FACULTEIT ECONOMIE
EN BEDRIJFSKUNDE

TWEEKERKENSTRAAT 2
B-9000 GENT
Tel. : 32 - (0)9 – 264.34.61
Fax. : 32 - (0)9 – 264.35.92

WORKING PAPER

Crystallization – the Hidden Dimension of Hedge Funds' Fee Structure

Gert Elaut, Ghent University, BELGIUM

Michael Frömmel, Ghent University, BELGIUM

John Sjödin, Ghent University, BELGIUM

and RPM Risk & Portfolio Management AB, SWEDEN

February 2014

2014/872

D/2014/7012/01

Crystallization – the Hidden Dimension of Hedge Funds' Fee Structure*†

Gert Elaut, Ghent University, BELGIUM

Michael Frömmel, Ghent University, BELGIUM

John Sjödin, Ghent University, BELGIUM

and RPM Risk & Portfolio Management AB, SWEDEN‡

First version: 30/11/2013

This version: 09/01/2014

Abstract

We investigate the implications of variations in the frequency with which hedge fund managers update their high-water mark on fees paid by investors. We first document the crystallization frequencies used by Commodity Trading Advisors (CTAs) and then perform simulations and a bootstrap analysis. We find a statistically and economically significant effect of the crystallization frequency on the total fee load. Hedge funds' total fee load increases significantly as the crystallization frequency increases. As such, our findings indicate that the total fee load not only depends on the management fee and incentive fee, but also on the crystallization frequency set by the manager.

JEL classification: G11; G23

Keywords: Hedge fund; CTA; Incentive fee; Fee structure; Gross returns; Crystallization

*Funding from the EC Grant Agreement n. 324440 (Futures) Marie Curie Action Industry-Academia Partnership and Pathways Seventh Framework Programme is gratefully acknowledged.

†We thank Alexander Mende, Péter Erdős, and Magnus Kottenauer of RPM for the many stimulating discussions and comments.

‡*Corresponding author*: Direct: +46 8 440 69 21; Fax: +46 8 440 69 01. Email address: `John.Sjodin@rpm.se`

1 Introduction

The impact of the two components of hedge funds' fee structure, the incentive fee and the high-water mark clause, on hedge fund behaviour has been discussed extensively in the academic literature. Especially their effect on fund managers' risk-taking behaviour has received considerable attention[1]. However, the fee structure also has more direct consequences for investors, apart from changing the risk profile of the investment. Fees impact long-term wealth and investors are more and more starting to realize this, not in the least because of the current low yield environment. Consequently, hedge funds' fees are now subject to closer scrutiny and are negotiated more often than in the past.

To illustrate this downward pressure on hedge funds' headline fee levels, we report the management fee and incentive fee of newly launched CTAs in the BarclayHedge database in Table 1. The table illustrates that, while there has been no significant change in the incentive fee levels, the average management fee has been decreasing steadily over time.

[Table 1 about here.]

The typical fee structure of hedge funds and CTAs is made up of a management fee (usually 2% of assets under management) and an incentive fee (usually of 20% of profits). This 2/20-fee structure is and has been the standard cost for allocations in the hedge fund industry. The incentive fee is generally supplemented with a high-water mark, such that investors only pay an incentive fee once any previous underperformance has been made up for.

However, the headline fee levels are only one aspect of the fee structure that should be considered. Other elements of the fee structure also have a significant impact on the total fee load. One element that is usually not taken into consideration when discussing hedge funds' fees is the frequency with which a fund updates its high-water mark. This frequency is commonly referred to as the *crystallization frequency* or the *incentive fee payment schedule*. The fees an investor pays to the fund manager does not only depend on the management and the performance fee, but also on the crystallization frequency.

[1] Studies include Goetzmann, Ingersoll, and Ross (2003), Hodder and Jackwerth (2007), Kouwenberg and Ziemba (2007), Panageas and Westerfield (2009), and Agarwal, Daniel, and Naik (2009).

The crystallization frequency is the point in time when the fund manager updates the high-water mark and is paid the incentive fee. The crystallization frequency differs from the accrual schedule, which is the schedule used to calculate and charge the fee to the fund's profit and loss account. Whereas the process of fee accrual does not impact investor returns, the same is not true for the fee crystallization. As the incentive fee crystallization frequency increases, the expected total fee load charged by the hedge fund manager increases as well.

The main contribution of our study to the existing literature on hedge funds' fee structure is that we highlight and analyse the impact of the crystallization frequency on hedge funds' fee load. To the authors' best knowledge, no study has yet investigated this aspect to hedge funds' fee structure. By itself, this finding is compelling. The crystallization frequency forms the basis for the incentive fee calculation and the way hedge funds update their high-water mark. It has a material effect on the fees investors pay and, thus, should also influence hedge funds' risk-taking behaviour.

Our findings have several implications. First, we show that the crystallization frequency has both a statistically and economically significant impact on fees paid by investors. In the case of CTAs, and assuming a 2/20 fee structure, shifting from annual to quarterly crystallization leads to a 49 basis points increase in the fee load (as a percentage of assets under management). This corresponds to a 12.2% increase of an investor's total fee load. Furthermore, an incentive fee of 15% combined with monthly crystallization leads to the same total fee load as an incentive fee of 20% under annual crystallization. Based on these findings, we conclude that there are *three* factors that investors need to consider when evaluating the expected fee load; the level of management fee, the level of incentive fee and the crystallization frequency of the incentive fee. In addition, in an environment where especially management fee levels are under pressure, the relative importance of the incentive fee and crystallization in the total fee load increases.

Second, our study also has implications for the academic literature that estimates hedge funds' gross returns and fee loads as well as research on hedge funds' risk-taking behaviour. To construct gross returns, previous studies in most cases assume that incentive fees are paid at year-end (e.g. Brooks, Clare, and Motson (2007), French (2008) and

Agarwal, Daniel, and Naik (2009)), although some authors assume quarterly payment (see Bollen and Whaley (2009) and Jorion and Schwarz (2013)). Some authors also calculate hedge funds' historical fee loads in their analysis. For example, French (2008) estimates that the typical investor in U.S. equity-related hedge funds has paid an annual combined fee or total expense ratio of 3.69% p.a. over the period 2000-2007. Brooks, Clare, and Motson (2007) find that between 1994 and 2006 hedge fund fees averaged 5.15% annually. Ibbotson, Chen, and Zhu (2011) suggest a lower estimate of 3.43% p.a. for the period 1995 to 2009. Similarly, Feng, Getmansky, and Kapadia (2011) report total fees over the period 1994-2012 to be on average 3.36% of gross asset value. However, these studies do not consider the impact of the crystallization frequency on these figures. To investigate the importance of the crystallization frequency on fees, we require an accurate method to move from net returns to gross returns and vice versa, while simultaneously allowing us to vary the crystallization frequency. For this purpose we develop an algorithm for calculating gross returns in which we let the crystallization frequency vary. This way, we can evaluate the impact of the crystallization frequency on hedge funds' fee load. As for hedge funds' risk-taking behaviour, our analysis has implications for the time frame over which previous results on hedge funds' risk-taking behaviour apply. If a hedge fund manager update its high-water mark more than once a year, the trading horizion is shortened accordingly.

