Modelling the Bid and Ask Prices of Illiquid CDSs

Michael B. Walker[1,2]

Working Paper
First version: May 28, 2010[3]

Abstract

CDS (credit default swap) contracts that were initiated some time ago frequently have spreads and/or maturities that are not available on the current market of CDSs, and are thus illiquid. This article introduces an incomplete-market approach to valuing illiquid CDSs that, in contrast to the commonly used complete-market risk-neutral approach to valuation, takes into account the risky nature of the illiquid CDS, and allows a dealer who buys an illiquid CDS from an investor to determine ask and bid prices (which differ) in such a way as to guarantee a minimum positive expected return on the deal. An alternative procedure, which replaces the expected return by an analogue of the Sharpe ratio, is also discussed. The approach to pricing just described belongs to the good-deal category of approaches, since the dealer decides what it would take to make an appropriate expected return, and sets the bid and ask prices accordingly. A number of different hedges are discussed and compared within the general framework developed in the article. These include the hedge that enforces the no-arbitrage bounds, a vanilla hedge making use of a single CDS from the market having the same maturity and notional as the illiqid CDS (but a different spread), and an optimal hedge that minimizes the capital at risk for the dealer (who is identified above), conditional on the dealer achieving a desired minimum expected return and on the bid or ask price for the transaction having a certain definite value. The approach is implemented numerically, and example plots of important quantities are given.

keywords: credit defaults swaps, CDSs, hedging, valuation, incomplete markets

[1]Department of Physics, University of Toronto, Toronto, ON M5S 1A7, CANADA; email: walker@physics.utoronto.ca; telephone: (416) 978-3821.

[2]I would like to thank Tomasz Bielecki, Jeremy Staum, Dan Rosen, and other participants of the Fields Institute Thematic Program on Quantitative Finance, Toronto, January-June, 2010, for helpful discussions. The support of the Natural Sciences and Engineering Research Council of Canada is acknowledged.

[3]Preprint of an article submitted for consideration in [IJTAF] [2010] [copyright World Scientific Publishing Company] [http://www.worldscinet.com/ijtaf/]

1 Introduction

Seasoned CDS contracts (i.e. CDS contracts that were initiated some time ago) that have termination dates and/or spread payments that differ from the termination dates and spread payments of the contracts on the current liquid CDS market are illiquid. These illiquid contracts can be approximately hedged in terms of portfolios of liquid CDS contracts, but the value of the hedged position is uncertain, due to its dependence on the realized values of the default time and the recovery rate, both of which are random variables. The existence of illiquid contracts that can not be perfectly hedged in the market of liquid CDSs means that the CDS market is incomplete. The common complete-market risk-neutral approach to valuing illiquid CDSs does not take into account the risks that remain after hedging, nor does it take into account the requirement a potential purchaser of the illiquid CDS should make a positive expected profit on the transaction. The purpose of this article is to describe a new approach which takes into account the factors just indicated in the valuation procedure, gives bid and ask prices that are, in general, different, and gives the complete-market risk-neutral result in cases where it is applicable (e.g. in reproducing the market prices of the liquid CDS contracts).

A credit default swap (CDS) is a credit derivative that provides insurance against the loss of notional of a corporate bond on default. The details of standard CDS contracts changed at the time of the so-called CDS Big Bang in April, 2009, with these details being described more fully in Section 2. Under the post-CDS-Big-Bang contracts, the CDS protection buyer makes an upfront payment at the contract's inception, and a standardized quarterly premium payment (also called the spread payment) up until default or the termination date (also called the maturity) of the contract, whichever occurs first; in return, the protection buyer receives the loss of the corporate bond on default. Section 2 also describes the market of liquid CDSs in more detail, and gives examples of how illiquid CDSs arise as seasoned CDS contracts that no longer have running spreads and/or termination dates corresponding to those of contracts on the current liquid CDS market. Thus, this article adopts an idealized picture in which there exist CDS contracts for which there is a liquid market (and for which the liquidity is, in reality, not perfect) and CDS contracts for which there is no market (but of which dealers may have an inventory and which thus would perhaps not be considered perfectly illiquid).

Illiquid CDSs are commonly priced in the finance industry using an approach that will be called the complete-market risk-neutral approach. This approach, the details of which can be found in Beumee et al. (2009), on page 12 of BNP (2004), in Felsenheimer et al. (2006), and in Chapter 6 of O'Kane (2008), is based on the choice of a risk-neutral measure that is uniquely determined by calibration to all maturities on the current market and by the choice of a fixed value for the recovery rate. The value of the illiquid CDS is then established as the expected present value of the payoff stream of the illiquid CDS, calculated using the uniquely-determined risk-neutral measure. In the special case that there is a CDS on the current market having the same time to maturity T_M as that of the illiquid CDS, the value of the illiquid CDS is given by the well-known result

$$u_M^{Old} = u_M + (w_M - w_M^{Old}) \times RPV01(T_M) \tag{1}$$

where u_M in the upfront payment that is made to acquire a long-protection CDS contract with a time to maturity of T_M years (measured from today) on the current market at at spread of w_M, and w_M^{Old} is the spread of the illiquid contract. Also, $RPV01(T_M)$ is said to be the present value of a risky 1 bp/yr paid quarterly until maturity at T_M or until default, whichever is sooner. The quantity $RPV01(T_M)$ could also be called a risky duration. In the complete-market risk-neutral approach, the bid and ask prices for the illiquid CDS are the same, both being equal to u_M^{Old}. The difficulty with this approach is that is does not adequately take account of the risk of the illiquid CDS, and furthermore, a dealer who takes over an illiquid CDS contract from an investor at the price given by Eq. 1 is not even guaranteed to make a positive expected profit on the deal.

In valuing an illiquid CDS, this article considers the payment required by or to a dealer who takes over an investor's illiquid position; this provides what is called the unwind value for the position. (The mark-to-market value of an illiquid CDS contract can be defined to be its unwind value.) It is assumed that when the dealer takes over an illiquid CDS, the dealer will hedge this position to reduce risk. In the case of an illiquid CDS contract, perfect hedging is not possible. Thus, a dealer who pays an investor a definite amount in taking over the investor's illiquid CDS position, and then hedges the position in order to reduce the risk, nevertheless has the possibility of ending up having a positive realized loss on the transaction. The present value of the maximum possible loss on the transaction is called, in this article, the capital at risk. Another quantity of interest to the dealer will be the expected return on the capital at risk, also called simply the expected return, and defined more precisely in Eq. 18 below. The deal will require the expected return to be positive, and sufficient to compensate for the risks undertaken.[4] The approach of this article requires that the dealer make a subjective decision that all deals providing an expected return greater than some lower bound are acceptable. This then allows the dealer to establish a lower good-deal bound for the ask price, and an upper good-deal bound for the bid price. Pricing according to good-deal criteria has been described, in other contexts, in Staum (2008), and more specifically, in Bernardo and Ledoit (2000); Carr et al. (2001); Černỳ and Hodges (2002); Cochrane and Saà-Requejo (2000); Staum (2004). The establishment of good-deal bid and ask bounds for an illiquid CDS is discussed in detail in Sections 5, 8, and 9. The use of an analogue of the Sharpe ratio in establishing good-deal bounds, as an alternative to the use of the expected return, is also discussed.

The first step towards finding the good-deal bid and ask bounds for a given

[4]This article assumes that counterparty risk has been nullified by sufficiently strong collateral arrangements. If this is not the case, the procedure of the present paper will represent only a first step in the process of determining bid and ask prices, and the results obtained will have to adjusted to take into account the counterparty risk.

illiquid CDS contract, will be to establish, for CDSs, a procedure for determining the no-arbitrage bounds on these prices, and the hedging portfolios that enforce these bounds. For a general description of the problem of no-arbitrage bounds, see, for example, Staum (2008); Pliska (1997). The challenge here is to find a model and a procedure that works well for CDSs (see Section 3 for this). The random nature of the recovery rate has an important impact on the risk associated with the hedged position, and must be taken into account. (Many commonly used models for CDS valuation take the recovery rate to be a constant.) The hedging portfolios that enforce the no-arbitrage bounds will be called here the LUB (least upper bound) and GLB (greatest lower bound) hedging portfolios, and are also referred to elsewhere as super-replicating portfolios. Section 4 gives numerical examples for no-arbitrage bounds and the corresponding hedging portfolios. Although the no-arbitrage bounds do place limits on the bid and ask prices, these limits, as expected, are not strong, and, as noted above, a good-deal criterion will be added to further refine the bid and ask prices. The hedging portfolios that enforce the no-arbitrage bounds are nevertheless found to be useful hedges for reducing the risk of the illiquid position.

The model and the basic techniques developed in this article to determine bid and ask prices are all described and illustrated in some detail in Sections 3 to 8 with examples in terms of hedging using the porfolios that enforce the no-arbitrage bounds. As an example, an illiquid CDS with a time-to-maturity of 5 yrs is hedged with a multi-CDS portfolio that includes CDSs from the market with times-to-maturity of 1, 2, 3, 4 and 5 yr. In the multi-CDS hedge, the losses on default are not exactly hedged (as they are with the vanilla hedge described in the following paragraph), but one does a better job of hedging the default-time uncertainty.

A second type of hedge, called the vanilla hedge, uses a single CDS from the current market with the same maturity and notional as the illiquid CDS, together with a bank deposit, as the hedging portfolio. The inclusion of the bank deposit allows the vanilla hedge to be treated by the approach of this paper, which determines different bid and ask prices to account for the risk of the hedged position. The vanilla hedge has the attractive feature that, in the combined position (e.g. long illiquid CDS, short hedging portfolio) the loss on default of the hedging portfolio exactly cancels the loss on default of the illiquid CDS. There is however still an uncertainty in the realized value of the combined position due to the uncertainty of the default time (and hence the cumulative spread payment). The vanilla hedge described in this article is the same hedge as is used to establish the complete-market risk-neutral pricing result of Eq. 1, except for the inclusion of the bank deposit. It will be of interest to the reader familiar with the complete-market risk-neutral approach that leads to Eq. 1, that the same hedge, except for the bank deposit, leads to different bid and ask spreads and a necessarily positive expected profit when the approach of this article is used. The performance of the vanilla hedge will be compared with that of the multi-CDS hedge that enforces the no-arbitrage bounds in many of the deveopments of Sections 3 to 8.

A small modification of the methods developed in Sections 3 to 8 allows the determination of various 'optimal' hedges. For example one can find a hedge that provides the dealer (identified above) with a definite expected return and allows the determination of bid and ask prices in such a way that the capital at risk is minimized. Or one can find a hedge such that the capital at risk is minimized conditional on the bid or ask price and the expected return have desired values. These optimal hedges are described in Section 9.

For all types of hedges, a 'good-deal' approach to valuation has been adopted. Good deals are specified both in terms of a minimum expected return required by the buyer of an illiquid position, or in terms of a minimum value of an effective Sharpe ratio required by the buyer.

Another type of multi-CDS hedge, which is not discussed in this article, is described, together with its pros and cons, in Chapter 6 of O'Kane (2008). The complete-market risk-neutral approach to valuation is used in conjunction with this hedge.

Because the hedging of an illiquid CDS is not exact, the hedged position (long CDS, short hedge) has a spectrum of realized present values. Sections 5 and 6 develop a method of describing this spectrum and the corresponding physical probability density, and give example numerical results for these. This physical probability density of the realized present value of the hedged position is an important quantity since it gives a broad overall description of risk for the hedged position. This probability density has a discrete component as well as a piecewise continuous component. A qualitative analytic explanation of the discontinuities in the piecewise continuous componest is given; this is useful because when discontinuities occur in the numerical implementation, one would like to make sure that one understands their origin.

Section 10 develops scaling relations showing how certain quantities (e.g. the bid and ask spreads established for the illiquid CDS, the spectrum of realized present values of the hedged position for the illiquid CDS, the physical probability measure for this spectrum, etc.) scale with the difference $(w_M - w_M^{Old})$, which quantity also appears in Eq. 1. These results hold only for the special case that there is a CDS on the current market with the same maturity as that of the illiquid CDS. Also, scaling does not apply when the hedge is the optimal hedge of subsection 9.2. Subsection 10.1 gives numerical results contrasting the cases where scaling is present, and where it is absent.

Even though the assumed physical probability distributions for the default time and the recovery rate are expected to be established, in so far as is possible, by careful research, there will be an element of subjectiveness in arriving at precise values for these quantities. Because of this, it is useful to have a method of determining the robustness of the results with respect changes in the assumed physical probability distribution. A method of doing this is described in Section 11, and numerical examples are given. For the numerical example followed throughout the course of the article, the calculated bid and ask prices are found to be reasonably stable with respect to changes in the physical probability distribution.

Section 12 presents a brief summary of the main results and conclusions.

The literature on the application of incomplete market ideas to problems in the valuation of illiquid derivatives is too large to review in this article focused on CDSs. Fortunately, a detailed and thorough recent review, from which the author has profited, is available in Staum (2008). An elementary introduction to incomplete market ideas, sufficient for an understanding of the present article, can be found in Pliska (1997). Articles on good-deal bounds for incomplete-market problems include Bernardo and Ledoit (2000); Carr et al. (2001); Černỳ and Hodges (2002); Cochrane and Saà-Requejo (2000); Staum (2004). Bielecki et al. (2004a,b) describe approaches to pricing defaultable claims in incomplete markets. However, there is nothing in the literature which describes how to use incomplete-market ideas to obtain bid and ask prices (with a non-zero bid-ask spread) for illiquid CDSs. In summary, the general ideas which form the basis for this article are well-known from the theory of incomplete markets, so that the originality of the presentation comes from finding a simple way of implementing these ideas for the case of CDSs. An inovation that might be of interest beyond the CDS problem tackled here is the use of a risk measure that tends to zero risk as the arbitrage-free bounds are approached.

