



**CARRY MEASUREMENT FOR
CAPITAL STRUCTURE
ARBITRAGE INVESTMENTS**

Jan-Frederik Mai
XAIA Investment GmbH
Sonnenstraße 19, 80331 München, Germany
jan-frederik.mai@xaia.com

July 10, 2015

Abstract An expected return (“carry”) measurement for capital structure arbitrage positions is discussed. Even though such a measurement is extremely difficult and suffers from numerous theoretical drawbacks, the present article should help to provide a good feel for the mechanics of classical capital structure arbitrage investments.

1 Introduction The initial goal of the research leading to this document was to introduce a reasonable, yet efficiently implementable, expected return measurement for capital structure arbitrage (cap struc arb) investments. A cap struc arb investment is a portfolio consisting of several financial instruments investing into different tranches of a single company’s capital structure, with the intention to have minimal exposure to the company’s default risk. A combination of long credit instruments and short equity instruments of the company results in cancellation effects minimizing the default risk exposure. The default risk being minimized by the equity hedge, such a position can still exhibit a significant, model-dependent, expected income. Depending on the applied model, this might be explained by an exposure to mark-to-market risk, liquidity risk, maturity mismatch risk, recovery risk, legal risk, etc. - or, more favorably, by price inconsistencies between credit and equity markets. A cap struc arb position is entered into when portfolio management decides that taking the involved risk is well-paid – even so well that one calls it euphemistically an “arbitrage”. This decision is based on quantitative methods, one being an expected income estimate as discussed in the present article.
How to define an expected return measure for such a position? Let us briefly recall the idea of the negative basis measurement (HY) described in Bernhart, Mai (2012), because it is educational and the idea underlying the cap struc arb trade is similar. The core idea in the (HY) neg basis measurement was to define an arbitrage-free model explaining both bond and CDS prices jointly, but one model parameter – a parallel shift of the reference discounting curve – was interpreted as a measure for the mispricing between the two. This parameter, called negative basis (HY), could also be interpreted as a performance measure. Ideally, a similar procedure is possible for cap struc arb positions as well. However, several reasons make the present task more difficult:

- (a) *Sound theoretical underpinning?* It is not trivial to define an arbitrage-free model that can explain all prices of credit and



equity instruments jointly. In particular, one criterion that portfolio management applies in order to detect lucrative investments is the impossibility of explaining all given prices by a (large) cosmos of credit-equity models. Unlike in the negative basis situation, the present author is not aware of a method to add an additional free and intuitive parameter to a cosmos of credit-equity models that makes all prices explainable jointly.

- (b) *Invested capital?* It is common to interpret and express a performance measure in terms of an annualized expected return. In order to do this, required is typically something like an invested capital amount that serves as denominator. However, in a pure derivative position one might encounter situations without initial investment, so it is difficult to tell how much the position earns in relative terms (relative to what?).
- (c) *Indispensability of standard pricing routines?* It is desirable to automate the performance measurement on a fund level, i.e. to compute this measure for each single position in the fund on a regular basis, e.g. daily. However, sophisticated credit-equity models rely on a battery of parameters whose calibration to market data is difficult to automate in a fast way. For huge computations on a fund level, e.g. like a Value-at-Risk computation, the industry typically applies techniques which work instrument-by-instrument (opposed to position-by-position), and for each single instrument a low-parametric, “standard” pricing routine is chosen. The reason for this is that (i) the implementation of the pricing routines must be efficient and (ii) the (few) model parameters must be intuitive and robustly (and automatically) retrievable from available data. Property (ii) typically stands in glaring contrast to the model being capable of explaining as many stylized empirical facts as possible, simply because the world’s complexity can often not be measured satisfactorily in terms of one or two parameters.

2 Carry measurement

The most intuitive and common expected return measurement approach is to compute an *annualized, expected return on investment* μ_{1y} , which is generically defined as

$$\mu_{1y} = \frac{\text{expected income (within next year)}}{\text{invested capital}}.$$

That’s precisely the strategy that is pursued in the present document. In the sequel, it will be described how the expected income and the invested capital can be computed. With regards to the aforementioned difficulties (a), (b), (c), the method is classified as follows:

- (a) *Sound theoretical underpinning?* Not really. The method is not based on a global financial model explaining all instruments in one position jointly, and additionally filtering out the remaining “arbitrage carry” in terms of a single model parameter. Rather the method relies on a rudimentary aggregation of globally inconsistent, instrument-specific return measurements.