Third, crystallization frequencies of hedge funds have not been documented. In addition, it is possible that differences in the crystallization frequency of hedge funds are to some extent related to differences in the ability of hedge funds subcategories to value their underlying positions. Industry standards on crystallization for different hedge fund categories might therefore also differ. To avoid this possibility from having an impact, we restrict our analysis to one particular hedge fund category, namely Commodity Trading Advisors (CTAs). Unlike some other hedge fund categories, CTAs trade almost exclusively highly liquid instruments. As such, CTAs are able to value their positions on a daily basis and, thus, do not have any practical limitations regarding the calculation of their NAVs. A second reason why we focus on CTAs is the way the money is invested in this industry. In many cases, investments are done through managed accounts. Clients then negotiate directly with the fund manager, which suggests that fee terms are more easily and more often negotiated. To document CTAs' crystallization practices, we perform a survey among the constituents of the Newedge CTA Index as well as an analysis of the fee notes of CTAs in the Tremont Advisory Shareholder Services (TASS) database. We find that, at least in the case of CTAs, high-water marks are most often updated quarterly, rather than annually. This contrasts the view commonly held in the academic literature that the high-water mark is set at the end of the year.

Even in cases where it is unlikely that the crystallization frequencies can be negotiated, as is the case in comingled investment vehicles, knowledge about the effect of the crystallization frequency is important for an allocator when evaluating and comparing different fund investments. We stress that, while we focus on CTAs, the implications of our study also apply to any investment vehicle whose fee structure depends on a high-water mark provision.

In this study we focus on the impact of the crystallization frequency of the incentive fee, and we do not go into the payment frequency of the management fee. We do this mainly because the payment of the management fee does not depend on a fund's high-water mark. Additionally, the vast majority of the hedge funds charge the management fee monthly[2].

The rest of the paper proceeds as follows. In Section 2 we describe the data. Section 3 reports the result of our survey on crystallization and industry practices. In Section 4 we construct gross returns from observed net-of-fee returns. Section 5 presents the main body of the analysis, where we quantify the impact of the crystallization frequency on the total fee load and compute the trade-off that exists between different levels of the incentive fee and the crystallization frequency. In Section 6, we perform a number of robustness checks. Section 7 concludes.

[2]For the Tremont Advisory Shareholder Services (TASS) database, we find that 78% of the CTAs in the database charge the management fee on a monthly basis. 13% charges the management fee quarterly and 8% charges the management fee annually.

2 Data

We analyse the impact of the crystallization frequency on fees paid by investors by using monthly net-of-fee returns of live and dead funds labelled CTA in the BarclayHedge Database. We use a sample that covers the period January 1994 to December 2012 to mitigate a potential survivorship bias, since most databases only started collecting information on defunct programs from 1994 onwards[3]. As BarclayHedge does not report a first reporting date, we cannot eliminate the backfill bias entirely. We therefore opt for an alternative approach and remove the first twelve observations of a fund's return history, following Teo (2009)[4,5].

We further require at least twelve return observations for a fund to be included, and only include funds whose monthly returns are denominated in USD or EUR. The EUR-denominated returns are converted to USD-denominated returns, using the end-of-month spot USD/EUR exchange rate. As the analysis also requires information on the funds' management fee and incentive fee, we remove cases where at least one of the two variables is unreported[6].

We then filter the resulting sample of funds by looking at their self-declared strategy description and remove funds whose description is not consistent with the definition of CTAs. In the process, we also determine whether the program under consideration is the fund's flagship program and discard duplicates. To ensure that our results apply to funds that can be considered part of the investable universe for most CTA investors, we remove funds whose net-of-fee returns exhibit unusually low- or high levels of variation. To this end, we discard funds when the standard deviation of the observed net-of-fee returns is lower than 2% or exceeds 60% p.a. After applying these restrictions, our sample consists

[3]Gross returns are first calculated using the funds' entire return history, after which the pre-1994 period is dropped.

[4]We first calculate gross returns (see Section 4) using the fund's entire track record, and afterwards drop the first twelve observations of the fund's net-of-fee and gross returns.

[5]By keeping track of the amount of months that are backfilled when a fund is first included to BarclayHedge database, we tracked backfill bias for the period 2005-2010. For that sample period, the median (average) backfill bias was twelve (fourteen) months.

[6]Additionally, we also exclude cases where both types of fee are zero or and cases where the fee levels are deemed unreasonable low or high (management fee in excess of 5% p.a., incentive fees below 5% or above 50% p.a.).

of 1,616 unique CTA programs. Table 2 reports summary statistics for the final set of funds.

[Table 2 about here.]

3 Crystallization and Industry Practices

Since public hedge fund databases do not keep track of funds' incentive fee crystallization frequency[7], we perform a survey among the constituents of the Newedge CTA index (as of May 2013). The Newedge CTA index is designed to track the largest CTAs and aims to be representative of the managed futures space. The index is comprised of the 20 largest managers (based on AUM) who are open to new investment and that report performance on a daily basis to Newedge. We complete the results of the survey with information available on the website of the U.S. Securities and Exchange Commission (SEC)[8].

The results of the survey are reported in Figure 1. Figure 1 indicates that, in the case of CTAs, the most commonly used crystallization frequency is quarterly. In those instances where the crystallization frequency is not quarterly, we find that the frequency generally tends to be higher, rather than lower. The left panel of Figure 2 reports the share of total assets under management (AUM) of the CTAs to which every frequency applies. While quarterly crystallization remains the most commonly applied crystallization frequency, monthly crystallization applies to a larger share of total AUM than one would suspect from Figure 1. Finally, we also quantify the scope of the survey vis-à-vis total AUM by the CTA industry. The right panel of Figure 2 plots the results, which shows that the constituents of the Newedge CTA index cover 43% of assets managed by CTAs that report to BarclayHedge.

[Figure 1 about here.]

[7] TASS's questionnaire only inquires about the management fee's payment frequency; the other widely used databases' questionnaires and manuals (Hedge Fund Research (HFR), CISDM, and BarclayHedge) indicate that the databases do not keep track of the fee payment frequencies.

[8] In particular, we make use of the SEC's Investment Adviser Public Disclosure (IAPD) and the Electronic Data-Gathering, Analysis, and Retrieval (EDGAR) database.