As noted above, what is referred to in this article as the liquid market of CDSs will not in fact be perfectly liquid. The effects of a relatively small degree of illiquidity in determining the bid-ask spreads in the liquid market is not discussed here, but is an important question. Brigo et al. (2010) give a recent survey of this literature. Section 4 of this article describes how a small modification of the procedures described below can take into account this 'liquid-market' bid-ask spread.

2 The CDS Market and Illiquid CDSs

In April 2009, in what is known as the CDS Big Bang, certain changes to the standard CDS contract were introduced. Pre-CDS-Big-Bang contracts were priced according to a running par spread with zero upfront payment. Post-CDS-Big-Bang contracts are priced according to an upfront payment with a standardized running spread. In North America, there is liquidity in contracts with spreads having the two standard values of $w = 100$ bp/yr (for investment-grade names), and $w = 500$ bp/yr (for high-yield names). In Europe, the standardized spread values are $w = 25, 100, 500, 1000$ bp/yr. For a detailed discussion of post-CDS-Big-Bang contracts see (Markit, 2009a,b).

Standard ISDA CDS contracts have termination dates of March 20, June 20, September 20, and December 20 for any given year (e.g. see Markit (2009a)). Typically, for all contracts concluded during the quarter ending on March 20th of a given year, there is liquidity only in contracts which have a termination date of March 20th some integral number of years in the future, and there is no liquidity in contracts terminating on one of the other standardized termination dates, except, perhaps, for a contract termination date of September 20, 6 months in the future. On March 21, there is a quarterly roll at which time the liquid contracts become those with a termination date of June 20. Furthermore, the number of annual maturities with liquidity is limited. For example, for a reasonbly liquid name, liquidity might be available in contracts with a time to maturity of 1,2,3,5,7 and 10 years.

It is now clear that, for a given reference name at a given time during a particular quarterly roll, there is liquidity in contracts of a limited number of termination dates, and of a single standardized spread. Other contracts concluded at previous times (called seasoned CDS contracts) may have termination dates and/or spreads that are not equal to those of the currently liquid contracts just described. The spreads of these other contracts could be different from those currently on the market for two reasons: a) the contracts could be legacy contracts concluded in the pre-CDS-Big-Bang era at a par spread determined by the market, and b) the contract could be post-CDS-Big-Bang North American contract concluded some time ago at a standardized spread of 100 bp/yr when the name was investment grade; if this name is today considered to be in the high yield category, liquidity might exist only in the 500 bp/yr spreads. Thus, seasoned CDS contracts could well be illiquid, which will occur when there is no contract on the current market having the same spread and termination date as the seasoned contract. Thus, it will be assumed in this article that investors and dealers can hold two essentially different CDS contracts, those that are currently on the liquid market, and those that are illiquid. This is admittedly an idealization, since even contracts that are considered to be on the market can be more or less liquid, depending on the name and the maturity, and dealers might have an inventory of illiquid CDSs.

The holder of a CDS contract for a given name that is on the currently liquid CDS market can easily and accurately hedge this contract simply by purchasing the offsetting contract, which will also be on the current market. On the other hand, the holder of an illiquid CDS contract will generally be unable to purchase an offsetting contract. Thus, the hedging of an illiquid contract will be carried out by purchasing a portfolio containing (at least one) liquid CDS contracts, as well as a bank deposit. Furthermore, the hedging of an illiquid contract will be only approximate, so that the hedger's hedged postion is risky, i.e. it has a realized present value that is uncertain (depending on the default time and recovery rate, both of which are random variables). This difference in the hedging procedures for liquid and illiquid CDS contracts must be reflected also in different valuation procedures. Liquid CDSs have values determined by the market. On the other hand, there is no market price for an illiquid CDS. As noted in Section 1, this article develops a framework for establishing good-deal bounds for the bid and ask prices of an illiquid CDS.

3 No-Arbitrage Bounds and Hedging Portfolios

Recall that an investor who buys a CDS providing protection against the loss on default of a given corporate bond will be said to hold a long CDS contract, and an investor who sells protection will be said to hold a short CDS contract.

This section begins by introducing notation describing the model used in this article. Following this, procedures for obtaining the no-arbitrage bounds imposed by the model on the bid and ask prices of an illiquid CDS contract, and for obtaining the hedging portfolios that enforce these bounds, will be outlined. The hedging portfolio that enforces the upper no-arbitrage bound will be called the LUB (least upper bound) hedging portfolio, and will be used to hedge short CDS contracts. The portfolio that enforces the lower no-arbitrage bound will be called the GLB (greatest lower bound) hedging portfolio, and will be used to hedge long CDS contracts. Typically, the illiquid CDS will have a running spread that differs from the running spreads of contracts on today's market, or a termination date that differs from the termination dates of contracts on today's market, or both. The general approach to hedging and valuation of illiquid contracts developed below, works equally well in all three cases.

The goal will be to construct useful hedging portfolios for an illiquid CDS contract from CDSs of maturities currently on the market, together with a cash deposit in a bank account. The CDSs selected from the market for a given hedging portfolio are labelled by the index $p = 1, 2, ..., K$, in order of increasing maturity. The maturity of the p-th CDS is called $T_{n(p)}$, and $n(p)$ is the number of premium payment times to maturity for this CDS. The illiquid CDS to be hedged will be assumed to have unit notional. The notional of the p-th CDS in the hedge is called $\alpha_{n(p)}$; a long CDS position is described by $\alpha_{n(p)} > 0$, whereas $\alpha_{n(p)} < 0$ describes a short CDS position. The total CDS notional present in the hedging portfolio at its inception, $\alpha_{total} = \sum_{p=1}^{K} \alpha_{n(p)}$, is an important characteristic (e.g. see Section 6). In the new CDS contract convention established at the time of the CDS Big Bang (Markit, 2009a), CDS contracts on the current market for a given name all have the same spread, but differ in their upfront payments. For greater generality, the approach as developed below will (except for the numerical examples) allow both the spreads and the upfront payments to depend on maturity.

Let $N = max\{n(K), M\}$. Then T_N is greater than or equal to the maturities of all CDSs in the hedging portfolio, as well as the maturity of the illiquid CDS. The $t = 0$ cash deposit in the bank account is written $\beta_{total} = \sum_{i=1}^{N} \beta_i$. If there has been no default up to and including the i-th premium payment time T_i, the amount $\beta_i \exp(r_F T_i)$ is paid at that time from the bank account to the holder of the hedging portfolio. Here, r_F is the risk-free interest rate, assumed constant in this paper. If default occurs at some time before maturity, then the entire amount remaining in the bank account is paid to the holder of the hedging portfolio at the default time.

Consider a long CDS contract from the current market having spread w_n and maturity T_n. The present value of the spread payment made by the contract holder at time T_i, provided no default has occurred up to and including time T_i, will be called $g_{n,i}$. Also, if default occurs at time $\tau = T_{i-1} + \delta$, $\delta \in (0, \delta_i]$, $\delta_i \equiv (T_i - T_{i-1})$, the present value of the loss payment made to the contract buyer minus the spread payment made by the contract buyer at the default time τ is called $h_{n,i}(\delta, \rho)$. These two present values are given explicitly as

$$g_{n,i} = w_n \delta_i d_i \delta_{i \leq n}, \quad h_{n,i}(\delta, \rho) = (1 - \rho - w_n \delta) d_i(\delta) \delta_{i \leq n}. \qquad (2)$$

Here, $d_i \equiv \exp(-r_F T_i)$ and $d_i(\delta) \equiv d_i \exp(r_F(\delta_i - \delta))$ are discount factors, and $\delta_{i\leq n}$ is unity if $i \leq n$ and zero otherwise. Note carefully that the quantities δ, δ_i and $\delta_{i\leq n}$ are all different, as are the quantities d_i and $d_i(\delta)$. The quantities $g^{Old}_{M,i}$ and $h^{Old}_{M,i}$ relating to the illiquid CDS are similarly defined except that w^{Old}_M replaces w_n, and M replaces n.

Today's value of a hedging portfolio constructed, as described above, from K CDSs on the market having notionals $\alpha_{n(p)}$, $p = 1, ..., K$, together with an initial bank deposit, is

$$V = \sum_{i=1}^{N} \beta_i + \sum_{p=1}^{K} \alpha_{n(p)} u_{n(p)} \tag{3}$$

Allowing both the upfront payment $u_{n(p)}$ (as in Eq. 3) and the spread w_n (as in Eq. 2) to depend on maturity gives a formulation of the problem that is applicable to both pre-CDS-Big-Bang and post-CDS-Big-Bang contracts. For the pre-CDS-Big-Bang contracts, $u_{n(p)} = 0$, whereas for the post-CDS-Big-Bang contracts $w_n = w$, independent of n.

The LUB (or GLB) hedging portfolio is defined to be the hedging portfolio of the smallest (or greatest) value V which satisfies the condition that the payoff by the hedging portfolio to its holder at any given time is greater than or equal to (or less than or equal to) the payoff of the illiquid long CDS to its holder at that time. These conditions are imposed in an optimization procedure carried out today, and must be satisfied for all values of the default time between 0 and T_N, and for all values of recovery rate between 0 and 1. The index σ is used to distinguish between the LUB and GLB hedging portfolios, with $\sigma = +$ for LUB and $\sigma = -$ for GLB; also, $(\sigma 1)$ is $(+1)$ when $\sigma = +$ and (-1) when $\sigma = -$. The procedure that determines the LUB or GLB hedging portfolio can therefore be written

$$V^{(\sigma)} = \text{minimize}[(\sigma 1)V] \tag{4}$$

subject to constraints (5) and (6) for $i = 1, 2, \ldots, N$. The constraint on the payoffs at the premium payment time T_i, provided that default occurs later than this time, is

$$(\sigma 1)\left(\beta_i - \sum_{p=1}^{K} \alpha_{n(p)} g_{n(p),i}\right) \geq (\sigma 1)\left(-g^{Old}_{M,i}\right). \tag{5}$$

The constraint on the payoff when default occurs at time $\tau = T_{i-1} + \delta$, $\delta \in (0, \delta_i]$, is

$$(\sigma 1)\left(\sum_{k=i}^{N} \beta_k + \sum_{p=1}^{K} \alpha_{n(p)} h_{n(p),i}(\delta, \rho)\right) \geq (\sigma 1) h^{Old}_{M,i}(\delta, \rho); \tag{6}$$

the latter constraint must hold for all δ satisfying $\delta \in (0, \delta_i]$, and for all ρ satisfying $\rho \in [0, 1]$.

Carrying out the procedure indicated in the sentence containing the instruction (4), subject to the contstraints that follow, gives the numerical values

maturity	1	2	3	4	5	7
par spreads (bp)	1068	1241	1282	1253	1219	1149
upfront (%)	5.25	12.47	18.08	21.56	24.05	27.00

Table 1: The given par spreads for General Motors Corporation senior CDSs on 20 March 2008 are as quoted by Thomson Datastream. These par spreads were converted to upfronts by using the ISDA CDS Standard Model as described in ISDA (2009) and in Beumee et al. (2009), and taking the running spread and the recovery rate to have the standard values of $w = 500$ bp/yr and $\rho = 20\%$ generally used for high-yield names with this model. Upfronts obtained in this way were used for the numerical examples of this report since recent upfront quotes from the market were not available to the author. General Motors filed for bankrupcy protection on June 1, 2009. The CDS recovery rate, determined by auction, was $\rho = 12.5\%$ (Reuters, 2009).

$\beta_i^{(\sigma)}, i = 1, \ldots, N$ and $\alpha_{n(p)}^{(\sigma)}, p = 1, 2, \ldots, K$ which define the LUB and GLB hedging portfolios, as well as the cost, $V^{(\sigma)}$, of purchasing these portfolios on today's market.

It is easy to see that $V^{(-)}$ is greatest lower bound of the range of arbitrage-free bid prices for the illiquid CDS. Suppose that an investor holding a long-protection illiquid CDS contract sells this contract to a dealer for a bid price of $u^{Old(-)} < V^{(-)}$. The dealer hedges this contract by shorting the GLB hedging portfolio on the market, thereby receiving the amount $V^{(-)}$. The dealer uses $u^{Old(-)}$ of this amount to pay the investor and pockets the remainder. Furthermore, the dealer's net position (long illiquid CDS contract and short GLB hedging portfolio) has only non-negative payoffs. Thus the dealer will make a riskless profit on this transaction and the bid price of $u^{Old(-)}$ is not arbitrage-free. Also, $V^{(-)}$ is the price of the most expensive portfolio that can be used to hedge a long position in the illiquid CDS. Thus, $V^{(-)}$ is the greatest lower bound of the range of arbitrage-free bid prices. Similarly, it can be shown that $V^{(+)}$ is the least upper bound for range of arbitrage-free ask prices.

4 Numerical Example for No-Arbitrage Bounds and Hedging Portfolios

This section describes numerical results for the no-arbitrage bounds and for the LUB and GLB hedging portfolios. Since other numercal results that make use of the same set of numerical input parameters will be given below, these parameters will be stated once and for all as follows:

Standard Parameter List

- illiquid CDS maturity: $T_M = 5$ years from today

- notional of illiquid CDS: unity
- the multi-CDS hedging portfolios contain maturities 1, 2, 3, 4, 5 years with upfronts as given in Table 1
- risk-free interest rate: $r_F = 2\%$
- running spread of CDSs of all maturities on the current market: 500 bp/yr
- running spread of illiquid CDS: $w_M^{Old} = 100$ bp/yr
- probability of default within 1 year: $PD_1 = 30\%$
- recovery rate probability density: γ_A of Fig. 2.
- expected return on capital at risk: $R_T = 25\%$.

A given numerical example does not necessarily use all of these parameters. For example, the numerical example of this section uses only the parameters specified in the first five items.