- (b) *Invested capital?* An appropriate definition for this denominator is provided in Subsection 2.2. In particular, capital that is not invested but potentially at risk will be incorporated appropriately.
- (c) *Indispensability of standard pricing routines?* For each and every instrument type the method applies a “standard” pricing routine, e.g. Black-Scholes for equity options. This is what makes the approach viable and efficient. The default risk component is taken into account in a simplified – and theoretically unjustified – form.

2.1 Basic customizations

We call μ_{1y} *carry* instead of *expected return*, even though this is just semantics. Moreover, instead of computing the expected income within the next year, we compute an expected income within the next month, and therefore obtain a monthly carry measure μ_{1m} that needs to be annualized afterwards, e.g. by a simplified computation such as

$$\mu_{1y} = 12 \times \mu_{1m}, \quad \text{or} \quad \mu_{1y} = (1 + \mu_{1m})^{12} - 1.$$

The reason for choosing a one-month period instead of a one-year period is two-fold: On the one hand, many equity instruments in a cap struc arb position mature long before the end of the next year. Therefore, in order to compute a one-year expected return one would need to make assumptions about how these maturing positions are re-invested, which is a very difficult task. A definition of such appropriate “rolling algorithms” can be circumvented by considering only a one-month period. If positions end within the next month, it is simply assumed that the final proceeds from those (if any) are re-invested into the risk-free bank account. This is a simplifying but in our view acceptable solution, because it does not happen too often in practice, as hedges are typically rolled prior to one month before maturity of the respective instrument. On the other hand, the method that is applied in order to compute the expected income (see Subsection 2.3) relies on future value projections of “world-drivers”, i.e. expected future values of interest rates, credit spreads and equities are assumed. The shorter the considered projection period, the more reliable this assumption becomes, since in all commonly applied financial models uncertainty of “world drivers” grows over time (we know the world in one month better than we understand the world in one year).

2.2 Invested capital

First of all, it is important to notice that the carry measurement should be an indication of how lucrative the position is in the future, but it should not take into account PnL from the past. This means it should be defined forward-looking. In particular, the capital we have spent into the position at original trade inception plays no role in the following definition of the invested capital. The invested capital of a cap struc arb position is generally defined as follows:

$$\text{invested capital} = \text{market value} + \text{capital at stake}.$$

The *market value* of the position is simply the sum over all market values of the single instruments in the position. Notice that for



derivatives the market value might also be negative, e.g. CDS or equity forwards. The *capital at stake* of the position is the sum over all capitals at stake for every individual instrument in the position. The latter are defined as the cash amounts required in order to make any potential payments which are triggered by an immediate default event. For example, the capital at stake is zero for bonds or equity options because – apart from a potential decrease in market value – no payment is triggered in case of a default event. However, for short CDS a default event triggers a default compensation payment that needs to be made. The same is true for a (long) equity forward, when a default event implies that the stock price drops to a value close to zero but the forward contract forces us to buy the stock at a higher price.

Example 2.1 (Capital at stake for short CDS)

For a short CDS position¹ we define the capital at stake as the nominal of the CDS. Assuming we sell CDS protection for which we receive an upfront payment of 30% of the CDS nominal, the market value of the CDS equals -30% of the CDS nominal, and the capital at stake equals the CDS nominal. Consequently, the invested capital is defined as 70% of the CDS nominal. In theory, when selling CDS protection there is no need to hold the capital at stake (i.e. the nominal) liquidly available in cash. Hence, short CDS positions could be used to build up leverage. However, it would be more conservative to create default compensation reserves, and one proposal for the size of this reserve is the capital at stake. Defining the latter as the nominal of the CDS is conservative in the sense that it prepares one for the worst possible scenario: a credit event happens, the realized recovery rate equals zero, and there are no proceeds from potential hedging positions that are held with the intention to offset some of the losses (e.g. put options). In practice, depending on one's personal risk appetite, we think it is reasonable to keep as liquid capital reserve any amount between the CDS nominal (= our definition of capital at stake, which is very conservative) and the expected capital which is required in case of a credit event, which is typically much lower, depending on offsetting hedging positions and recovery assumptions. Keeping fewer capital reserves than this results in a truly leveraged position.