As mentioned above, public databases do not keep track of the crystallization frequency in a systematic way. However, the fee notes in the Tremont Advisory Shareholder Services (TASS) database in a number of cases do provide a sufficient amount of information to pinpoint the crystallization frequency. Therefore, and in addition to the above survey, we also examine the fee notes of defunct and live CTAs reported in the TASS database. The results are also reported in Figure 1. Comparing these results with those of our own survey suggests that the sample of funds from TASS is characterised by higher crystallization frequencies. These differences could be due to survivorship bias as well as differences in fund size. Nevertheless, the results for the TASS sample corroborate our earlier finding that quarterly is the most common crystallization frequency. When funds use a crystallization frequency other than quarterly crystallization, the frequency tends to be higher rather than lower.

[Figure 2 about here.]

For completeness, we also look at the relationship between the reported fee levels and the crystallization frequency of the funds. It could be that funds with lower crystallization frequencies have higher incentive fee levels, such that the total fee load is comparable. To verify that this is not the case, we group the sample of funds in TASS based on their reported crystallization frequency and analyse the average incentive and management fee of the different groups. The results, reported in Table 3, indicate that funds with a higher crystallization frequency tend to have higher headline incentive fee levels. For example, the average incentive fee level for funds with monthly crystallization (22.38%) is significantly higher than that of funds that employ a quarterly crystallization frequency (21.05%), with a p-value of 0.0775. In addition, we also find that the headline management fee level tends to increase as the crystallization frequency increases. These results suggest that funds that have higher a crystallization frequency on average also exhibit higher headline fee levels.

[Table 3 about here.]

4 Construction of Gross Returns

As mentioned in the introduction, analysing the impact of the crystallization frequency on hedge funds their fee load requires a reasonable method to calculate hedge funds' gross returns and to charge fees to investors under various crystallization frequencies. To this end, we develop an algorithm that achieves this objective. We provide a thorough description of the algorithm in the appendix.

To calculate gross returns for the sample of CTAs, we assume that CTAs apply quarterly crystallization to charge incentive fees. Our survey results and the results from TASS's fee notes suggest that this is the most commonly used crystallization frequency. In addition, when CTAs apply another crystallization frequency, they generally tend to use higher crystallization frequencies. As such, the assumption of quarterly fee crystallization should lead to fairly conservative estimates of the funds' gross returns.

In Table 4 we compare the observed net-of-fee CTA returns with the obtained gross CTA returns. Funds appear to earn significantly higher risk-adjusted returns – measured by the annualized Sharpe ratio – based on gross returns, as compared to net-of-fee returns. Also, both skewness and kurtosis are significantly higher for the gross returns. Consequently, we find a higher proportion of cases in which the Jarque-Bera test for normality rejects the null hypothesis of normality. Finally, we find that both net-of-fee returns and gross returns of CTAs exhibit negative autocorrelation.

[Table 4 about here.]

5 Incentive Fee Crystallization and Fee Load

5.1 Analysis of the Historical Effect

As an introduction to our main analysis, we first estimate the crystallization frequency's potential historical effect on investor wealth. This way, we can get a feel of the economic significance of the effect of crystallization. Using the data set of gross returns obtained in Section 4, we re-apply the fund's reported headline fee levels under different crystallization frequencies. This way we obtain net-of-fee returns under different crystallization frequencies as well as the corresponding fee load.

In Table 5 we report the average gross return, average net-of-fee return, and the average fee load under the different fee crystallization schemes. The reported average net-of-fee returns are all statistically different from each other at the 1% level of significance (p-values unreported for conciseness). Furthermore, the results suggest that investors whose investment is subject to quarterly (monthly) crystallization, will earn net-of-fee returns which are on average 25 (42) basis points per year lower than in the case of annual crystallization. To put these figures into perspective, an annual difference of 42 basis points over a 10-year period will compound to a difference of 9.32% in the expected capital gain. For a MUSD 1 initial investment, this difference equals USD 63,303.

Even more important than these absolute numbers, is the impact on the risk-adjusted performance. Our results suggest that when investors move from annual to monthly crystallization, the Sharpe ratio deteriorates from 0.4 to 0.34, a 15.65% decrease.

[Table 5 about here.]

We also observe from Table 5 that management fees are slightly lower than 2% p.a., despite the positive drift in CTAs their returns. This is consistent with our finding that management fees, at least for newly launched funds, tend to be below 2% p.a. on average (see Table 1).

5.2 Simulation-based Analysis

To study the effect of the crystallization frequency on the level of fees investors pay, we now analyse the effect of crystallization in a controlled environment. In particular, we simulate monthly gross returns assuming they follow a normal distribution. We use the data set of gross returns calculated above to determine the appropriate parameters for the normal distribution. As such, we set the mean gross return equal to 0.768% per month and we assume a standard deviation of 4.683% per month (see Table 4). Next, we generate 10,000 sample paths of monthly gross returns and we apply a standard 2/20-fee structure under different crystallization frequencies. The risk-free rate used in the calculations is the average monthly US risk-free rate over the period 1994-2012, 0.28% per month. We use this framework to examine the impact of the crystallization frequency on the total fee load.

Panel A of Table 6 reports the results for one-year, three-year, and five-year investment horizons. We consider periods up to five years as this corresponds to the average age of the CTAs in the sample (see Table 2). As such, our analysis covers the relevant horizion over which the effect of crystallization applies for the majority of hedge fund investors. To gauge the significance of the results, we indicate whether the obtained fee level differs significantly from the fee load under annual crystallization. We set annual crystallization as the benchmark since most previous research made the assumption that the incentive fee is paid at the end of the year. Our results illustrate that a higher crystallization frequency always leads to a higher average fee load. Management fees are slightly higher than 2% and increasing in time due to the positive drift in the simulated gross asset values (GAV). For ease of comparison, Figure 3 and Figure 4 plot the fee load under different crystallization frequencies and the difference with annual crystallization, respectively.

[Table 6 about here.]

It is evident from panel A of Table 6 that increasing the investment horizon dampens the impact of a higher crystallization frequency on fee load. We can explain this finding by the fact that the fee loads reported for the three- and five-year investment horizons are an average across the individual years. In years where a fund is not able to charge incentive fees, the total fee is the same under different crystallization frequencies. Despite this downward drag on the total fee load, caused by years in which only a management fee is paid, the difference in fee load for the different crystallization frequencies remains significant.

Another important factor that impacts the total fee load paid by investors is the volatility level of the program. To illustrate the impact of higher volatility on the differences in fee load, we redo the simulation but change the standard deviation of the gross returns. In particular, we analyse the total fee load in the case of a 10%, 20%, and 30% volatility p.a. At the same time, we hold the expected return fixed to single out the effect of higher volatility.

The results are reported in Table 7. What is interesting to note is that the difference in total fee load for different crystallization frequencies is increasing in the volatility. As an example, consider the difference between quarterly and annual crystallization. Assuming an annual volatility of 10%, the difference in total fee load is 15 basis points, which suggests that quarterly crystallization leads to a fee load on average 4.75% higher than annual crystallization. However, assuming an annual volatility of 20% this difference increases to 25 basis points (7.08% increase in fee load). If we increase annual volatility to 30% p.a., the difference becomes 54 basis points, or an annual fee load that is 12.50% higher than under annual crystallization.