Numerical results for two different sets of LUB and GLB hedging portfolios have been obtained by following the procedure of Section 3, and are presented in Fig. 1. In the first case considered, the current CDS market will be assumed to consist of maturities of 1, 2, 3, 4, and 5 years as indicated in the Standard Parameter List. The LUB and GLB and hedging portfolios constructed in this case are called multi-CDS hedges. The values of the multi-CDS hedges are called $V^{(\sigma)}$. These values represent the costs of constructing the hedges from their components on that market, and are also the no-arbitrage bounds on the price of the illiquid CDS in the given market. (Also, in contrast to the values of the following paragraph, the quantities $V^{(\sigma)}$ do not carry a subscript v).

For the second case considered, the current CDS market will be assumed to consist of a single CDS having the same notional and maturity as those of the illiquid CDS to be hedged, and the GLB and LUB portfolios constructed under this assumption will be called vanilla hedging portfolios. This is similar to a common practice (e.g. see p. 98 of O'Kane (2008)) of hedging with a single CDS of the same maturity from the market, except that a bank deposit has been added to turn the hedging portfolio into an LUB or GLB portfolio. The values of the vanilla LUB and GLB portfolios (i.e. the cost of putting these portfolios together by buying their components on the market) are called $V_v^{(\sigma)}$ (with a subscript v). It turns out that in this case, the optimization procedure finds that the single market CDS in the hedging portfolio has the same notional as that of the illiquid CDS to be hedged.

Note in Fig. 1 that when $w_M^{Old} = 500$ bp/yr, the values of all hedging portfolios coincide. For this value of w_M^{Old}, all GLB and LUB portfolios consist of a single CDS from the market having a maturity of 5 years and a running spread of 500 bp/yr, equal to the running spread of the illiquid CDS, and the hedging is perfect. In fact, for this particular value of w_M^{Old}, the supposedly illiquid CDS is on the liquid market.

maturity	1	2	3	4	5	α_{total}	β_{total}
$\alpha^{(+)}$	-0.0319	-0.0342	-0.0368	-0.0395	1.0000	0.8576	0.1720
$\alpha^{(-)}$	-0.0405	-0.0434	- 0.0464	-0.0497	1.1800	1.0000	0.0000

Table 2: The table gives the notionals for each maturity of the CDS purchased on the market, as well as the initial cash deposit β_{total}, in the multi-CDS LUB and GLB hedging portfolios. The sum of the individual notionals in the LUB or GLB portfolio is α_{total}. The illiquid CDS being hedged has maturity $T_M = 5$ years, spread $w_M^{Old} = 100$ bp/yr, and notional 1.0000.

The notionals of all market CDSs present in the multi-CDS hedging portfolios when the illiquid CDS has a spread of $w_M^{Old} = 100$ bp/yr (as it would have been if the illiquid CDS contract was a North American post-CDS-Big-Bang contract concluded at a time when the reference name was investment grade) are given in Table 2.

Note in Table 2 that some notionals are negative (short protection positions) and some are positive (long protection positions). Long and short positions should have different upfront prices. The data from Table 1 shows only a single price, the mid-point price, since this was all that was available from the data source. Different bid and ask prices from the market can be taken into account by giving the long and short positions from Table 2 their appropriate prices and redoing the calculation of the notionals. The bid-ask spreads for CDSs on the liquid market are expected to be significantly smaller than the bid and ask spreads of illiquid CDSs which are the subject of this article.

Note from Fig. 1 that, for any value of w_M^{Old} (except 500 bp/yr), the value $V_v^{(+)} > V^{(+)}$. Recall that the value $V_v^{(+)}$ is obtained by minimizing this value subject to certain constraints. When the portfolio over which one is minimizing is extended from a single CDS on the market to include 4 more market CDSs, one expects to get a lower minimum. A corresponding remark explains the result that $V_v^{(-)} < V^{(-)}$ at any given w_M^{Old}. When the market consists of CDSs with maturities 1, 2, 3, 4, 5 yrs, the no-arbitrage bounds on the illiquid CDS price are given by $V^{(\sigma)}$, and the prices $V_v^{(\sigma)}$ are outside the no-arbitrage bounds. One can still use a vanilla hedge to hedge an illiquid CDS if one wishes, but one should be careful to have the sale price of an illiquid CDS lying within the no-arbitrage bounds the are appropriate for the existing CDS liquid market.

5 Valuing an Illiquid CDS Contract

An illiquid short CDS contract can be approximately hedged by combining it with a long LUB hedging portfolio. The hedged short illiquid CDS position (long LUB plus short illiquid CDS) then has only non-negative payoffs to the position holder. Similary, a long illiquid CDS contract can be approximately hedged with a short GLB hedging portfolio. The hedged long illiquid CDS position (long illiquid CDS and short GLB hedging portfolio) also has only

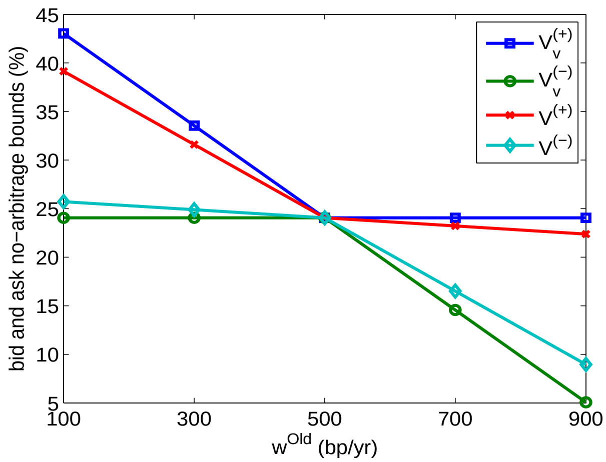

Figure 1: Plot of bid and ask no-arbitrage bounds for an illiquid CDS versus w_M^{Old}. $V^{(\sigma)}$ refers to the multi-CDS hedge, while $V_V^{(\sigma)}$ refers to the vanilla, single-CDS hedge. Note that the separation between the upper (ask) bound and the lower (bid) bound is greater for the vanilla hedge than it is for the multi-CDS hedge. The data used to obtain this plot was taken from the Standard Parameter List found at the beginning of Section 4.

non-negative payoffs. The hedging is only approximate in both cases however because the payoff stream of the illiquid CDS is only approximately replicated by the payoff stream of the appropriate hedging portfolio. Note carefully in this paragraph the use of the term 'hedging portfolio' to denote either the LUB or the GLB portfolio, as well as the use of the term 'hedged illiquid CDS position' to describe either the combination of a short illiquid CDS contract with a long LUB hedging portfolio, or the combination of a long illiquid CDS contract with a short GLB hedging portfolio.

The problem now is to value a given hedged postion. The hedged postion has many different possible payoff streams. A payoff stream is defined by giving the the realized default time τ and the realized recovery rate ρ for the payoff stream. The present value of the payoff stream characterized by a particular $\tau \leq T_M$ and a particular ρ, is called $\Delta(\tau, \rho)$. For $\tau > T_M$, the present value of the payoff stream is called Δ_0. The quantities $\Delta(\tau, \rho)$ and Δ_0 can be calculated from the information in Table 2 by summing up the present values of the premium payments made to the holder of the hedged position up to and including time τ and subtracting the present values of the loss payments made by the holder of the hedged position at τ, assuming a recovery rate of ρ, for all of the CDSs in the hedged position. Also, the initial bank deposit must be added. More explicitly, the present value of the payoff stream of the hedged postion σ, in the case that the default time $\tau = T_{i-1} + \delta$, $\delta \in (0, \delta_i]$ $(i = 1, 2, \ldots, N)$ and that the recovery rate is ρ is

$$\begin{aligned}\Delta_i^{(\sigma)}(\delta,\rho) &= (\sigma 1)\left[\sum_{j=1}^{N}\beta_j^{(\sigma)} + \sum_{j=1}^{i-1} g_{M,j}^{Old} - h_{M,i}^{Old}(\delta,\rho)\right] \\ &+ (\sigma 1)\left[\sum_{p=1}^{K}\alpha_{n(p)}^{(\sigma)}\left(-\sum_{j=1}^{i-1} g_{n(p),j} + h_{n(p),i}(\delta,\rho)\right)\right]. \end{aligned} \tag{7}$$

If no default occurs up to and including the time T_N, the present value of the payoff stream is

$$\Delta_0^{(\sigma)} = (\sigma 1)\sum_{i=1}^{N}\left(\beta_i^{(\sigma)} - \sum_{p=1}^{K}\alpha_{n(p)}^{(\sigma)} g_{n(p),i} + g_{M,i}^{Old}\right). \tag{8}$$

Eq. 7 can be used to define the function $\Delta^{(\sigma)}(\tau,\rho)$ (which differs from the function $\Delta_i^{(\sigma)}(\delta,\rho)$) depending on the default time, by

$$\Delta^{(\sigma)}(\tau,\rho) = \Delta_i^{(\sigma)}(\delta,\rho), \quad \tau = T_{i-1}+\delta, \quad \delta \in (0,\delta_i]; \quad i = 1,2,\ldots,N. \tag{9}$$

The function $\Delta^{(\sigma)}(\tau,\rho)$ can be shown to be a continuous function of τ at fixed ρ, except for possible discontinuities at the default time $\tau = 0$, and at the times of maturity T_M and $T_{n(p)}$, $p = 1,2,\ldots,K$ of the CDSs present in the hedged position (e.g. see Fig. 3 below).

The holder of the hedged position wishing to proceed to an estimate of its statistical properties must establish physical measures describing holder's views, as a result of detailed research on the question, of the default-time probability density and the recovery-rate probability density appropriate to the underlying reference name. The default characteristics will be assumed to be specified in terms of a probability density $\Upsilon(\tau)$ for the default time τ, such that $\Upsilon(\tau)d\tau$ is the probability that default occurs at some time in the interval $(\tau,\tau+d\tau]$. The quantity $\Upsilon(\tau)$ must be specified for $0 \le \tau \le T_N$, and Υ_0, the probability of survival until T_N, then follows from the condition that the sum of all probabilities must be unity. In the numerical examples discussed below, the default probabilities are assumed to be specified in terms of a constant hazard rate h, so that

$$\Upsilon(\tau) = h\exp(-h\tau),\ 0 \le \tau \le T_N \ \text{ and } \ \Upsilon_0 = \exp(-hT_N). \tag{10}$$

In this article, the hazard rate is given as

$$h = -log(1-PD_1); \tag{11}$$

where PD_1 is an estimated probability for default within the first year. Also, the probability density for the recovery rate given default, $\gamma(\rho)$, $0 \le \rho \le 1$, must be specified; by definition, the probability that the recovery rate is in the interval $(\rho,\rho+d\rho]$ is $\gamma(\rho)d\rho$. Fig. 2 shows three functions $\gamma_A(\rho)$, $\gamma_B(\rho)$, and

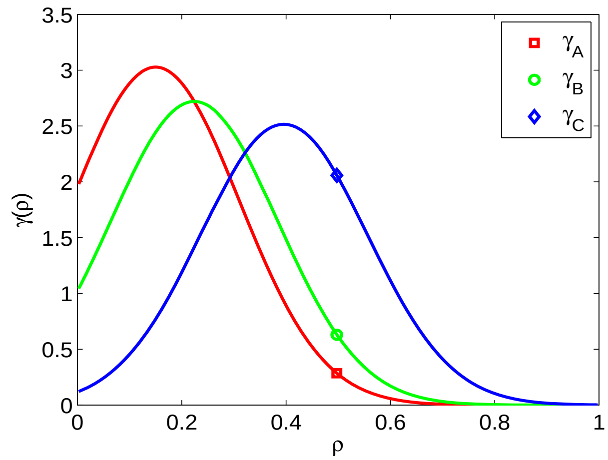

Figure 2: Plot of three different probability densities for the recovery rate ρ. The probability densities $\gamma_A(\rho)$, $\gamma_B(\rho)$ and $\gamma_C(\rho)$ are defined on the interval $0 \le \rho \le 1$, and on this interval are proportional to normal probability densities having means of 0.15, 0.224 and 0.396 respectivelly, with all having standard deviations of 0.16 . The mean values of ρ for the three probability densities are $\overline{\rho}_A = 20$ %, $\overline{\rho}_B = 25$ %, and $\overline{\rho}_C = 40$ %.

$\gamma_C(\rho)$, used in the numerical examples below. The probability density $\gamma_A(\rho)$ was chosen to have a mean recovery rate $\overline{\rho_A} = 20\%$, because the ISDA CDS Standard Model (ISDA, 2009) uses a fixed recovery rate of 20% for high yield names. The probability densities $\gamma_B(\rho)$ and $\gamma_C(\rho)$ are used below to examine the robustness of the bid and ask prices found using the procedure of this article, with respect to changes in the recovery-rate probability density.

The extension to the case of a time-dependent hazard rate, and a probability density for the recovery rate that is conditional on the default time is straight forward. A user having a forcast for the evolution of the future prospects, including credit worthiness, of the reference name could use such an extension to evaluate the impact of changes in the future prospects on the valuation process.

