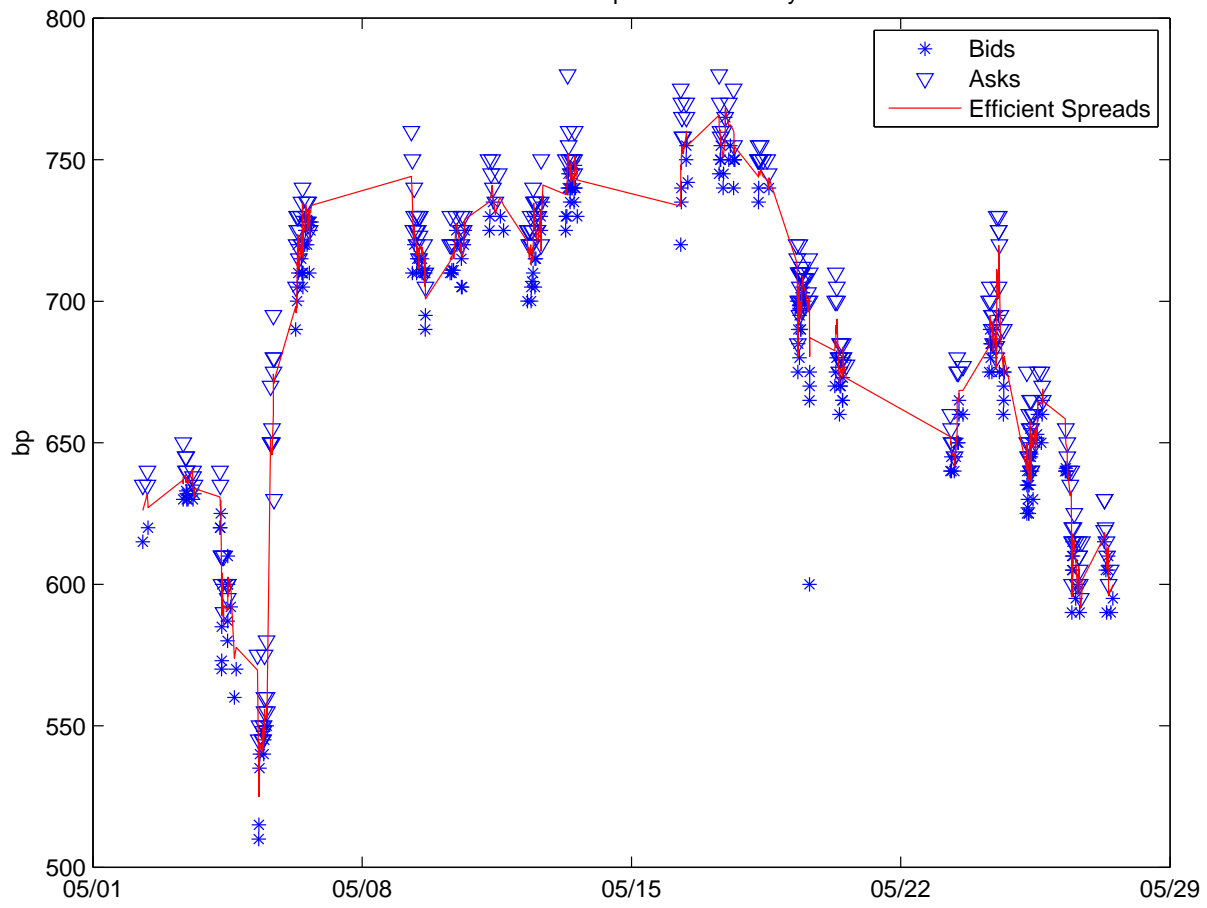


Figure 7:
Ford

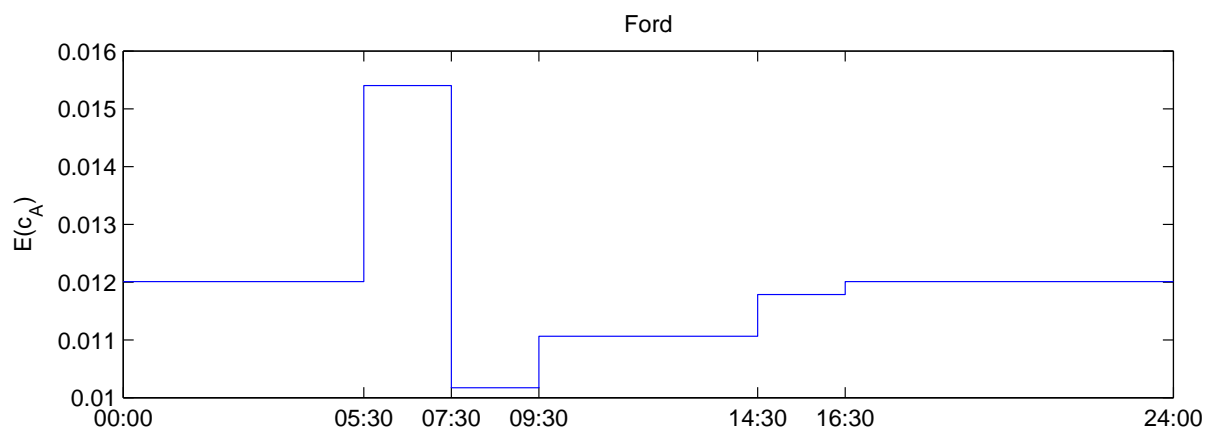
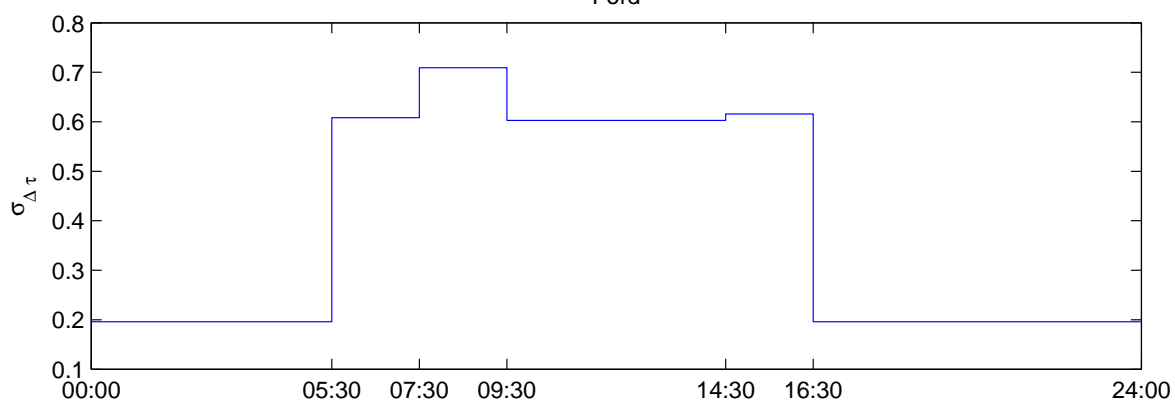


Figure 8:
GMAC

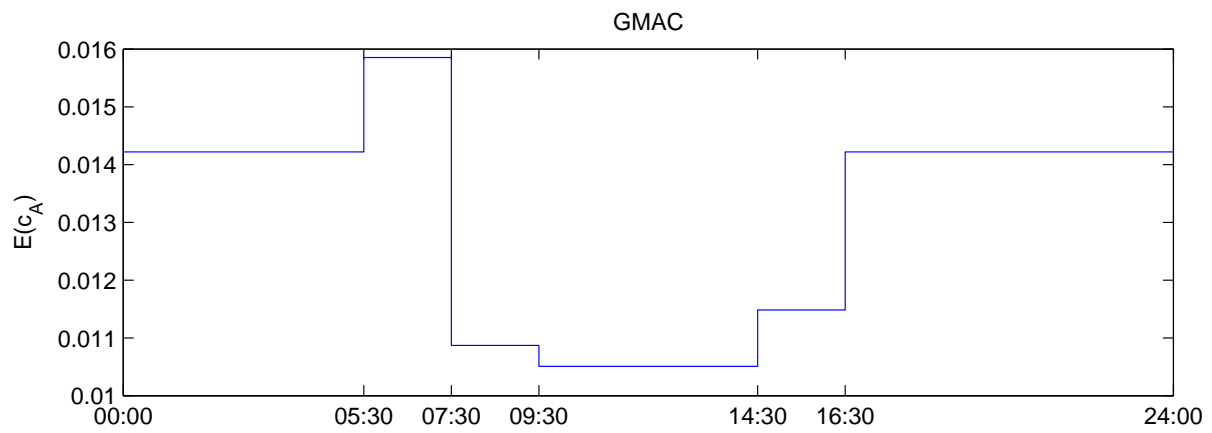



Figure 9:
Sears





Figure 10:
FT





Figure 11:
Ford CDS spreads 2004–2006




Figure 12:
GMAC CDS spreads 2004–2006




Figure 13:
Sears CDS spreads 2004–2006




Figure 14:
FT CDS spreads 2004–2006