[Figure 3 about here.]

[Figure 4 about here.]

[Table 7 about here.]

5.3 Block Bootstrap Analysis

To analyse the impact of the crystallization frequency more empirically, we apply a block bootstrap by randomly sampling gross return histories and calculating the fee load under different crystallization regimes. The advantage of this approach is that it allows us to relax the distributional assumptions made with regard to the return generating process in Subsection 5.2. A block bootstrap allows us to account for higher moments in monthly returns (e.g. CTAs' returns exhibit positive skewness) and to preserve any autocorrelation present in the gross data. These properties of the return generating process can have a material impact on the results of the analysis and investors' total fee load.

In performing the block bootstrap, we consider all the potential 12/36/60-month samples in the data set of gross returns and pick 10,000 12-months, 36-month and 60-month samples. To avoid look-ahead bias, we allow the sampling procedure to select incomplete samples occurring at the end of a fund's track record. In those cases where a fund terminates before the end of the sample period, we assume that investors redeem[9]. We also assume that every draw starts the beginning of a calendar year (i.e. from January onwards). As before, we consider the case of a 2/20-fee structure, such that the obtained fee loads can be compared to the results reported in Subsection 5.2. Differences in the fee load and differences between the various crystallization frequencies should then be a consequence of features of CTAs' return generating process – including fund termination – we have not modelled in Subsection 5.2.

We report the results for the bootstrap in panel B of Table 6. Similarly to the simulation results, we find significantly higher fee loads as the crystallization frequency increases. The effect is also economically significant. For the one-year investment horizon, the total fee load is 49 (82) basis points p.a. higher in the case of quarterly (monthly) crystallization when compared to annual crystallization. This suggests that, under a 2/20-fee structure, the fee load is expected to be 12.2% (20.5%) higher if a manager charges the incentive fee quarterly (monthly), rather than annually. If the investment horizon is extended to five years, the difference decreases 23 (40) basis points p.a., a difference of 6.5% (11.4%).

Another notable finding is that for the one-year investment horizon, the management fee in the case of monthly crystallization is significantly lower than that under annual crystallization. This illustrates the fact that a higher crystallization frequency lowers the NAV on which funds can charge the management fee, since an incentive fee payment lower the investor's NAV. However, in economic terms this effect is small. As such, it is more than offset by the higher fee load that results from the higher incentive fees paid.

Next, we have a look at the distribution of the differences in fee loads. From the above analysis, we collect the set of differences in incentive fee under annual and quarterly crystallization. We do this both for the simulated sample and the bootstrapped sample. The results, reported in Figure 5 , illustrate that a higher crystallization frequency always leads to a higher incentive fee load (which is also what to expect). This is evident from the fact that all obtained differences are nonzero. Consequently, the distribution is highly

[9]While most of these occurrences will correspond to fund terminations due to bad performance, we nevertheless treat the fund's exit as full redemption. If there is a positive accrued interest fee at the time of the last observation, it will be charged to the investor's account.

skewed to the right[10]. The Figure also shows that in approximately 35.63% (41.77%) of the simulated (bootstrapped) cases, the two crystallization frequencies do not show any difference in fee load. This is the case whenever (a) a fund does not get over its initial high-water mark, (b) when new highs are reached but not crystallized and (c) when the fund sets new high-water marks at every crystallization date.

[Figure 5 about here.]

In the first two instances, investors only pay the management fee, which is the same for both crystallization frequencies. Of course, investors invest with a positive view on the investment's future performance. An unintended consequence of a higher crystallization frequency is therefore that the investors will pay more (i.e. there will be a positive difference in the fee load) at times when investors are generally less satisfied with the fund's performance.

To see this, consider the following case. When a fund manager, during a particular year, performs very well and continuously sets new highs until the end of the calendar year, it does not matter what crystallization frequency is applied. However, in cases where the fund's NAV at year-end drops below a high-water mark set during the year – the difference in fee load under different crystallization frequencies will be positive. In those cases, investors will be paying higher fees while at the same time the fund's newly crystallized high-water mark will actually be above the NAV at the end of the year (i.e. a drop in NAV). This makes it clear that a higher crystallization frequency will tend to decrease the fund manager's investment horizon and lower the incentive to perform subsequent to the crystallization.

When we condition on those bootstrapped cases where an incentive fee is actually payable, the difference in incentive fee load is 78 basis points higher under quarterly crystallization, as compared to annual crystallization. Comparing this result to the unconditional average, a 49 basis points difference, suggests that in those cases that investors actually pay an incentive fee, the fee load will be higher than our main results would suggest.

[10]This particular distribution is also the reason is why all tests of statistical significance are done using an empirical t-distribution (bootstrap).

5.4 Trade-off between Incentive Fee and Payment Frequency

Thus far, we have assumed a standard 2/20-fee structure to analyse the impact of different payment frequencies. The analysis has shown that, when investors want to compare the (expected) fee load between different funds, such a comparison will be inaccurate if funds differ in terms of the incentive fee payment frequency. In this subsection, we quantify the trade-off that exists between the incentive fee and the crystallization frequency, keeping fixed the level and payment frequency of the management fee. This trade-off might be relevant if the crystallization frequency and incentive fee level are considered negotiable factors.

To ensure that our obtained estimates of the fee load are close to what an investor can expect in reality, the figures are also based on the block bootstrap outlined above. In particular, we calculate the fee load for 10,000 randomly drawn three-year sample paths of gross returns and vary the crystallization frequency and the incentive fee level.

Table 8 reports the size of the effect for different combinations of both negotiable factors. Unlike what incentive fee headline levels would suggest, the table illustrates that changes in the crystallization frequency lead to considerable differences in total fee load. For example, the results suggest that a 15% incentive fee with monthly crystallization leads to a similar total fee load as a 20% incentive fee with annual crystallization (not significantly different).

[Table 8 about here.]

6 Robustness Checks

We now perform a number of robustness checks with regard to the level of the effect. Relaxing or imposing additional restrictions on the dataset used in the analysis will not change our finding that higher crystallization frequencies increase investors' fee load. However, it might have an influence on level of the fee loads and the economic significance of the effect of crystallization.

6.1 Impact of Backfill Bias

In our baseline analysis we account for backfill bias by discarding the first twelve observations of a fund's track record. Here we investigate the importance of this assumption for our baseline results.

To this end, we perform the following analysis. We redo the bootstrap analysis used in section 5.3 a 100 times, both for the baseline gross return data set and the newly obtained gross return data that does not correct for backfill bias. Then, we test whether the results in both cases differ significantly. Panel A of Table 9 reports the result. In line with our expectations, we find that a potential backfill bias tends to upward bias the obtained incentive fee loads. Nevertheless, the size of the difference in fee loads remains similar in both instances, both in magnitude and statistically significance (not reported).

[Table 9 about here.]