Given the probability densities described above, the mean value of the possible payoff streams for the hedged positions can be evaluated as

$$\overline{\Delta}^{(\sigma)} = \int_0^{T_N} d\tau \int_0^1 d\rho \Delta^{(\sigma)}(\tau, \rho) \Upsilon(\tau) \gamma(\rho) + \Upsilon_0 \Delta_0^{(\sigma)}. \tag{12}$$

Now consider the case of an investor who wishes to unwind a unit-notional short seaasoned CDS contract. The dealer contacted by the investor in this respect envisages taking over the short illiquid CDS contract[5], and hedging

[5]This transaction is equivalent to the dealer selling protection to the investor under the terms of the illiquid contract.

it with a corresponding long LUB portfolio purchased on the current market for $V^{(+)}$. The dealer's combined position (short illiquid CDS and long LUB portfolio) then has a non-negative payoff stream with an expected value of $\overline{\Delta}^{(+)}$. The dealer proposes to charge the investor the amount $V^{(+)}$ to pay for the hedge, and might also be expected to give back to the investor the fraction $\lambda^{(+)}$ of the positive expected payoff of her hedged position. The fraction $\lambda^{(+)}$ should be positive, otherwise the dealer would make a positive profit with no risk[4]. On the other hand, $\lambda^{(+)}$ should not be greater than unity, otherwise the dealer would have a positive net expected loss on the transaction. The net payment of the investor to the dealer for this transaction (i.e. the ask price) is $u^{Old(+)} = V^{(+)} - \lambda^{(+)}\overline{\Delta}^{(+)}$. A similar argument gives the amount $u^{Old(-)}$ that would be paid by the dealer to an investor (i.e. the bid price) for the investor's long protection illiquid CDS contracts. In summary

$$u^{Old(\sigma)}(\lambda^{(\sigma)}) = V^{(\sigma)} - (\sigma 1)\lambda^{(\sigma)}\overline{\Delta}^{(\sigma)}, \quad 0 < \lambda^{(\sigma)} < 1. \tag{13}$$

Bid and ask prices that are in the range defined by the restrictions on $\lambda^{(\sigma)}$ in Eq. 13 with are said to be potentially acceptable. A further restriction (see below) of the selection criterion will be necessary to identify the range of prices that are acceptable to the dealer, i.e. those that represent a good-deal for the dealer. The parameters $\lambda^{(\sigma)}$ determine the unwind values of the illiquid CDS contracts in question. Because it is not easy to gain intuition about the values of $\lambda^{(\sigma)}$ that would constitute a good deal for the dealer or the investor, these parameters will be replaced by the quantities $R_T^{(\sigma)}$ defined below.

The random variable giving the present value of the dealer's realized net loss on the transactions just described is, for a given value of $\lambda^{(\sigma)}$,

$$L^{(\sigma)}(\tau, \rho) = -\Delta^{(\sigma)}(\tau, \rho) + \lambda^{(\sigma)}\overline{\Delta}^{(\sigma)} \tag{14}$$

if default occurs at time $\tau \leq T_M$, and

$$L_0^{(\sigma)} = -\Delta_0^{(\sigma)} + \lambda^{(\sigma)}\overline{\Delta}^{(\sigma)} \tag{15}$$

if the default time is greater than T_M.

It is useful to have various quantitative measures of the risk incurred by a dealer who holds a hedged position. This risk will depend on the agreed upon price $u^{Old(\sigma)}(\lambda^{(\sigma)})$ (or, equivalently, the agreed upon value for $\lambda^{(\sigma)}$). The most detailed knowlege concerning the risk associated with the hedged position is contained in the knowledge of the probability density for random variable Δ, as given, for example in Figs. 4 and 5. Two other measures, which are useful because of their simplicity, are $L_{Max}^{(\sigma)}$, the maximum possible loss to the holder of the hedged position, and $E^{(\sigma)}(L|L > 0)$, the expected value of the loss, conditional on the loss being greater than zero. The maximum possible loss is

$$L_{Max}^{(\sigma)} = \lambda^{(\sigma)}\overline{\Delta}^{(\sigma)}. \tag{16}$$

Since the maximum possible loss is the amount of capital that the dealer must have available to cover losses in a worst case scenario, it is called the capital at risk for the purposes of this article. Note that, when $\lambda^{(\sigma)}$ is restricted to the range of potentially acceptable values indicated in Eq. 13, the capital at risk is positive. Because it has a straight-forward economic interpretation and because its use leads to simple analytic formula for quantities used in the process of establishing good-deal bid and ask prices, it has been chosen to be a key variable in future developments.

Because using any single parameter (such as the capital at risk) to characterize the risk of the hedged position is an oversimplification, a dealer evalating a potential transaction may also wish to examine the probability density for present values of the realized losses. This is easily constructed from the probability density for Δ, a numerical example of which is illustrated below in Fig. 4. Note, for example, in the left hand panel of the figure (which is relevant for the dealer taking over a short-protection illiquid position at an ask price) that there is about a 17% chance of Δ having the discrete value zero, and hence of the loss having its maximum possible value. In this case the capital at risk is a reasonable good single-parameter risk measure. For the case of the dealer taking over a long-protection illiquid CDS (illustrated in the right-hand panel of the figure) the probability of Δ being exactly zero is zero, but the probability density for Δ being zero is non-zero, so the probability of Δ having some value not too far from zero is significant, as is the probability of having a loss not too far from its maximum value. Finally, cases where the probability density for Δ is zero at $\Delta = 0$ are not excluded, and for these, the probability of having a loss close to the maximum is relatively small. Simply giving the value of the capital at risk does not distinguish between these three different cases. Although the dictates of simplicity require the use of the capital at risk as a central parameter in the approach used in this article, some adjustment in the expected return (defined and discussed below) required by the dealer could be made to try to account for the different cases discussed in this paragraph.

Note that pricing at the no-arbitrage bound is characterized by $\lambda^{(\sigma)} = 0$, and hence by the capital at risk $L_{Max}^{(\sigma)} = 0$, in agreement with the idea that there is no risk for a price at the no-arbitrage bound, since the payoffs to the holder of the hedged position are non-negative. This property of the capital at risk, that it tends to zero when the price for the illiquid CDS tends to that of the no-arbitrage bound, is a desirable property for a risk measure to have, and is found here for the case when hedging is carried out in terms of the porfolio that enforces the no-arbitrage bound. By way of contrast, the overall spread in the realized present values of the hedged position, as characterized, for example, by the variance of the random variable Δ, does not have this property.

Another relevant quantity is the dealer's expected profit, called $\overline{P_{FT}}^{(\sigma)}$. Since the realized profit is the negative of the realized loss, Eqs. 14 and 15 give

$$\overline{P_{FT}}^{(\sigma)} = (1 - \lambda^{(\sigma)})\overline{\Delta}^{(\sigma)}. \tag{17}$$

Note that the dealer's expected profit is positive when $\lambda^{(\sigma)}$ is restricted to the range of potentially acceptable values indicated in Eq. 13.

Given that the capital at risk, $L_{Max}^{(\sigma)}$, is positive, the dealer should have a reserve of capital from which this maximum possible loss could be covered, if necessary. The dealer is entitled to compensation for holding this capital and therefore might wish to calculate an expected return $R_T^{(\sigma)}$ on this capital from the relation

$$R_T^{(\sigma)} = \frac{\overline{P_{FT}}^{(\sigma)}}{L_{Max}^{(\sigma)}} = \frac{1-\lambda^{(\sigma)}}{\lambda^{(\sigma)}}; \text{ which implies } \lambda^{(\sigma)} = \frac{1}{1+R_T^{(\sigma)}}. \tag{18}$$

Note that the range of potentially acceptable values of $R_T^{(\sigma)}$ is given by

$$0 < R_T^{(\sigma)} < +\infty \tag{19}$$

The quantity $R_T^{(\sigma)}$, which is called the expected return on the capital at risk, is also called, for short, the expected return.

Eqs. 13 can now be rewritten in such a way that the bid and ask prices are determined in terms of the expected return, $R_T^{(\sigma)}$, rather than the parameter $\lambda^{(\sigma)}$. First define the bounds on the potentially acceptable bid and ask prices by

$$\begin{aligned} u_{min}^{Old(-)} &= V^{(-)}; \ u_{max}^{Old(-)} = V^{(-)} + \overline{\Delta}^{(-)}; \\ u_{min}^{Old(+)} &= V^{(+)} - \overline{\Delta}^{(+)}; \ u_{max}^{Old(+)} = V^{(+)}. \end{aligned} \tag{20}$$

With these definitions, Eqs. 13 can be rewritten as

$$\begin{aligned} u^{Old(-)}(R_T^{(-)}) &= u_{min}^{Old(-)} + (u_{max}^{Old(-)} - u_{min}^{Old(-)})\frac{1}{1+R_T^{(-)}}; \\ u^{Old(+)}(R_T^{(+)}) &= u_{min}^{Old(+)} + (u_{max}^{Old(+)} - u_{min}^{Old(+)})\frac{R_T^{(+)}}{1+R_T^{(+)}}. \end{aligned} \tag{21}$$

The solution of these equations for the expected return as a function of the bid or ask price is

$$\begin{aligned} R_T^{(-)}(u^{Old(-)}) &= \frac{u_{max}^{Old(-)} - u^{Old(-)}}{u^{Old(-)} - u_{min}^{Old(-)}}; \\ R_T^{(+)}(u^{Old(+)}) &= \frac{u^{Old(+)} - u_{min}^{Old(+)}}{u_{max}^{Old(+)} - u^{Old(+)}}. \end{aligned} \tag{22}$$

Note, from Eqs. 22, that when the ask price approaches its maximum potentially acceptable value, $u_{max}^{Old(+)}$ (which is the upper no-arbitrage bound) from below, the expected return tends to $+\infty$. This is because the capital at risk tends to zero at the upper no-arbitrage bound. Thus, the expected return would be considered to be excessive for an ask price close to its upper no-arbitrage bound. On the other hand, when the ask price approaches its minimum potentially acceptable value $u_{min}^{Old(+)}$ from above, the expected return tends to zero. For zero expected return, the dealer will nevertheless have a positive capital at risk equal to $L^{(+)} = \overline{\Delta}^{(+)}$. However, the dealer would expect to earn a positive return on any capital at risk. To accomplish this the dealer can decide to accept only those deals that would give an expected return higher than a subjectively chosen positive good-deal lower bound (for the expected return). This good-deal lower bound on the expected return gives, from Eqs. 21, a good deal lower-bound on the ask price. The dealer would require the investor to pay more than the good-deal ask-price lower bound in order for the dealer to take over the investor's short-protection position. In so far as the dealer is concerned, there is of course no upper bound on the ask price. If an investor is happy to pay more than the no-arbitrage upper bound of $u_{max}^{Old(+)}$, the dealer will be happy to accept the arbitrage profit. The investor, however, may wish to establish a good-deal upper bound on the ask price. Similarly, the dealer will establish a good deal upper-bound on the bid price: the dealer will buy a long-protection contract from an investor only if the bid price is less that this upper bound. The establishment by the dealer of lower and upper good-deal bounds for the ask and bid prices, respectively, is further illustrated in Section 8 and Fig. 7.

Another possible choice for a quantity for which high values would signal a good deal is the Sharpe ratio. This ratio has been consider for good-deal indentification in Carr et al. (2001); Černỳ and Hodges (2002); Cochrane and Saà-Requejo (2000). For the purposes of this article, an analogue of the Sharpe ratio, called the effective Sharpe ratio, is defined as

$$S_R^{(\sigma)} = \frac{R_T^{(\sigma)}}{L_{Max}^{(\sigma)}}. \qquad (23)$$

The use of the effective Sharpe ratio to establish bid and ask good-deal bounds is illustrated in Fig. 9.

6 Numerical Examples for $\Delta^{(\sigma)}(\tau, \rho)$ and $(\Gamma(\Delta)\ , \Upsilon_0)$

The realized present value functions $\Delta^{(\sigma)}(\tau, \rho)$ for the hedged illiquid CDS positions are plotted in Fig. 3 for the case where the LUB and GLB hedging portfolios are described in Table 2. For a fixed default time τ, $\Delta^{(\sigma)}(\tau, \rho)$ varies linearly with ρ for $\rho \in [0, 1]$. Also, $\Delta^{(\sigma)}(\tau, \rho)$ is a continuous function of τ and ρ in each region $\tau \in (j-1, j]$ and $\rho \in [0, 1]$, for $j = 1, 2, \ldots, 5$.

A knowledge of $\Delta^{(\sigma)}(\tau, \rho)$ and of the probability densities $(\Upsilon(\tau),\ \Upsilon_0)$ and $\gamma(\rho)$ allows the probability density $(\Gamma^{(\sigma)}(\Delta),\ \Upsilon_0)$ to be determined numerically. Here, $\Gamma^{(\sigma)}(\Delta) d\Delta$ is the probability that the realized present value Δ of the hedged position lies in the interval $(\Delta, \Delta + d\Delta]$ and Υ_0 is the probability that the value Δ_0 occurs. The probability density $(\Upsilon(\tau),\ \Upsilon_0)$ is given by Eq. 10 with the hazard rate h given in terms of the probability of default within 1 year, PD_1, by Eq. 11. The quantities $\Gamma^{(\sigma)}(\Delta)$ and Υ_0 are shown in Fig. 4 for

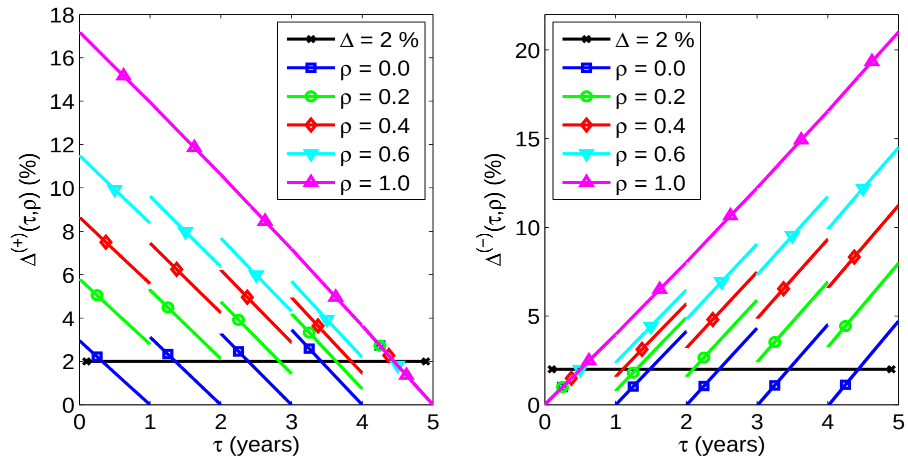

Figure 3: Plots of $\Delta^{(\sigma)}(\tau, \rho)$, the realized present value of the payoff stream of the hedged illiquid CDS when default occurs at time τ and the recovery rate in ρ, plotted versus τ for a number of values of ρ. The parameter values given in the caption to Fig. 1 were used to obtain this figure, as was the data presented in Table 2.

the case where the input parameters are taken from the Standard Parameter List shown at the beginning of Section 4, and where $\Delta^{(\sigma)}(\tau, \rho)$ is as shown in Fig. 3. For comparison purposes, the same quantities are shown plotted in Fig. 5 for the case where the hedging has been carried out in terms of the vanilla hedging portfolio. Note that when hedging is carried out with the vanilla hedging portfolio, the values of realized present value Δ of the hedged position which occur with significant probability are spread out over a range of nearly 20 %, whereas when a multi-CDS hedge is used, these same values are spread out over a range of only about 10 %, although in the GLB case, there is a significant contribution from Δ_0 at a Δ value of somewhat more than 20 %; this last contribution is not a concern from the point of view of the downside risk to the holder of the hedged position, however, since it represents a positive profit for this party.