2.3 Expected income

For each and every instrument in the position we compute an expected income within the next month. Given the probability of a default event within the next month, denoted by p_{1m} , this expected income within the next month of a single instrument is computed via the formula

$$\begin{aligned} \text{expected income (1m)} &= p_{1m} \times \text{jump-to-default of the instrument} \\ &\quad + (1 - p_{1m}) \times \text{expected income (1m) given no default}. \end{aligned}$$

Since it is often more natural to consider an annualized carry measurement, we also compute an expected income within the

¹Without additional recovery swap. See Ferger (2015) for a treatment of recovery swaps.



next year by the formula

$$\begin{aligned}\text{expected income (1y)} &= p_{1m} \times \text{jump-to-default of the instrument} \\ &+ 12 \times (1 - p_{1m}) \times \text{expected income (1m) given no default.}\end{aligned}$$

The idea of the latter formula is that the expected one-month income is earned twelve consecutive months. Admittedly, alternative definitions – e.g. involving a one-year survival probability – are imaginable as well, but we consider the presented formula a reasonable approximation. For credit instruments, the default probability p_{1m} is obtained from the given quotes of the respective instruments. For equity derivatives, the default probability p_{1m} is implied from the credit instruments in the position, assuming that each cap struc arb position contains at least one credit instrument. If a CDS curve is available, then this is the preferred choice for retrieving the default probability, otherwise p_{1m} is obtained as the mean over all bond-implied default probabilities. Clearly, if several bonds imply different default probabilities then the aggregation of these instrument-by-instrument computations is inconsistent theoretically. However, it is the best we can do under the side constraint of requiring a high level of computational efficiency. Depending on one's interest in the analysis it might also make sense to assume $p_{1m} = 0$, which results in a “carry measurement given no default”.

The expected income given no default is computed for each instrument using a “standard” pricing method. These standard pricing routines are listed in Table 1, together with their underlying parameters. For credit instruments, the realized recovery rate assumption R is exogenous input obtained from the front office system, i.e. they are manual input by the traders. The volatility parameter σ in the simple JTD-model² used for convertible bonds is retrieved likewise. For bonds, convertible bonds and equity options, this leaves us with one free parameter, which is chosen in such a way that the observed market value of the respective instrument is matched perfectly. The piecewise constant default intensities for CDS are bootstrapped from a CDS credit curve that is once again retrieved from the front office management system (and used there to price all the CDS, so these prices are perfectly explained).

For each instrument the expected income given no default is computed via

$$\begin{aligned}\text{expected income given no default} &= \text{cashflow earnings} \\ &+ \text{expected value in one month} - \text{current market value.}\end{aligned}$$

Cashflow earnings comprise received coupon payments from regular bonds, convertible bonds or short CDS positions. In theory, they can also be negative in case of a long CDS position. Equity instruments do not have cashflow earnings, because dividends are ignored throughout. The expected value in one month is computed using the pricing models depicted in Table 1, which are determined from exogenous input parameters and current market values, as explained above. However, we additionally require

²See Mai (2012) for a description of the simple JTD-credit-equity model.



instrument	pricing algorithm	parameter(s)
bond	constant intensity λ	R, λ
convertible bond	simple JTD-model	R, λ, σ
CDS, recovery swap	piecewise const. intens. $\lambda(t)$	$R, \lambda(t)$
equity option	Black-Scholes	σ
equity forward	model-free	n.a.
stock	model-free	n.a.

Table 1: Pricing methodologies and parameters for each instrument type. The red parameters are exogenous user input by the traders, while the remaining parameters are fitted to the observed market prices of the respective instruments.

expectations about how some of the input parameters change within the next month. More precisely, we define expectations about (i) interest rate curves, (ii) stock prices, and (iii) default intensities within one month. Our base scenario assumptions for these parameters are as follows: (i) the interest rate curves are projected into the future according to today's forward rates, (ii) all stock prices are assumed to remain unaltered in value, and (iii) default intensity curves are assumed to be exactly the same in one month as they are now.