6.2 Impact of Fund Size

Another possible concern, raised by Kosowski, Naik, and Teo (2007), is that funds with assets under management below MUSD 20 might be too small for many institutional investors. To ensure that the magnitude of fee load differences is representative and do not deviate too much from the fee load institutional investors can expect, we perform the following robustness check.

Similar to the previous robustness check, we redo the bootstrap analysis a 100 times, but impose an additional restriction when selecting a sample path. In particular, we only select a sample path if – at the start – the corresponding fund's assets under management are above MUSD 20. To avoid look-ahead bias, the fund's size is allowed to drop below MUSD 20 in subsequent months. Results are reported in panel B of Table 9. Consistent with the finding that small funds tend to outperform more mature funds, we find that the fee load is lower when we omit smaller funds.

6.3 Impact of Risk-taking Behaviour

To perform the bootstrap in the baseline case, we assume that every sample path drawn from the gross return dataset starts in January. However, Nanda and Aragon (2012) show that hedge funds take part in tournament behaviour. Hedge funds tend to increase their risk-profile in the second half of the year when they are underperforming, relative to their peers. As such, the funds' risk-profile could differ throughout the calendar-year, and thus have an impact on our reported fee loads. To check whether this is the case, we redo the bootstrap and select sample paths that correspond to actual calendar-years.

The results are reported in panel C of Table 9. The p-values in Panel C indicate that in most cases, the total fee load is somewhat higher if we use actual calendar-years. We interpret this finding as being in line with the results of Aragon and Nanda (2012) their findings on risk-taking behaviour among hedge funds. Our results indicate that, if we take into account intra-year patterns in the funds' returns, we find higher total fee loads. This result therefore suggests that funds actively change their exposure to safeguard accrued incentive fees, causing our results to exhibit slightly higher fee loads if we take these intra-year patterns into account.

7 Conclusion

The fee load of investors does not depend on the headline fee levels alone. Other aspects of the fee structure should also be considered when analysing fee structures that include incentive fees and a high-water mark provision. One such factor is the frequency with which hedge funds update their high-water mark.

To our best knowledge we are the first to document the impact of the crystallization frequency on hedge funds' fee loads. Using simulations and a bootstrap based on a comprehensive data set of CTAs, our main finding is that, under a 2/20-fee structure, quarterly crystallization leads to a fee load which is on average 49 basis points p.a. higher than under annual crystallization. This difference is economically large and should be a relevant consideration when discussing the fee structure. Our results are relevant for allocators who want to assess the fee load of fee schemes which differ in terms of crystallization frequency. Moreover, we find that different headline fee levels can lead to similar total fee loads, once the crystallization frequency is taken into consideration.

In addition, a failure to take into account the frequency with which the high-water mark is updated leads to erroneous estimates of funds their gross returns. In particular, assuming an annual payment of the incentive fee when the industry standard of a number of hedge fund categories is akin to quarterly crystallization, will lead to the underestimation of the gross returns of those hedge fund categories. As such, while annual crystallization might be common among some hedge fund categories, we document that quarterly crystallization is the most common crystallization frequency among CTAs.

Our analysis of the crystallization frequency suggests several avenues for future research. First, we did not go into the implications of the payment frequency on the risk-taking behaviour of hedge funds and CTAs. Changes in the crystallization frequency alter the horizon over which the implications of the high-water mark on risk-taking behaviour should be evaluated. As such, it can be expected that a higher crystallization frequency leads to a shorter trading horizon, and thus might conflict with a fund's stated strategy horizon. Second, we only cover one hedge fund category. As such, there might be considerable differences in the crystallization frequencies applied by different hedge fund categories. These differences might be related to hedge fund characteristics such as the liquidity of the strategy.

References

Agarwal, V., N. D. Daniel, N. Y. Naik, 2009. Role of Managerial Incentives and Discretion in Hedge Fund Performance. The Journal of Finance 64(5), 2221–2256.

Aragon, G. O., V. Nanda, 2012. Tournament behavior in hedge funds: high-water marks, fund liquidation, and managerial stake. Review of Financial Studies 25(3), 937–974.

Bollen, N. P., R. E. Whaley, 2009. Hedge Fund Risk Dynamics: Implications for Performance Appraisal. The Journal of Finance 64(2), 985–1035.

Brooks, C., A. Clare, N. Motson, 2007. The gross truth about hedge fund performance and risk: the impact of incentive fees. Available at SSRN.

Feng, S., M. Getmansky, N. Kapadia, 2011. Flows: The 'invisible hands' on hedge fund management. in *Midwest Finance Association 2012 Annual Meetings Paper*.

French, K. R., 2008. Presidential Address: The Cost of Active Investing. The Journal of Finance 63(4), 1537–1573.

Goetzmann, W. N., J. E. Ingersoll, S. A. Ross, 2003. High-Water Marks and Hedge Fund Management Contracts. The Journal of Finance 58(4), 1685–1718.

Hodder, J. E., J. C. Jackwerth, 2007. Incentive Contracts and Hedge Fund Management. Journal of Financial and Quantitative Analysis 42, 811–826.

Ibbotson, R. G., P. Chen, K. X. Zhu, 2011. The ABCs of hedge funds: alphas, betas, and costs. Financial Analysts Journal 67(1), 15–25.

Jorion, P., C. Schwarz, 2013. Are hedge fund managers systematically misreporting? Or not?. Journal of Financial Economics (0), –.

Kosowski, R., N. Y. Naik, M. Teo, 2007. Do hedge funds deliver alpha? A Bayesian and bootstrap analysis. Journal of Financial Economics 84(1), 229–264.

Kouwenberg, R., W. T. Ziemba, 2007. Incentives and risk taking in hedge funds. Journal of Banking and Finance 31(11), 3291–3310.

Panageas, S., M. M. Westerfield, 2009. High-Water Marks: High Risk Appetites? Convex Compensation, Long Horizons, and Portfolio Choice. The Journal of Finance 64(1), 1–36.

Teo, M., 2009. The Geography of Hedge Funds. Review of Financial Studies 22(9), 3531–3561.

Appendix: Description Algorithm for Gross Returns

Here we describe the algorithm we use to compute monthly gross returns from reported net-of-fee returns. Our approach allows for a monthly estimation of gross returns under different crystallization regimes (monthly or lower frequency).

The algorithm is based on the following set of assumptions:

1. The Gross Asset Value at the fund's inception (GAV_0) is equal to 100.
2. The algorithm is based on a single-investor assumption.
3. The management fee is paid monthly[11].

The Gross Return at time t ($GrossRet_t$) equals:

$$GrossRet_t = \frac{GAV_t}{GAV_{t-1}} - 1 \tag{1}$$

where GAV_t and GAV_{t-1} are the Gross Asset Value at month t and $t-1$, respectively.