It is of interest to develop a qualitative understanding of the discontinuities in $\Gamma^{(\sigma)}(\Delta)$ which occur in Figs. 4 and Fig. 5. For this purpose, it will be sufficient to focus on understanding how the form of $\Delta^{(-)}(\tau, \rho)$ exhibited in Fig. 3 produces the discontinuities observed in $\Gamma^{(-)}(\Delta)$ in Fig. 4.

The probability density $\Gamma^{(-)}(\Delta)$ will be computed as

$$\Gamma^{(-)}(\Delta) = \sum_{j=1}^{T_M} \Gamma_j^{(-)}(\Delta) \tag{24}$$

where $\Gamma_j^{(-)}(\Delta)$ is the contribution to $\Gamma^{(-)}(\Delta)$ from $\Delta^{(-)}(\tau, \rho)$'s for which $\tau \in$

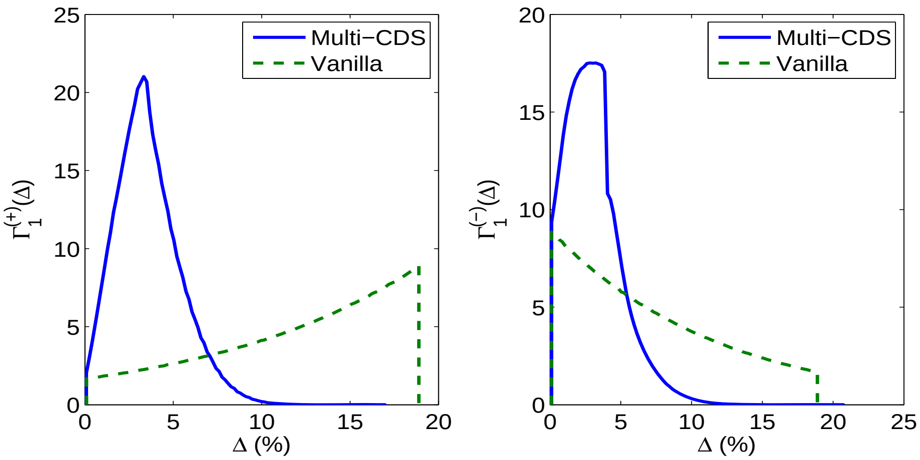

Figure 4: Plot of $\Gamma(\Delta)$ and Υ_0 for the case where a illiquid CDS of maturity $T_M = 5$ years and spread $w_M^{Old} = 100$ bp/year is hedged by the multi-CDS LUB and GLB hedging portfolios described in Table 2. The quantities DeltaBar appearing in the legends stand for $\overline{\Delta}^{(+)}$ and $\overline{\Delta}^{(-)}$. For all numerical examples in this article, the probability density ($\Upsilon(\tau)$, Υ_0) is given by Eq. 10 with the hazard rate h given by Eq. 11 and the probability of default within 1 year taken to be $PD_1 = 0.3$. Also, the recovery rate density γ_A of Fig 2 was used.

$(j-1, j]$ years, i.e. for which default occurs in the j-th year from the inception of the hedged long illiquid CDS position. Note first, from Table 2, that the total notional of the GLB hedging portfolio is $\alpha_{total} = 1.000$. This means that, if default occurs at any time during the first year after the inception of the hedged position, then the loss due to default of the GLB hedging portfolio exactly cancels the loss due to default of illiquid long CDS position: there is thus no dependence of $\Delta^{(-)}(\tau, \rho)$ on ρ during this first year, as shown in Fig. 3, and $\Delta^{(-)}(\tau, \rho)$ can be written as $\Delta^{(-)}(\tau)$, dependent only on τ, during this first year. Thus, during this first year, one can also solve to find $\tau = \tau(\Delta)$. One can now find

$$\Gamma_1^{(-)}(\Delta) = \Theta(\Delta) \left[\frac{\Upsilon(\tau)}{d\Delta/d\tau} \right]_{\tau=\tau(\Delta)}, \tag{25}$$

where $\Theta(\Delta) = 1$ for $\Delta \in (0, \Delta^{(-)}(\tau = 1)]$ and zero otherwise. The quantity $\Gamma_1^{(-)}(\Delta)$ thus increases discontinuously from zero at $\Delta = 0$ and decreases discontinuously at $\Delta = \Delta^{(-)}(\tau = 1)$. This accounts for the discontinuities observed in $\Gamma^{(-)}(\Delta)$ at $\Delta = 0$ and at $\Delta \approx 450\%$ in Fig. 4.

To find the contributions $\Gamma_j^{(-)}(\Delta)$, $j = 2, 3, 4, 5$, consider first the value $\Delta = 2\%$ indicated by the horizontal line in the right hand panel of Fig. 3. Note that for a given τ, the minimum non-zero value of $\Delta^{(-)}(\tau, \rho)$ is $\Delta_{min}^{(-)}(\tau) = \Delta^{(-)}(\tau, \rho = 0)$ and define $\tau_j(\Delta)$ to be the solution of $\Delta = \Delta_{min}^{(-)}(\tau_j)$ for τ_j

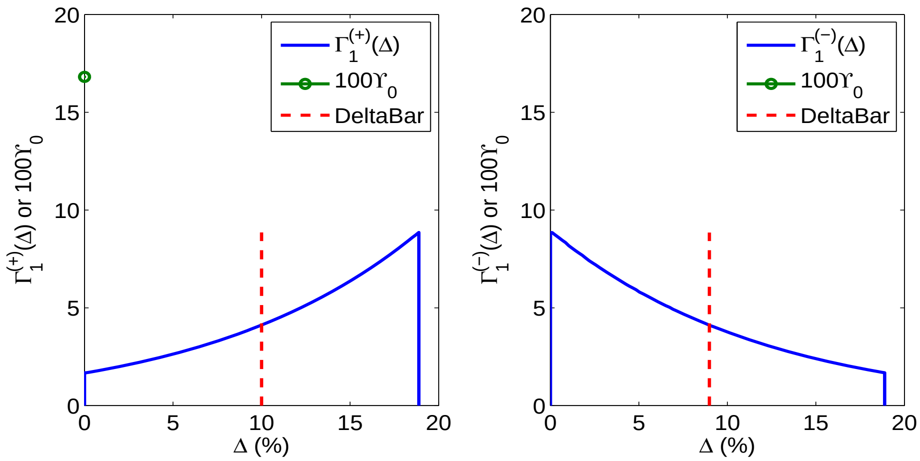

Figure 5: Plot of $\Gamma(\Delta)$ and Υ_0. This is similar to Fig. 4 except that the illiquid CDS of maturity $T_M = 5$ years and spread $w_M^{Old} = 100$ bp/year is hedged by the single-CDS vanilla hedging portfolio.

(keeping $\Delta = 2\%$). Then the value of $\Gamma_j^{(-)}(\Delta)$ is

$$\Gamma_j^{(-)}(\Delta) = \int_{j-1}^{\tau_j(\Delta)} d\tau \left[\frac{\Upsilon(\tau)\gamma(\rho)}{F(\tau)} \right]_{\rho=\rho(\tau,\Delta)} \tag{26}$$

where $F(\tau) = (\partial \Delta^{(-)}(\tau, \rho)/\partial \rho)_\tau$ (which is independent of ρ since $\Delta^{(\sigma)}(\tau, \rho)$ is linear in ρ for $\rho \in [0, 1]$). Also, the function $\rho(\tau, \Delta)$ is the solution of $\Delta = \Delta^{(-)}(\tau, \rho)$ for ρ (which is easily found due to the linearity of $\Delta^{(-)}(\tau, \rho)$ in ρ).

Now note from Fig. 3 that as Δ decreases from 2 % to zero, $\tau_j(\Delta)$ decreases continuously from $\tau_j(\Delta = 2\ \%)$ to $j - 1$, for $j = 2, 3, 4, 5$; thus $\Gamma_j^{(-)}(\Delta)$ decreases continuously to zero as Δ decreases to zero. Similarly, it is easily seen that $\Gamma_j^{(-)}(\Delta)$, $j = 2, 3, 4, 5$ decreases continuously to zero as the largest value of Δ making contributions to $\Gamma_j^{(-)}(\Delta)$ is approached from below. Thus, all contributions of $\Gamma_j^{(-)}(\Delta)$, $j = 2, 3, 4, 5$ to $\Gamma^{(-)}(\Delta)$ are continuous in Δ.

Similar arguments to those just given explain all of the discontinuities in $\Gamma^{(\sigma)}(\Delta)$ observed in Figs. 4 and 5. (As a result of the discretization procedure in the numerical evaluation, the discontinuities in Fig. 4 in $\Gamma^{(+)}(\Delta)$ at $\Delta \approx 4\%$ (too small to be observed), and in $\Gamma^{(-)}(\Delta)$ observed at $\Delta \approx 4.5\%$ are not as abruptly vertical as one would expect if calculated with perfect precision.)

7 Multi-CDS Versus Vanilla Hedges

Consider again the case of an investor who wishes to unwind an illiquid CDS position by selling it to a dealer. Because the investor will no longer hold

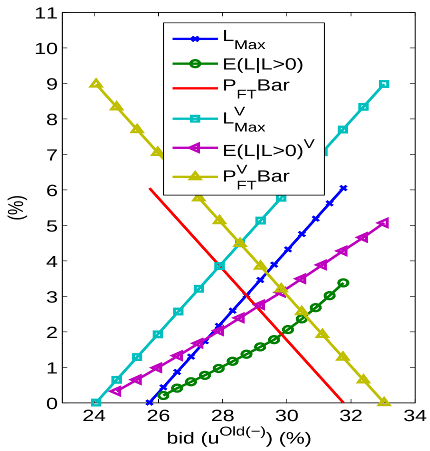

Figure 6: Plot of two measures of risk to the dealer, (L_{Max} and $E(L|L>0)$), and the dealer's expected profit ($P_{FT}Bar = \overline{P_{FT}}$) for the case (a) (right panel) where a short illiquid CDS of maturity $T_M = 5$ years and spread $w_M^{Old} = 100$ bp/year is hedged by the long vanilla (superscript V in the legend) LUB portfolio, and by the long multi-CDS (no supercript V) LUB portfolio, and case (b) (left panel) where a long illiquid CDS of maturity $T_M = 5$ years and spread $w_M^{Old} = 100$ bp/year is hedged by the short vanilla (superscript V in the legend) GLB portfolio, and by the short multi-CDS GLB portfolio (no supercript V). The data is plotted for the entire range of potentially acceptable values of $u^{Old(+)}$ and $u^{Old(-)}$, defined by restricting λ to lie in $(0,1)$ in Eq. 13. These ranges are: for $u^{Old(+)}$, $33.03\% < u^{Old(+)} < 43.03\%$ for the vanilla hedge and $36.16\% < u^{Old(+)} < 39.13\%$ for the multi-CDS hedge; and, for $u^{Old(-)}$, $24.06\% < u^{Old(-)} < 33.03\%$ for the vanilla hedge, and $25.72\% < u^{Old(-)} < 31.77\%$ for the multi-CDS hedge.

the illiquid CDS after the sale, the only thing that matters to the investor is the negotiated price, $u^{Old(\sigma)}$. On the other hand, after the sale and the establishment of a hedge by the dealer, the dealer will hold a risky position. Thus the dealer will be interested not only in the selling price, but in the degree of risk associated with the hedged position. This section looks at a number measures that can be used to assess the risk and profitability of the hedged position of a dealer who takes over an illiquid CDS. The question of how to determine a preferred hedge is resolved more concisely in Section 8.

The left panel of Fig. 6 is constructed from the point of view of a dealer who has agreed to pay an investor the bid price of $u^{Old(-)}$ for the investor's long illiquid CDS contract, and has then hedged this contract with a short GLB hedging portfolio. The right panel takes the point of view of an investor who has agreed to pay a dealer the ask price of $u^{Old(+)}$ to take over the investor's short illiquid CDS contract, which the dealer then hedges with a long LUB hedging portfolio.

Consider the question of choosing a value for the ask price $u^{Old(+)}$ (see right panel of the figure). There are potentially acceptable ask prices for both the vanilla and the multi-CDS hedges when $36.16\% < u^{Old(+)} < 39.13\%$. In this range, the expected profit and the risk measures are all much greater when the vanilla hedge is used than when the multi-CDS hedge is used. Since a large expected profit is desirable, but a large value for the risk measure is undesirable, the plots of Fig. 6 do not give a definitive view on which hedge is preferred (for this, see the following section). For ask prices greater than 39.13%, the dealer who uses the multi-CDS hedge will make a risk-free profit[4], so the vanilla hedge will definitely not be used in this case. For ask prices less than 36.16%, the dealer using the multi-CDS hedge will make an expected loss. However, for $33.03\% < u^{Old(+)} < 36.16\%$, the vanilla hedge makes a positive expected profit, but at the cost of a realtively high capital at risk. Similar considerations apply to the choice of the bid price $u^{Old(-)}$.