3 Carry vs. mispricing check

When is a cap struc arb position an interesting investment for our fund? One criterion is that it has a (sufficiently large) positive carry, which is computed as described above. However, the previous sections indicate that the reliability of such a carry measure might not be too strong, due to theoretical flaws. A second criterion that is applied by portfolio management in order to detect lucrative cap struc arb positions is a *mispicing check*. A potential position features a favorable mispricing if the prices of the involved credit and equity instruments are inconsistent. This inconsistency is checked by a battery of credit-equity models.

A natural question is: how do the two criteria compare? We provide a seemingly puzzling example in the case of a short CDS versus put option position, which is similar to the situation of one of our positions in the fund XAIA Credit Debt Capital. Assume we sell 5 year CDS protection with a running coupon of 500 bps. The upfront payment we receive is 10%. We assume a CDS notional of 3 million EUR. For the hedging equity position we buy 1.5 year put options with strike 0.8 EUR, assuming the current stock price level is 1 EUR. The market quoted price for the put option is assumed to be 0.143 EUR. Based on our analysis using credit-equity pricing models, a put price that is arbitrage-consistent with the observed CDS curve should be at least 0.18 EUR, which is significantly higher than the observed 0.143 EUR. This means that we detect a favourable mispricing. However, the position has a slightly negative carry, when we buy 3,196,347 put options, so that we are instantaneously JTD-neutral under a 20% recovery assumption. How can that be?

Figure 1 illustrates the jump-to-default, the carry, and the maximal possible gain of the position, in percent of the invested ca-

pital. The maximal possible gain is achieved when immediately after entering into the position the CDS spread drops to zero and the stock price raises to infinity (CDS makes a massive gain exceeding the losses on the put hedge by far). One can observe that the position has a massive upside in case the creditworthiness of the company improves, but if no change (or even a worsening) of the creditworthiness occurs the position is almost flat. It is therefore justified to call this position an “arbitrage” because there is an upside, but (almost) no downside potential. For comparison, the white bars in Figure 1 illustrate the same values assuming the put price was actually 1.8 EUR instead of 1.43 EUR, implying an arbitrage-free setup. One observes that a downside appears in this case.

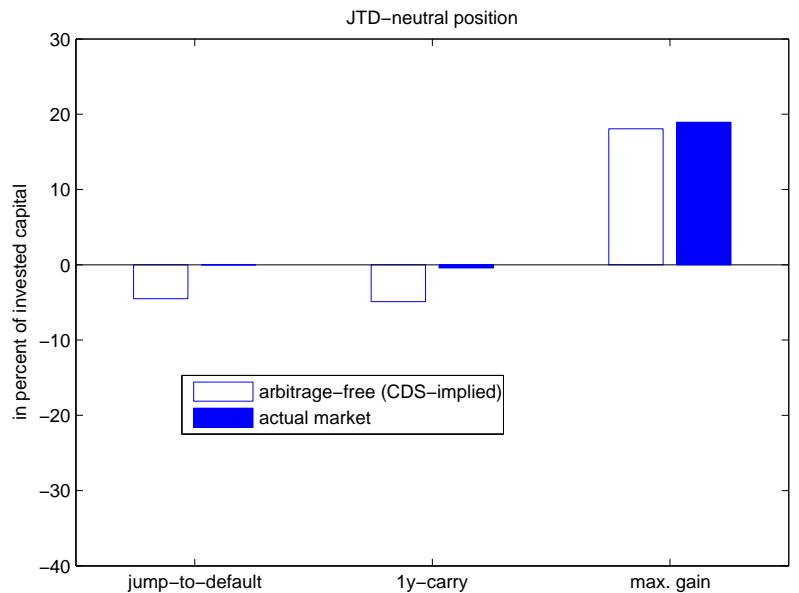


Fig. 1: Jump-to-default, annualized carry, and maximal possible gain of the described cap struc arb position.

If one wanted to make the position carry positive, say one would like to tune the position such that the carry equals 3%, one must buy fewer put options. This alternative position is visualized in Figure 2. It is observed that this trade is not a striking arbitrage, because now the downside risk increases dramatically (and the upside potential increases as well).