The Management Fee ($MgtFee_t$) paid in month t equals:

$$MgtFee_t = \frac{GAV_t}{GAV_{t-1}} \cdot NAV_{t-1} \cdot \frac{MF\%}{12} \tag{2}$$

where $MF\%$ is the management fee (p.a.). The Total Management Fee Paid up to month t ($TotalMgtFeePaid_t$) then equals

$$TotalMgtFeePaid_t = \sum_{i=1}^{t} MgtFee_i \tag{3}$$

Interest Earned ($InterestEarned_t$) equals the amount of interest earned by the manager on the excess cash and cash deposited in the margin account:

$$InterestEarned_t = NAV_{t-1} \cdot Rf_t \tag{4}$$

[11] This assumption can easily be relaxed to a different payment frequency by handling the payment of the management fee in the same way as the incentive fee. We nevertheless fix the payment frequency to monthly because an analysis of the managment fee is not the thrust of the analysis.

where Rf_t is the risk-free rate in month t. Total Interest Earned ($TotalInterestEarned_t$) is the sum of all interest earned on cash up to month t:

$$TotalInterestEarned_t = \sum_{i=1}^{t} InterestEarned_i \tag{5}$$

Using the above definitions, the Preliminary Net Asset Value at time t ($PrelNAV_t$) is then:

$$PrelNAV_t = \frac{GAV_t}{GAV_{t-1}} \cdot NAV_{t-1} - TotalMgmFeePaid_t - TotalIntEarned_t \tag{6}$$

We substract the management fee and the interest earned from the gross asset value. That way the manager will only earn an incentive fee on performance in excess of any management fee charged or any risk-free return earned on cash[12]. For the next set of equations, we introduce an indicator ($Cryst_t$) that takes on the value 1 if a crystallization occurs, and zero otherwise.

The Accrued Incentive Fee ($AccrIncFee_t$) is a percentage of the performance – the incentive fee $IF\%$ – in excess of the current high-water mark (HWM_{t-1}):

$$\begin{cases} max(0, PrelNAV_t - HWM_{t-1}) \cdot IF\% & \text{if } Cryst_t = 0 \\ 0 & \text{if } Cryst_t = 1 \end{cases} \tag{7}$$

Therefore, the incentive fee is accrued when no crystallization occurs and will equal zero when crystallization takes place. In that case, the incentive fee is paid to the fund manager. The accrued incentive fee over the period since the last crystallization is then added to the Incentive Fee Paid variable ($IncFeePaid_t$):

$$\begin{cases} max(0, PrelNAV_t - HWM_{t-1}) \cdot IF\% + IncFeePaid_{t-1} & \text{if } Cryst_t = 1 \\ IncFeePaid_{t-1} & \text{if } Cryst_t = 0 \end{cases} \tag{8}$$

The High-Water Mark at time t (HWM_t) is updated to the current preliminary Net

[12] We take this into consideration because CTAs typically hold up to 80% of the money in a cash account.

Asset Value *if* crystallization occurs and is not updated if no crystallization occurs:

$$\begin{cases} max(PrelNAV_t, HWM_{t-1}) & \text{if } Cryst_t = 1 \\ HWM_{t-1} & \text{if } Cryst_t = 0 \end{cases} \tag{9}$$

The Net Asset Value at time t (NAV_t) equals:

$$NAV_t = PrelNAV_t + TotalInterestEarned_t - AccrIncFee_t - IncFeePaid_t \tag{10}$$

Since no closed-form solution is available, we solve for GAV_t numerically. In particular, we determine the value of GAV_t that equates the NAV_t computed in equation (10) – based on GAV_t – to the observed NAV at time t. We then store the obtained value of GAV_t and move to the next month, calculating GAV_t in an iterative way. When we charge fees in the subsequent analysis, we also use the above equations to go from GAV_t to NAV_t.

Figure 1: Distribution of the Crystallization Frequencies of the Incentive Fee

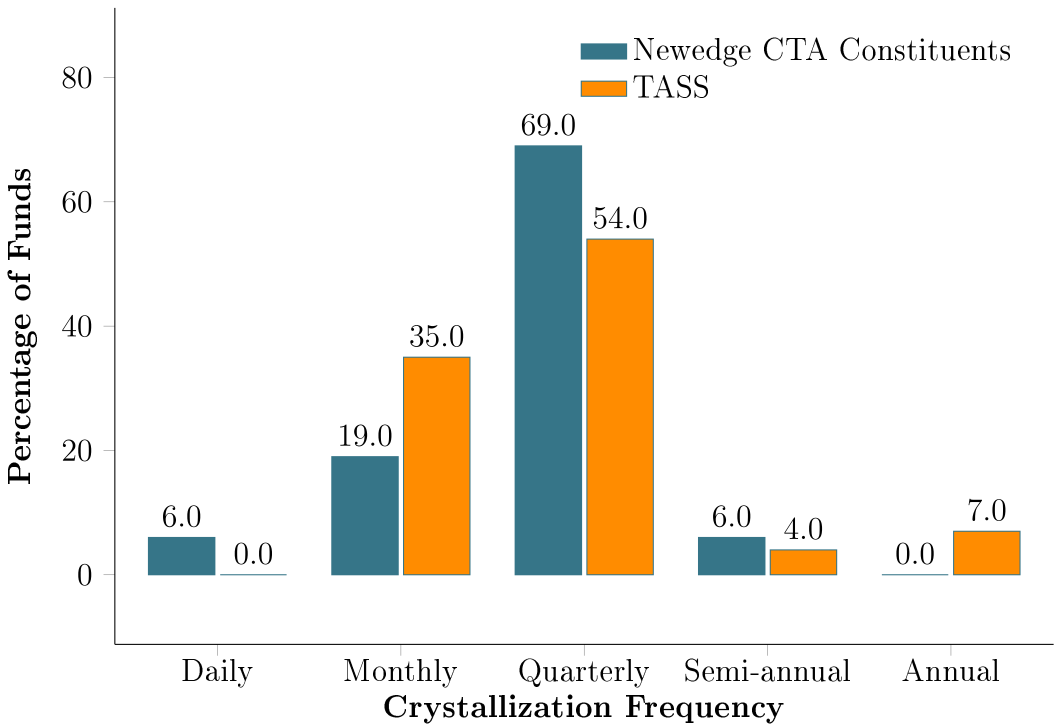

Based on survey conducted in May 2013 and Tremont Advisory Shareholder Services, 2012. For the Newedge CTA index, 4 funds did not disclose their payment frequency. In the case of TASS, the fee notes of 185 funds (out of a sample of 408 fee notes) contained a sufficient amount of information to determine the payment frequency of the incentive fee.