In general, it is expected that $E(L|L > 0) < L_{Max}$ Note, however, that in the right panel of the figure, that, as functions of the ask price, both L_{Max} and $E(L|L > 0)$ tend to zero at the value of the ask price corresponding to the upper no-arbitrage bound, and with the same value of the limiting slope. This is true for both the vanilla hedge and the multi-CDS hedge, and can be shown to be a consequence of the fact there is a non-zero probability Υ_0 of having the discrete loss of $\Delta = \Delta_0 = 0$ (see Figs. 4 and 5. In the left panel, it it is also the case that both L_{Max} and $E(L|L > 0)$ tend to zero at the same value of the price (in this case, the bid price), but with different limiting slopes. The fact that the limiting slopes are different in this case can be shown to be a consequence of the fact there is no discrete probability of having $\Delta = 0$.

The principal overall impression left by Fig. 6 is that both the risks and the expected profits tend to be larger when hedging is carried out with the vanilla hedge, that when hedging is carried out with the multi-CDS hedge.

8 Establishing Good-Deal Bounds for Bid and Ask Prices

8.1 Dealer Sets Expected Return

Consider a dealer who takes over an illiquid CDS contract from an investor and hedges this contract by buying the appropriate hedging portfolio on the current market as described above. Since the hedging is not perfect, there is the possibility that the dealer will realize a net loss on the transaction. The present value of the maximum possible loss was called above the capital at risk. The dealer will also make an expected profit on the transaction, and thus an expected return on the capital at risk (called, for short, the expected return) defined by Eq. 18. Section 5 described how the dealer could arrive at a good-deal lower bound on the ask price and a good-deal upper bound on the bid price in terms of a good deal lower bound on the expected return. Fig. 7 shows graphically the

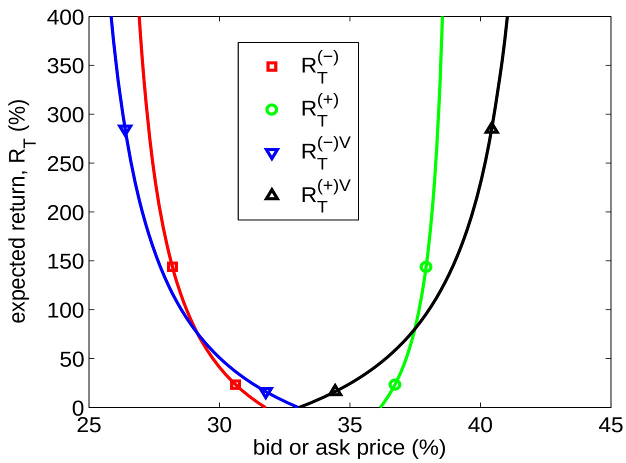

Figure 7: Plot of the expected return on capital at risk, $R_T^{(-)}$ (or $R_T^{(+)}$), versus the bid (or ask) price. A superscript V in the legend indicates results for the vanilla hedge; otherwise, the results are for the multi-CDS hedge. For the multi-CDS hedge, for bid prices greater than the maximum potentially acceptable bid price of 31.77 %, and for ask prices less than the minimum potentially acceptable ask price of 36.06 %, the expected profit and the expected return on capital at risk are both negative. For the vanilla hedge, the maximum potentially acceptable bid price and the minimum potentially acceptable ask price are both equal to 33.03%.

relationship between the expected return, and the bid and ask prices; and thus also the relationship between the good-deal bounds on the expected return, and the good-deal bounds on the bid and ask prices.

Note from the figure, that for the multi-CDS hedge, a lower bound on the potentially acceptable range for the bid-ask spread is 4.29 %. It is also clear for the case of the multi-CDS hedge that, if the bid and ask prices were equal, as in Eq. 1, the dealer would have a negative expected return for either the bid price, or the ask price, or both.

Note also in Fig. 7 that, for the vanilla hedge, the bid-ask spread is zero when the required expected return is zero. There is, however, a significant capital at risk when the expected return is zero, so the dealer who concludes a deal at zero expected return will be in the undesirable position of receiving zero expected return in compensation for the risk taken on. Thus, a non-zero bid-ask spread is expected in the case of the vanilla hedge also. It is of interest to contrast this result with the result of the complete-market risk-neutral approach to valuation, which also includes only a single CDS in the hedge, but which gives equal bid and ask prices (as, for example in Eq. 1).

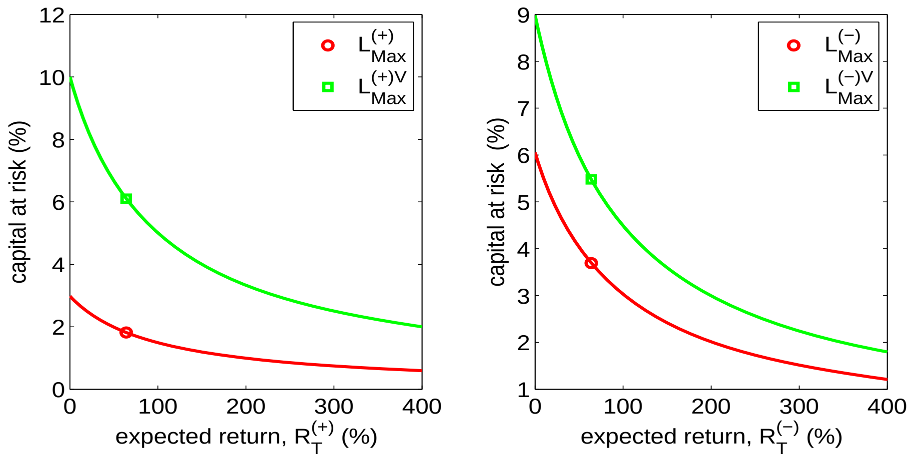

Figure 8: Plot of the capital at risk versus the expected return for both the vanilla hedge (superscript V in the legends) and the multi-CDS hedge. The left (right) hand panel is for the case where the dealer acquires an investor's short (long) illiquid position.

8.2 Choosing Between Two Hedges

This subsection assumes initially that, when the dealer who takes over an illiquid CDS has a choice of two or more hedges, the dealer will choose the hedge that minimizes the amount of capital at risk, assuming that a definite value for the expected return has been fixed, independently of the consequences that this might have for the bid and ask prices. Making this choice on the basis of the effective Sharpe ratio is discussed at the end of this subsection.

Fig. 8 shows the capital at risk plotted against the expected return for both the multi-CDS and the vanilla hedges for input parameters taken from the Standard Parameter List printed at the beginning of Section 4. Clearly, for any fixed value of the expected return, the capital at risk is significantly reduced by choosing the multi-CDS hedge. Also note that the capital at risk can be written in the form

$$L_{Max}^{(\sigma)} = \overline{\Delta}^{(\sigma))}/(1 + R_T^{(\sigma)}). \tag{27}$$

Figs. 4 and 5 show that the value of the quantity $\overline{\Delta}^{(\sigma)}$ for the vanilla hedge is significantly larger than the value of the corresponding quantity for the multi-CDS hedge, thus confirming the conclusion reached from looking at Fig. 8, that at any given value for the expected return, the multi-CDS hedge is preferred.

8.3 The Effective Sharpe Ratio as a Good-Deal Criterion

Above it was shown how the establishment of a good-deal lower bound for the expected return gave good-deal bounds for the bid and ask prices. The

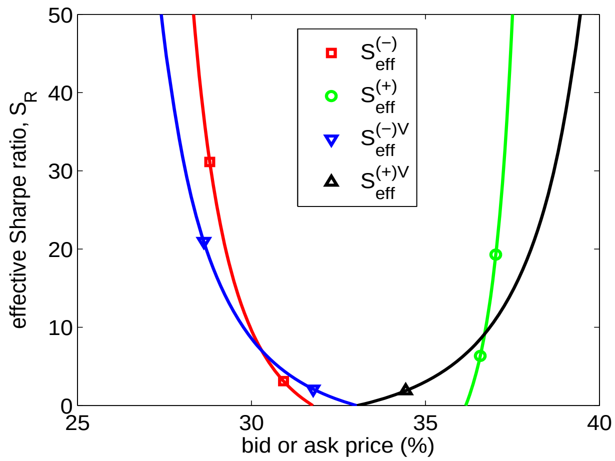

Figure 9: This figure is similar to Fig. 7 except that the expected return is replaced by the effective Sharpe ration, S_R, as the quantity that is selected by the dealer to define the good-deal lower bound.

expected return is a quantity for which large values mean good deals. Similarly, the effective Sharpe ratio, defined by Eq. 23, is a quantity for which large values indicate good deals. The choice of the effective Sharpe ratio over the expected return adds weight to deals which have a smaller capital at risk for a given expected return.

To find the dependence of the capital at risk on $\overline{\Delta}^{(\sigma)}$ for a given effective Sharpe ratio, begin by eliminating $R_T^{(\sigma)}$ from Eqs. 23 and 27, which gives $S_R(L_{Max}^{(\sigma)})^2 + L_{Max}^{(\sigma)} - \overline{\Delta}^{(\sigma)} = 0$. The positive root of this equation is

$$L_{Max}^{(\sigma)} = \frac{\sqrt{1 + 4S_R\overline{\Delta}^{(\sigma)}} - 1}{2S_R}, \tag{28}$$

which shows that, at fixed S_R, L_{Max} is a monotonically increasing function of $\overline{\Delta}^{(\sigma)}$. Thus, for a given effective Sharpe ratio, the preferred hedge is the one with the lowest value of $\overline{\Delta}^{(\sigma)}$, and, as in the previous section, the multi-CDS hedge enforcing the no-arbitrage bounds is preferred relative to the vanilla hedge.

Fig. 9 shows the effective Sharpe ratio plotted as a function of the bid and ask prices for the same set of input parameters as was used to obtain Fig. 7. The effective Sharpe ratio can be used in the same way as the expected return to determine good-deal bounds for the bid and ask prices.

9 Various Optimal Hedges

This section describes the construction of a number of different optimal hedges. The choice of the particular optimal hedge to use depends on the objectives of the user.

9.1 Minimizing the Capital at Risk

The preceding section has argued that, from the dealer's point of view, a good criterion for selecting between hedges is to choose the hedge that, for a given expected return, results in the lowest value of the capital at risk. On this basis, the multi-CDS hedge that enforces the no-arbitrage bounds on the value of the illiquid CDS is preferred to the vanilla hedge. However, only one multi-CDS hedge was considered, and the question thus arises as to whether or not a different multi-CDS hedge might give a lower value for the capital at risk. This subsection shows how to find the hedge that gives the lowest capital at risk for a given expected return.

Subsection 8.2 shows that finding the lowest capital at risk for a given expected return is the same as finding the lowest value for $\overline{\Delta}^{(\sigma)}$ for a given expected return. The quantity $\overline{\Delta}^{(\sigma)}$ can be written in the form

$$\overline{\Delta}^{(\sigma)} = (\sigma 1) \left[D_c + \sum_{i=1}^{N} \beta_i + \sum_{p=1}^{K} D_u(p) \alpha_{n(p)}, \right] \tag{29}$$

where D_c and the $D_u(p)$'s are known constants (which depend on the input parameters specified in the Standard Parameter List given at the beginning of Section 4, but not on the β_i's or the $\alpha_{n(p)}$'s). The procedure for determining the optimal hedge is thus to minimize $\overline{\Delta}^{(\sigma)}$ with respect to the β_i's and $\alpha_{n(p)}$'s, subject to the contraints of inequalities 5 and 6. These constraints are used here, as well as above for the no-arbitrage bounds calculation, so that the realized present values of the possible payoff streams of the dealer's hedged position will be non-negative. Thus, the feasible region for the present problem is the same as the feasible region for the problem of obtaining the no arbitrage bounds. Note also that no new constraint has to be added to account for the fact that the expected return should be kept constant during the optimization process, since, by Eq. 18, the expected return does not depend on the parameters with respect to which one optimizes.

Optimizing the objective function $\overline{\Delta}^{(\sigma)}$ of Eq. 29 to obtain the optimal hedge is very similar to optimizing the objective function of Eq. 3 to determine the no-arbibrage bounds and corresponding hedges, the only difference being that the upfronts $u_{n(p)}$ of Eq. 3 are now replaced by the parmeters $D_u(p)$ of Eq. 29. (The constant D_c in Eq. 29 does not affect the optimization.) The numerical values of the parameters $D_u(p)$ when the input parameter values given in the Standard Parameter List below are shown in Table 3. In spite of the significant difference between the values of the $D_u(p)$'s of Table 3 and the $u_{n(p)}$'s given

p	1	2	3	4	5
$D_u(p)$ (%)	19.56	32.98	42.19	48.51	52.84

Table 3: The parameters $D_u(p)$, $p = 1, 2, ..., K = 5$, evaluated for the values of the input parameters given in the Standard Parameter List. These parameters are given as percentages to facilitate comparison with the upfronts given in Table 1.

in Table 1, the two optimization procedures give the the same hedges (given in Table 2), and the same values for $\overline{\Delta}^{(\sigma)}$. (The specification of a hedge requires that the initial bank deposit β_{total} and the notionals $\alpha_{n(p)}$ of the market CDSs making up the hedge be specified.)

Since, in the numerical example under discussion, the optimal hedge described in this subsection is the same as the hedge that enforces the no-arbitrage bounds, the calculations for all quantities given from Sections 4 to 8 give the same results. Thus, for input parameters taken from the Standard Parameter List (e.g. $R_T = 25\%$ and $\overline{\rho} = 20\%$), the lowest values for the capital at risk are, 2.4% obtained when hedging the short illiquid CDS and resulting in an ask price of 36.74%, and, 4.8% obtained when hedging the long illiquid CDS and resulting in an bid price of 30.56%.