Figure 3 visualizes the model-implied probabilities of the stock price at the maturity of the put option in 1.5 years. Once these probabilities are derived from the observed CDS quotes, and once they are derived from the observed put option quotes. It is observed that the CDS quotes imply a higher default probability than the option quotes (7% compared to 4.8%). Moreover, it is also observed that the variance implied from the CDS is higher than the variance implied from the put options. These two observations explain why the CDS-implied put prices are higher than the market put prices. Moreover, please notice that the three extreme scenarios visualized in Figures 1 and 2 could be thought of as arising in Figure 3 on the left (jump-to-default), in the middle (ATM, carry), and on the right (maximal gain) of the x -axis.

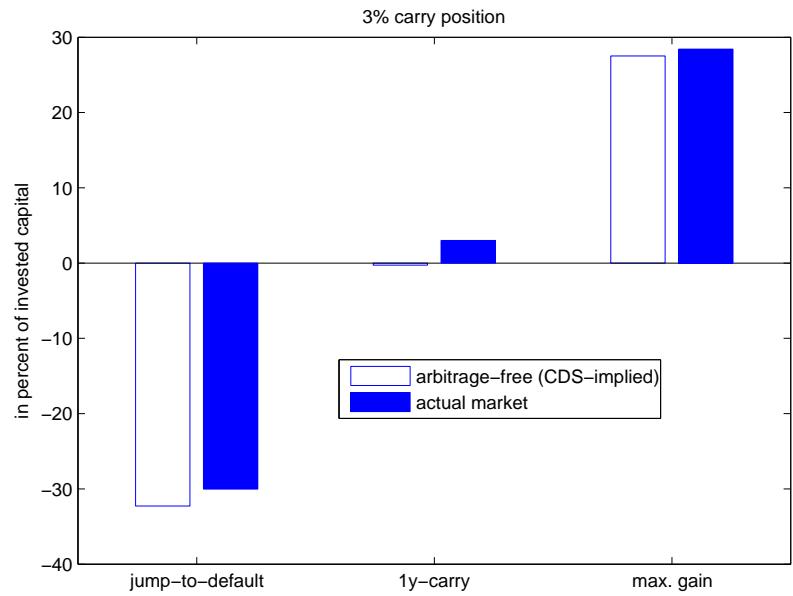


Fig. 2: Jump-to-default, annualized carry, and maximal possible gain of the described cap struc arb position, with fewer put options.

Hence, the CDS-implied probabilities imply a greater likelihood for the upside of the position than the put-implied probabilities do.

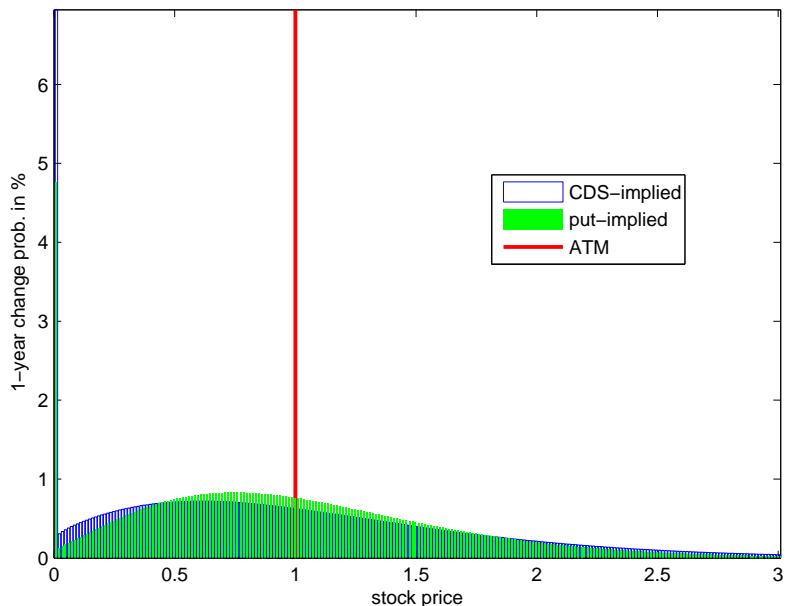


Fig. 3: Model-implied probabilities for the stock price in 1.5 years.

- References**
- G. Bernhart, J.-F. Mai, Negative basis measurement: finding the holy scale, *XAIA homepage article* (2012).
 - N. Ferger, Recovery swaps in practice, *XAIA homepage article* (2015).
 - J.-F. Mai, The joint modeling of debt and equity: an introduction, *XAIA homepage article* (2012).