Figure 2: Scope of the Survey

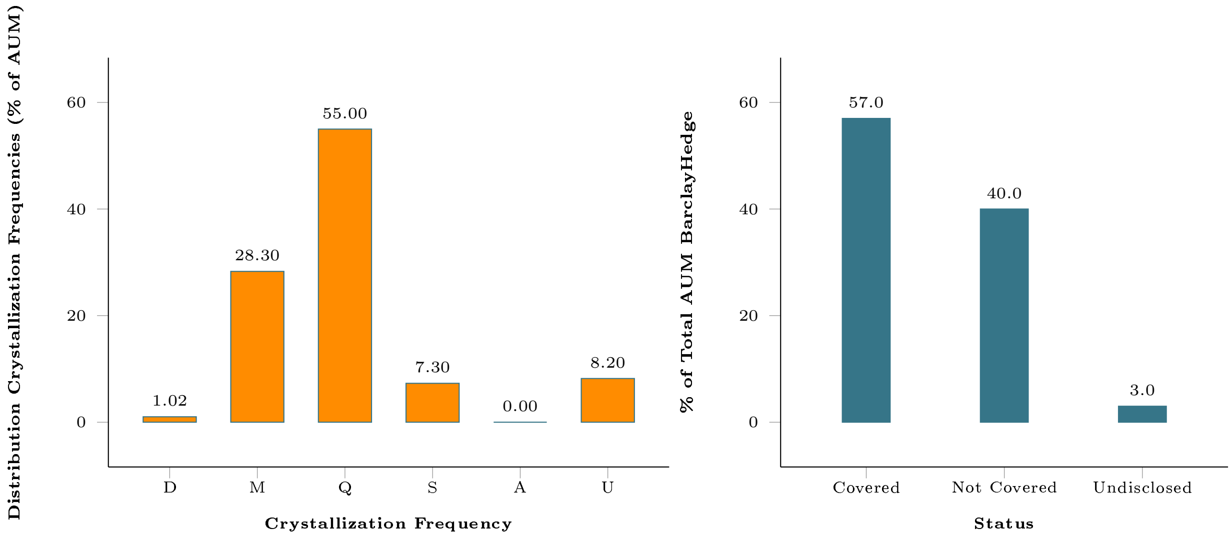

Based on survey results and the programs' AUM reported to BarclayHedge, as of May 2013.

Figure 3: Total Fee Load under Different Crystallization Frequencies

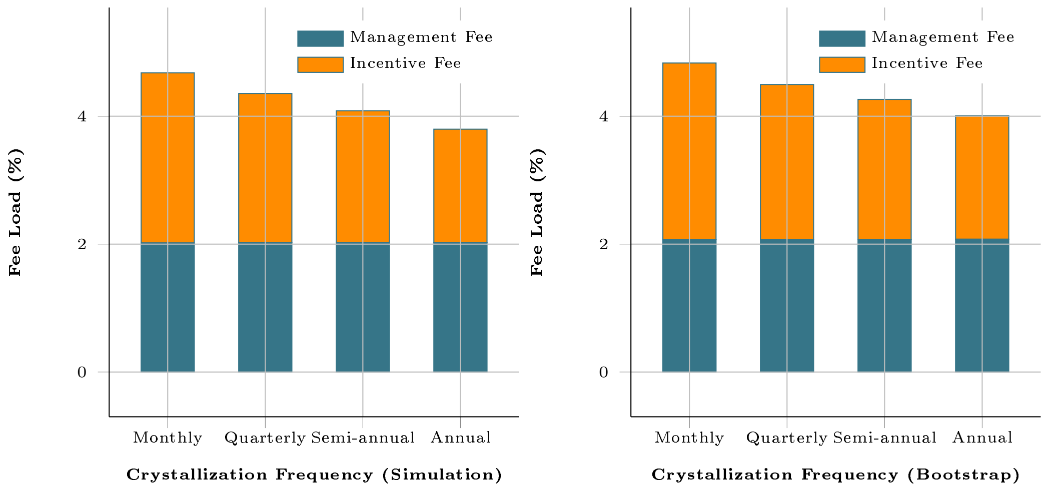

Figure 4: Comparing the Total Fee Load with Annual Crystallization

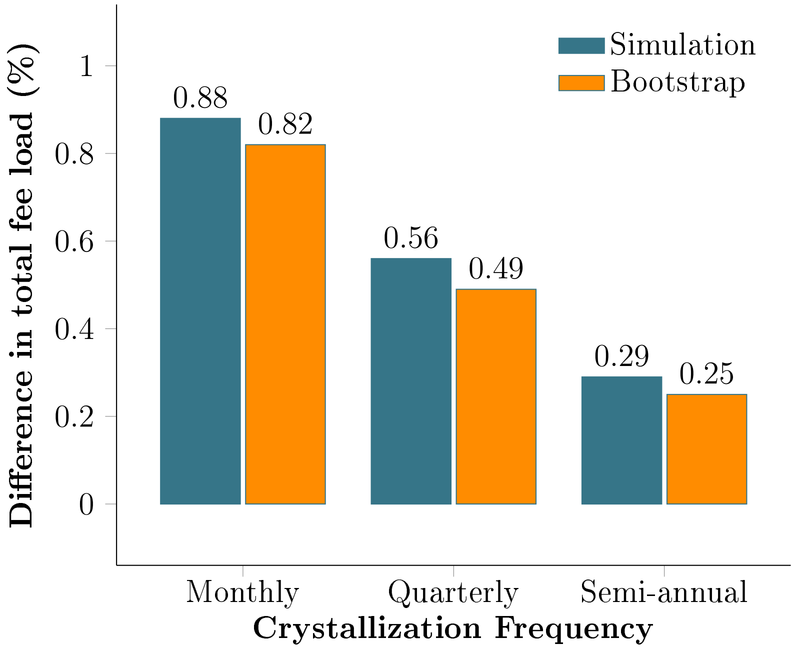

Figure 5: Distribution of the Difference in Incentive Fee Load – Quarterly Versus Annual Crystallization for One-Year Horizon