The quantity $\Delta^{(\sigma)}(\tau, \rho)$ depends linearly on the recovery rate ρ. Therefore, the quantity $\overline{\Delta}^{(\sigma)}$ depends linearly on the expected recovery rate $\overline{\rho}$. This dependence shows up in a reasonably strong dependence of the $D_u(p)$'s on $\overline{\rho}$. The calculation just described assumes $\overline{\rho} = 0.2$ (as in the Standard Parameter List). In order to try to find a situation in which the optimal hedge is distinct from the multi-CDS hedges that enforce the no-arbitrage bounds, the optimal hedge is obtained for all values of $\overline{\rho}$ such that $0 < \overline{\rho} < 1$. For the hedges where the dealer is selling protection (at an ask price) the optimal hedge is always the same as the multi-CDS hedge enforcing the no-arbitrage bounds, independently of the value chosen for $\overline{\rho}$. However, for the case where the dealer is buying protection (at a bid price) the two hedges are the same for $\overline{\rho} < 85.98\%$, and different for $\overline{\rho} > 85.99\%$. The changeover one regime to the other is association with a simultaneous change of sign of all $D_u(p)$'s. The optimal hedges found for $\overline{\rho} > 85.99\%$ can not be used, however, since all of the bid prices that can found using this hedge lie below the lower no-arbitrage bound.

9.2 Minimizing the Capital at Risk with Extra Conditions

The capital at risk was minimized in the previous section without regard to the consequences for the ask (or bid) prices. It might occur, for example, that when the dealer informs the investor of the ask price that was obtained in this calculation, that the investor will argue for a lower ask price. A dealer who wants the investor's business may wish to consider how much her capital at risk will rise if she sells at a lower ask price, but maintains the same value for the expected return. The bid and ask prices are given in Eq. 13. By making use of

Eq. 29, the bid and ask prices can be written in the form

$$u^{Old(\sigma)} = -\lambda D_c + (1-\lambda)\sum_{i=1}^{N}\beta_i^{(\sigma)} + \sum_{p=1}^{K}(u_{n(p)} - \lambda D_u(p))\alpha_{n(p)}^{(\sigma)}. \qquad (30)$$

Now suppose that one wishes to impose the constraint that the bid price is greater than or equal to the given value $u_0^{Old(-)}$, or that the ask price is less than or equal to a given value $u_0^{Old(+)}$, i.e. that

$$(\sigma 1)u^{Old(\sigma)} \leq (\sigma 1)u_0^{Old(\sigma)} \qquad (31)$$

This is easily done by adding the constraint of inequality 31 to the list of constraints used in the minimization of $\overline{\Delta}^{(\sigma)}$ in the previous section.

Fig. 10 shows results obtained using input parameters from the Standard Parameter List for the pricing of a deal where an investor asks a dealer to take over an illiquid CDS. The dealer initially prices the deal by setting a required expected return of $R_T = 25\%$ and then finding the hedge that gives the minimum capital at risk, with no constraint on the possible values for the bid or ask prices. The results for the minimum capital at risk and the corresponding ask price (investor holds a short-protection position) are given by point A in the figure, and the corresponding results for the bid price (investor holds a long protection position) are given by point B. If the investor protests that the ask price is too high, the dealer can calculate the increase in the capital at risk resulting from a given decrease in the ask price; the results of this are shown on the ask-price line. Similar considerations hold for the bid price. The dealer can then decide how much, if any, she is willing to reduce the ask price (or increase the bid price).

This section has described two different optimal hedges useful in obtaining bid and ask prices for an illiquid CDS. The first hedge described takes into account the buyer's (i.e. dealer's) preferences only (to have a given expected return and to minimize capital at risk), while the second hedge describes a way of adding the seller's (i.e. investor's) preferences (to have a given selling price) at the expense of increasing the buyer's capital at risk. Other types of hedges (described in terms of different objective functions and different extra constraints) can also be treated using similar methods.

Fig. 10 was obtained assuming that the dealer's minimum value of expected return required for a good deal is $R_T = 25\%$. If the dealer had decided to use the effective Sharpe ratio (defined in Eq. 23) as the measure for a good deal, similar results would have to be obtained at constant values of the Sharpe ratio. It is more difficult to obtain results at constant Sharpe ratio directly because the added constraint (corresponding to inequality 31 with Eq. 30) in this case is a non-linear function of the variables with respect to which the optimization is carried out. Note, however, that each point on a line in Fig. 10 corresponds to some effective Sharpe ratio (given by Eq. 23). By carrying out the calculation of minimizing the capital at risk for a given value of the bid or ask price $u_0^{Old(\sigma)}$

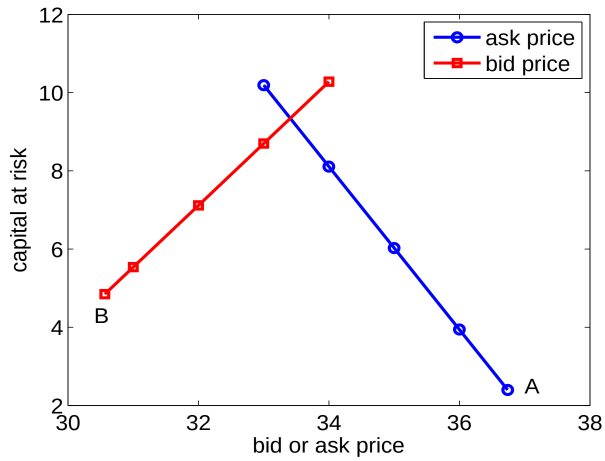

Figure 10: The plot shows the minimum capital at risk conditional on the dealer's expected return being $R_T = 25\%$, and the bid or ask price being as indicated by the horizontal coordinate.

and expected return R_T ($\lambda = 1/(1 + R_T)$) in constraint 31, and doing this for a range of values of R_T, one can establish the minimum capital at risk for a given $u_0^{Old(\sigma)}$ for a range of values of the effective Sharpe ratio. If the desired effective Sharpe ratio lies within the range of effective Sharpe ratios that is found by this calculation, then the minimum capital at risk for this effective Sharpe ratio and the given bid or ask price $u_0^{Old(\sigma)}$ has been found.

10 Scaling in a Special Case

The procedure described above for finding the GLB and LUB hedging portfolios used to hedge a illiquid CDS of maturity T_M and spread w_M^{Old} was to carry out the optimization operation (4) subject to the constraints (5) and (6). This procedure is applicable whether of not there is a CDS on the current market with either its maturity equal to T_M or its spread equal to w_{Old} or both. The scaling results derived in this section hold only in the special case that there is a CDS on the current market with the same maturity T_M as that of the illiquid CDS that is to be hedged, but which has a different spread w_M from that w_M^{Old} of the illiquid CDS. In this case, the quantities of interest (e.g. the values of the hedging portfolios $V^{(\sigma)}$, the probability density $\Gamma^{(\sigma)}(\Delta)$ for the realized present value Δ of the payoff stream of the hedged position, and the upfront values $u^{(\sigma)}(\lambda)$ associated with illiquid CDS) have well-defined scaling behaviours with respect to changes in scale of the variable $w_M - w_M^{Old}$. In the case that there is no CDS on the current market with the same maturity as that of the illiquid CDS to be hedged, there is obviously no such scaling behavior since there is no market determined quantity w_M. Also, the scaling relations do not hold when the constraint 31 is included.

Note that the optimization indicated in (4) is carried out with respect to the variables β_i, $i = 1, \ldots, N$ and $\alpha_{n(p)}$, $p = 1, \ldots, K$. To demonstrate the scaling behavior, the variables β_i and $\alpha_{n(p)}$ in the relations (4), (5) and (6) are replaced by the new variables

$$\begin{aligned}
\beta'_i &= \frac{\beta_i}{w_M - w_M^{Old}}, \quad i = 1, \ldots, N, \\
\alpha'_M &= \frac{\alpha_M - 1}{w_M - w_M^{Old}}, \\
\alpha'_{n(p)} &= \frac{\alpha_{n(p)}}{w_M - w_M^{Old}}, \quad p = 1, \ldots, K, \quad \text{except } n(p) \neq M.
\end{aligned} \tag{32}$$

When the optimization is carried out in terms of the new primed variables, it is found that these primed variables are determined as functions of the spreads $w_{n(p)}$ and upfronts $u_{n(p)}$, $p = 1, \ldots, K$, but are independent of the variable $w_M - w_M^{Old}$.

By using (32) to replace β_i, $\alpha_{n(p)}$ in (3) by β'_i and $\alpha'_{n(p)}$, one finds that $V^{(\sigma)}$ has the form

$$V^{(\sigma)}(w_M - w_M^{Old}) = u_M + (w_M - w_M^{Old}) F_V^{(\sigma)}, \tag{33}$$

for $(w_M - w_M^{Old}) > 0$, and

$$V^{(\sigma)}(w_M - w_M^{Old}) = u_M + (w_M - w_M^{Old}) F_V^{(\overline{\sigma})}, \tag{34}$$

for $(w_M - w_M^{Old}) < 0$. where all dependence on the variable $(w_M - w_M^{Old})$ is shown explicitly, and the quantity $F_V^{(\sigma)}$ has an implicit dependence on the variables $w_{n(p)}$ and $u_{n(p)}$, $p = 1, \ldots, K$, describing the current CDS market. The dependence on $(w_M - w_M^{Old})$ indicated in Eqs. 33 and 34 is seen clearly in Fig. 1. The quantities V^σ, $\sigma = +, -$ are linear functions of $(w_M - w_M^{Old})$ for both $(w_M - w_M^{Old}) > 0$ and $(w_M - w_M^{Old}) < 0$, with a discontinuity in slope at $(w_M - w_M^{Old}) = 0$. Futhermore, the slope of $V^{(+)}$ for $(w_M - w_M^{Old}) > 0$ is equal to the slope of $V^{(-)}$ for $(w_M - w_M^{Old}) < 0$.

Similarly, the quantities $\Delta^{(\sigma)}(\tau, \rho)$ have the dependence on $(w_M - w_M^{old})$ indicated in the equations

$$\begin{aligned}
\Delta^{(\sigma)}(w_M - w_M^{Old}, \tau, \rho) &= (w_M - w_M^{Old}) F_\Delta^{(\sigma)}(\tau, \rho), \\
\Delta_0^{(\sigma)}(w_M - w_M^{Old}) &= (w_M - w_M^{Old}) F_{\Delta_0}^{(\sigma)},
\end{aligned} \tag{35}$$

for $(w_M - w_M^{Old}) > 0$, and

$$\begin{aligned}
\Delta^{(\sigma)}(w_M - w_M^{Old}, \tau, \rho) &= -(w_M - w_M^{Old}) F_\Delta^{(\overline{\sigma})}(\tau, \rho), \\
\Delta_0^{(\sigma)}(w_M - w_M^{Old}) &= -(w_M - w_M^{Old}) F_{\Delta_0}^{(\overline{\sigma})},
\end{aligned} \tag{36}$$

for $(w_M - w_M^{Old}) < 0$. Also,

$$\Gamma^{(\sigma)}(\Delta, w_M - w_M^{Old}) = \frac{1}{w_M - w_M^{Old}} F_\Gamma^{(\sigma)}\left(\frac{\Delta}{w_M - w_M^{Old}}\right) \tag{37}$$

for $(w_M - w_M^{Old}) > 0$,

$$\Gamma^{(\sigma)}(\Delta, w_M - w_M^{Old}) = \frac{1}{w_M^{Old} - w_M} F_\Gamma^{(\overline{\sigma})}\left(\frac{\Delta}{w_M^{Old} - w_M}\right) \tag{38}$$

for $(w_M - w_M^{Old}) < 0$, and

$$\Gamma_0^{(\sigma)}(w_M - w_M^{Old}) = \Upsilon_0. \tag{39}$$

Finally,

$$u^{Old(\sigma)}(\lambda^{(\sigma)}, w_M - w_M^{Old}) = u_M + (w_M - w_M^{Old}) F_u^{(\sigma)}(\lambda^{(\sigma)}) \tag{40}$$

for $(w_M - w_M^{Old}) > 0$ and

$$u^{Old(\sigma)}(\lambda^{(\sigma)}, w_M - w_M^{Old}) = u_M + (w_M - w_M^{Old}) F_u^{(\overline{\sigma})}(\lambda^{(\sigma)}) \tag{41}$$

for $(w_M - w_M^{Old}) < 0$.

All of the $F's$ on the right-hand sides of the above equations, like $F_V^{(\sigma)}$, have an implicit dependence on the variables $w_{n(p)}$ and $u_{n(p)}$, $p = 1, \ldots, K$, describing the current CDS market, and do not depend on $(w_M - w_M^{Old})$.

10.1 Examples Showing Scaling and its Absence

This subsection gives numerical results for bid and ask prices for an illiquid CDS for the cases where scaling is present, and where it is absent. Hedging is assumed to have been carried using the hedge that enforces the no-arbitrage bounds.

The bid $(u^{Old(-)})$ and ask $(u^{Old(+)})$ prices for an illiquid CDS, computed for the parameter values in the Standard Parameter List printed at the beginning of Section 4, which includes a targeted expected return on capital at risk of $R_T^{(+)} = R_T^{(-)} = 25\%$, are plotted in Fig. 11. Note that there is a CDS contract on the current market having the same maturity as the illiquid CDS ($T_M = 5$ yrs) and which has an upfront value of $u_M = 24.05\%$. The Figure shows that the qualitative features implied by the scaling results of Eqs. 40 and 41 are obtained. The plots of bid and ask prices vary linearly with $w_M - w_M^{Old}$ for both $w_M - w_M^{Old} > 0$, and for $w_M - w_M^{Old} < 0$. Furthermore, for example, the slope of the ask price versus $w_M - w_M^{Old}$ when $w_M - w_M^{Old} > 0$ is equal to the slope of the bid price versus $w_M - w_M^{Old}$ when $w_M - w_M^{Old} < 0$. The qualitative behavior of the bid and ask prices shown in Fig. 11 is quite different from the behavior given by complete-market risk-neutral pricing as described in the paragraph surrounding Eq. 1, and for which there is no bid-ask spread.

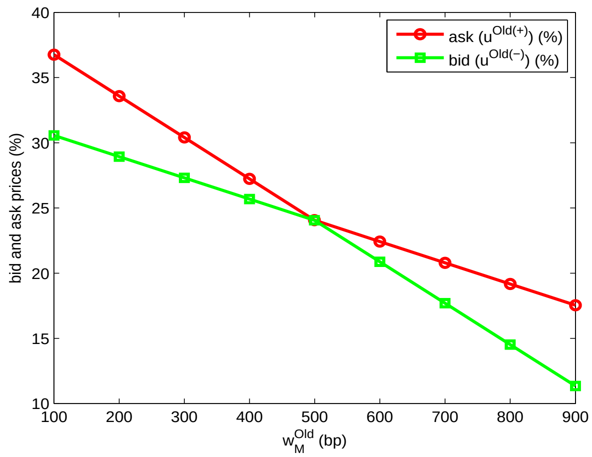

Figure 11: Plots versus w_M^{Old} of the bid and ask prices of a illiquid CDS. The data used to obtain this Figure can be found in the Standard Parameter List given at the beginning of Section 4. Thus, there is a CDS on the current market of the same maturity as the illiquid CDS and for an upfront price of $u_M = 24.05\%$.

The no-arbitrage bounds for the bid and ask prices for the same input parameters as used for Fig. 11 are shown in Fig. 1. In general, for incomplete markets, the lower and upper no-arbitrage bounds are considered to be too low and too high to serve as useful guides for the bid and ask prices (e.g. see Staum (2008)). Note that the bid and ask prices shown in Fig. 11 are significantly higher and lower, respectively, than the no-arbitrage bounds.

When $w_M^{Old} \neq w_M$, the illiquid CDS contract does not coincide with the contract of one of the CDSs on the current market, and there is a non-zero bid-ask spread reflecting the fact that the illiquid contract can not be perfectly hedged on the current market. On the other hand, for the special case $w_M = w_M^{Old}$, the illiquid CDS contract has the same future payoff stream as the CDS on the current market of the same maturity, and therefore can be perfectly hedged by buying the offsetting position from the market. The bid and ask prices of the illiquid CDS in this case are equal to each other, and equal to the cost u_M of buying the hedge on the market.

The input data used to obtain Fig. 12 is the same as that used to obtain Fig. 11, except that the maturity of the illiquid CDS being considered is $T_M = 6yr$, for which there is no CDS on the current market. However, a CDS with maturity 7 yr, with an upfront value given in Table 1, has been added to the current market and is a component of the hedging portfolios. The results shown in Figs. 11 and 12 are qualitatively similar except that there is no value of w_M^{Old} in Fig. 12 at which the bid-ask spread is zero, and the strict linear variation with w_M^{Old} exhibited in Fig. 11, and which is a consequence of the scaling results, is not seen in Fig. 12.

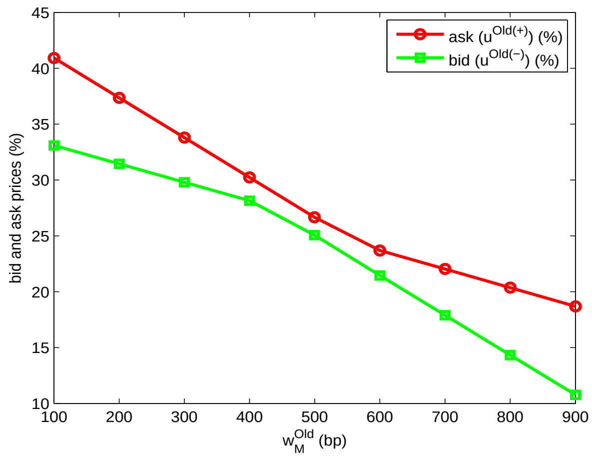

Figure 12: Plots versus w^{Old} of the bid and ask prices for a illiquid CDS of maturity $T_M = 6$ yr. The input parameters used to obtain this figure are the same as those used to obtain Fig. 11, except that a CDS of maturity 7 yr has been added to the current market. Note that there is no CDS of maturity $T_M = 6$ yr on the current market.

11 Robustness

The physical probability measure used to obtain the recommended values $u^{Old(\sigma)}$ of a illiquid CDS must be fixed by the individual or firm carrying out the valuation procedure, based on research to determine realistic physical probability distributions of default times and recovery rates given default. Thus, the assumed physical probability distributions, while based on research, also depends to a certain extent on judgement. Confidence in the procedure will be increased, therefore, if it can be shown that changes in the parameters determining the physical probability distributions, and thus, the distributions themselves, produce only relatively small changes in the recommended values of illiquid CDSs being valued (e.g. see Staum (2008)) . This section gives an example Figure showing how the robustness of the bid and ask prices with respect to large variations of the assumed physical probability distributions can be assessed, and also shows that in the example considered, the bid and ask prices are reasonably robust. The results described here are for the optimum hedge of Section 9.1 or the hedge that enforces the no-arbitrage bounds, which in the example studied are the same.

Fig. 13 describes the dependence of the bid and ask prices on the assumed physical default-time probability distribution and recovery-rate distribution. Note that substantial changes in the assumed physical probability distributions are considered, corresponding to a probability of default within 1 year varying between 20% to 60%, and a recovery rate probability density changing from

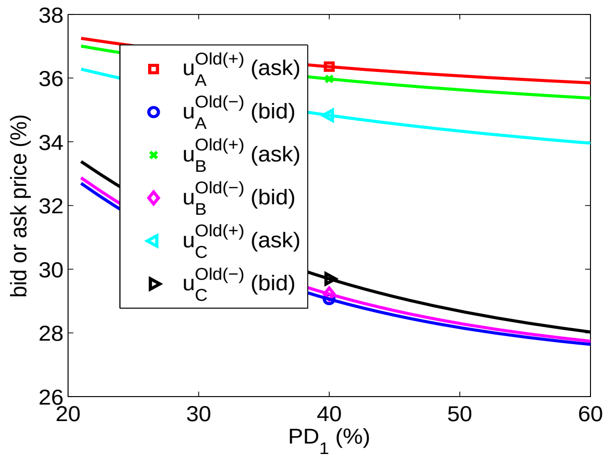

Figure 13: Plot of bid and ask prices of the illiquid CDS studied in Fig. 6 versus PD_1. The bid and ask prices labelled by the subscript A (or B, or C) have been calculated assuming the recovery rate probability density $\gamma_A(\rho)$ (or $\gamma_B(\rho)$, or $\gamma_C(\rho)$) of Fig. 2. Here $R_T = 25\%$.

γ_A to γ_B, or to γ_C, of Fig. 2. (In fact the results shown in Fig. 13 depend only on $\overline{\rho}$ and not on the details of the recovery-rate probability density.) The other parameter values assumed in obtaining these results are those from the Standard Parameter List given at the beginning of Section 4. In going from γ_A to γ_B, the expected recovery rate increases from $\overline{\rho}_A = 20\%$ to $\overline{\rho}_B = 25\%$, an increase of 25 %, while the ask price changes by only about 4 % of bid-ask spread. Also, in going from γ_A to γ_C, the expected recovery rate increases by 100 % (from 20 % to 40 %), while the change in the ask price is about 20 % of the bid-ask spread. The changes in the bid prices are roughly a factor of 2 smaller than the changes in the ask prices. Thus, the bid and ask prices are reasonably robust with respect to the choice of the recovery rate probability density and the default time distribution.

In addition, one can easily calculate the numerical sensitivities of the good-deal bid and ask prices with respect to small changes in the parameters defining the physical probability distributions, such as PD_1 for the default time distribution, and the mean and standard deviation parameters that define the recovery rate density. This is an efficient way of examining the robustness of the results with respect to small changes in the physical probability measure. However, looking at the sensitivities only will not detect a sudden change in the hedge that might occur when the parameters change by a finite values. In summary, there are procedures for evaluating the robustness of the bid and ask prices for small and larger changes of parameters.

12 Summary and Conclusions

CDS contracts purchased on the CDS market some time ago may have spreads and/or termination dates that do not correspond to CDS contracts on the current CDS market, and are thus illiquid. Such contracts are valued in this article by considering a transaction in which an investor sells the illiquid contract to a dealer. Although the dealer hedges the contract as well as possible in terms of CDSs on the current liquid CDS market, she is still left with a hedged position that is risky. This means that the realized value of the dealer's hedged position is uncertain, and in particular, that the dealer has the possibility to realize a positive loss on the transaction, the maximum value of which is called the capital at risk.

The article sets up a detailed procedure for arriving at what are called good-deal bounds (from the point of view of the dealer) for the bid and ask prices of the illiquid CDS. For ask prices greater than a lower ask-price good-deal bound, and for bid prices less than an upper good-deal bid-price bound, the dealer is guaranteed to make an expected return on the capital at risk which is greater than the lower good-deal bound on the expected return on the capital at risk. This lower good-deal bound on the expected return is set by the dealer as the lowest value of expected return that would be acceptable to her. Similarly, the good-deal bounds could be described in terms of an analogue of the Sharpe ratio, rather than the expected return.

The implementation of this program is carried out by first setting up a model that can be solved to determine no-arbitrage bounds for the bid and ask prices of the illiquid CDS. Athough these bounds are too wide to be of use as bid and ask prices for an illiquid CDS, detailed investigation shows that the hedging portfolios that enforce these bounds are useful hedges for the illiquid CDS. Also, a physical probability measure is established which requires the dealer to specify a probability density for the default times of the reference name (which can be simply done in terms of an estimate of the probability of default of reference name of the illiquid CDS within one year) as well as a probability density for the recovery rate. Once the hedging portfolios and the physical probability measure have been determined, and the dealer has established a lower good-deal bound for the expected return, the good deal bounds on the bid and ask prices can be calculated. Although much of the article is spent in studying valuation in terms of the hedges enforcing the no-arbitrage bounds, the developments in this study are easily extended so as to give various optimal hedging procedures. For example, one can determine the best hedge in the sense of minimizing the dealer's capital at risk for a given expected return (or for a given effective Sharpe ratio). Also, one can determine the hedge that gives the minimum capital at risk for a given expected return (or for a given effective Sharpe ratio) and for given values for the bid or ask prices. The procedures described have been implemented numerically, and numerical plots illustrating the behavior of a number of important quantities have been included in the article. Also, a procedure for examining the robustness of the bid and ask prices with respect to a mis-specification of the physical probability measure is described, and its application to a numerical example shows that the good-deal bid and ask prices for that example are reasonably robust.

References

Bernardo, Antonio and Olivier Ledoit, 2000, "Gain, Loss and Asset Pricing," *Journal of Political Economy*, **108**, 144

Beumee, Johan, Damiano Brigo, Gareth Stoyle and Daniel Schiemert, 2009, "Charting a Course Through the CDS Big Bang," available at http://www.defaultrisk.com/pp_crdrv176.htm

Bielecki, Tomasz R, Monique Jeanblanc and Marek Rutkowski, 2004, "Mean-Variance Hedging of Defaultable Claims," available at: http://www.defaultrisk.com/pp_other_22.htm

Bielecki, Tomasz R, Monique Jeanblanc and Marek Rutkowski, 2004, "Indifference Pricing and Hedging of Defaultable Claims," available at: http://www.defaultrisk.com/pp_price105.htm

BNP, 2004, "Understanding Credit Derivatives Vol.4, CDS Pricing", available at http://www.classiccmp.org/transputer/finengineer/

Brigo, Damiano, Mirela Predescu and Agostino Capponi, 2010, "Credit Default Swaps Liquidity Modeling: A survey," available at: papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID1564327_code1178784.pdf?...

Carr, Peter, Helyette Geman, and Dilip B Madan, 2001, "Pricing and Hedging in Incomplete Markets," *Journal of Financial Economics,* **62**, 131

Černỳ, Aleš and Stewart Hodges, 2002, "The Theory of Good Deal Pricing in Financial Markets," In: Geman, H, D Madan, S Pliska and T Vorst (Eds.), *Mathematical Finance - Bachelier Congress 2000*, p. 175 Springer-Verlag,Berlin

Cochrane, J H and J Saà-Requejo, 2000, "Beyond Arbitrage: Good-deal Asset Price Bounds in Incomplete Markets", *Journal of Political Economy*, **108**,79

Felsenheimer, J, P Gisdakis and M Zaiser, 2006, *Active Credit Portfolio Management: A Practical Guide to Credit Risk Management Strategies*, Wiley-VHC.

ISDA CDS Standard Model, 2009, available at http://www.cdsmodel.com/cdsmodel/

Markit Group Limited, 2009a, *The CDS Big Bang: Understanding the Changes to the Global CDS Contract and North American Conventions*, available at http://www.markit.com/cds/announcements/resource/cds_big_bang.pdf

Markit Group Limited, 2009b, *CDS Small Bang: Understanding the Global Contract and European Convention Changes* available at http://www.iinews.com/site/cds_small_bang_07202009.pdf

O'Kane, Dominic, 2008, *Modelling Single-name and Multi-name Credit Derivatives*, Wiley, Chichester.

Pliska, S R, 1997 *Introduction to Mathematical Finance*, Blackwell, Oxford.

Reuters, June 12, 2009, "General Motors CDS Worth 12.5% in Auction", news report, available at www.reuters.com/article/idUSTRE55B5JO20090612

Staum, Jeremy, 2004, "Fundamental theorems of asset pricing for good deal bounds," *Mathematical Finance,* **14** (2), 141.

Staum, Jeremy, 2008, "Incomplete Markets", in *Handbooks in OR and MS, Vol. 15, Chapter 12*, J.R. Birge and V. Linetsky (Eds.), Elsevier.