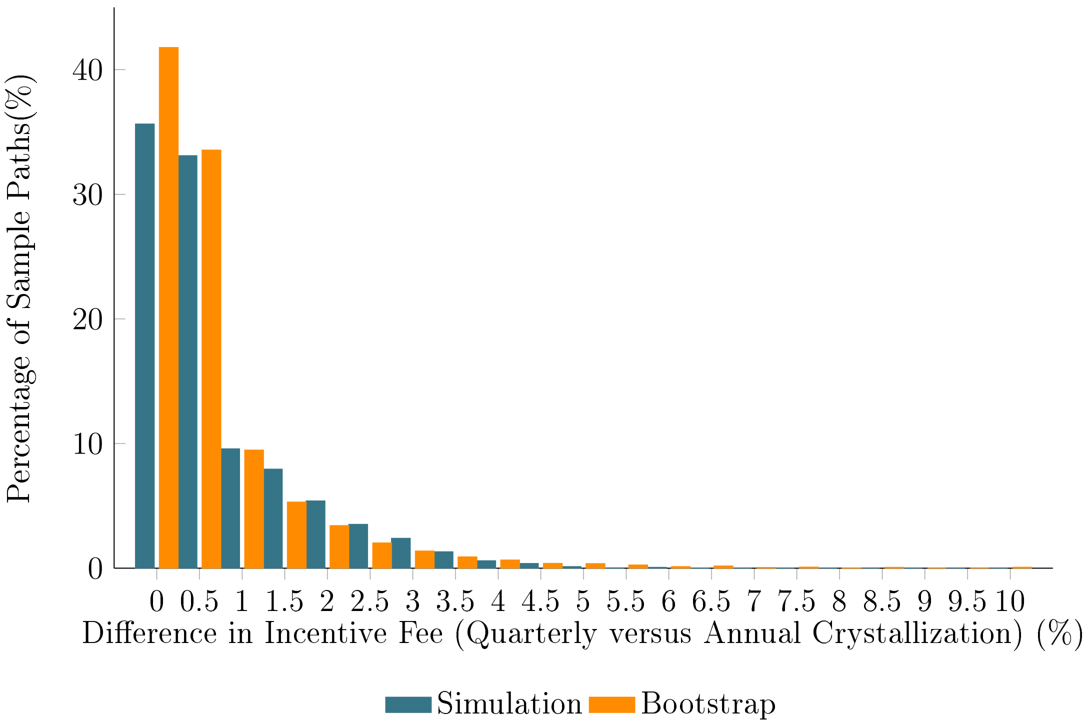

Table 1: Evolution in CTA Headline Fee Levels

	Number of funds	Management fee	Bootstrapped 95% CI	Incentive fee	Bootstrapped 95% CI
1994-1998	295	1.97%	[1.88%;2.07%]	20.63%	[20.29%;20.98%]
1999-2003	394	1.71%	[1.65%;1.78%]	20.51%	[20.23%;20.83%]
2004-2008	377	1.67%	[1.6%;1.73%]	20.71%	[20.3%;21.16%]
2009-2012	163	1.62%	[1.51%;1.72%]	20.64%	[19.91%;21.42%]

This table reports the number of launched funds, average incentive fee and management fee for CTAs in BarclayHedge for different sub-periods.

Table 2: Summary Statistics CTAs

	Mean	Min	P25	P50	P75	Max
Monthly net-of-fee return	0.567%	-6.474%	0.061%	0.505%	0.988%	9.524%
Standard deviation of monthly net-of-fee returns	5.084%	0.610%	2.753%	4.273%	6.586%	17.173%
Age (years)	5.4	1.0	2.1	3.8	7.0	19.0
Management fee	1.869%	0%	2%	2%	2%	5%
Incentive fee	20.561%	5%	20%	20%	20%	50%

This table reports summary statistics for the sample of 1616 CTAs from the BarclayHedge database.

Table 3: Relationship between Crystallization Frequency and Fee Levels

Crystallization Frequency	Incentive Fee	Bootstrapped 95% CI	Management Fee	Bootstrapped 95% CI
Monthly	22.38%	[20.72%;24.23%]	1.63%	[1.36%;1.91%]
Quarterly	21.05%	[20.35%;21.8%]	1.64%	[1.48%;1.79%]
Semi-annual	20.00%	[20%;20%]	1.93%	[1.79%;2%]
Annual	19.62%	[17.69%;21.15%]	1.47%	[1.17%;1.81%]

This table reports the average incentive fee level and management fee level under different crystallization frequencies for sample of CTAs in TASS.

Table 4: Comparison of Net-of-fee Returns and Gross Returns

	Net-of-fee returns	Gross returns	p-value
Average return	0.567%	0.768%	0
Standard deviation of monthly returns	5.084%	4.683%	1.00E-04
Annualized Sharpe Ratio	0.483	0.691	0
Skewness	0.313	0.454	0
Kurtosis	4.820	5.126	0.0133
First order serial autocorrelation	-0.011	-0.004	0.1382
JB-Statistic (Percentage of rejections)	47.215%	52.228%	

This table compares net-of-fee returns with the estimated gross returns based on the algorithm described above for the set of 1616 CTAs.
The reported p-values test the difference in means using the empirical t-distribution (bootstrap).

Table 5: Summary Statistics historical fee-loads

	Average	Standard Deviation	Sharpe Ratio
Gross Return	8.65%	16.22%	0.61

	Net-of-fee Return	Standard Deviation	Sharpe Ratio	Management Fee	Incentive Fee
Monthly	4.90%	16.75%	0.34	1.933%	2.411%
Quarterly	5.07%	16.33%	0.37	1.934%	2.260%
Semi-annual	5.20%	16.05%	0.38	1.934%	2.159%
Annual	5.32%	15.75%	0.40	1.935%	2.141%

This table reports the average annual gross return, average standard deviation and average Sharpe ratio for the set of 1616 CTAs. The second part of the table reports the corresponding statistics for the net-of-fee returns, as well as the average management fee and incentive fee.

Table 6: Impact of Crystallization on Fee Load

Panel A

Crystallization Frequency		Incentive Fee	Management Fee	Total Fee Load
1-year horizon	Monthly	2.661%***	2.018%***	4.679%***
	Quarterly	2.334%***	2.022%**	4.356%***
	Semi-annual	2.061%***	2.024%	4.085%***
	Annual	1.771%	2.026%	3.797%
3-year horizon	Monthly	1.913%***	2.029%*	3.942%***
	Quarterly	1.770%***	2.030%	3.800%***
	Semi-annual	1.649%***	2.030%	3.680%***
	Annual	1.506%	2.031%	3.538%
5-year horizon	Monthly	1.671%***	2.035%	3.706%***
	Quarterly	1.573%***	2.035%	3.609%***
	Semi-annual	1.492%***	2.036%	3.528%***
	Annual	1.390%	2.035%	3.426%

Panel B

Crystallization Frequency		Incentive Fee	Management Fee	Total Fee Load
1-year horizon	Monthly	2.762%***	2.070%***	4.836%***
	Quarterly	2.423%***	2.073%	4.500%***
	Semi-annual	2.188%***	2.075%	4.266%***
	Annual	1.932%	2.077%	4.012%
3-year horizon	Monthly	2.063%***	2.059%	4.127%***
	Quarterly	1.862%***	2.060%	3.926%***
	Semi-annual	1.728%***	2.061%	3.793%***
	Annual	1.606%	2.061%	3.671%
5-year horizon	Monthly	1.839%***	2.050%	3.894%***
	Quarterly	1.667%***	2.0511%	3.723%***
	Semi-annual	1.554%***	2.052%	3.670%***
	Annual	1.440%	2.052%	3.4960%

This table reports the average incentive fee, average management fee, and average total fee load for two analyses. Panel A reports the results for a simulation of funds' gross returns, under the assumption of normal distribution, with a mean of 0.768% per month and a standard deviation of 4.683% per month. Panel B shows the results from performing a block bootstrap where 12, 36, or 60 month blocks of gross returns are drawn from the obtained sample of CTAs. Fee load equals the average annual fee load over the investment horizon, as a percentage of initial NAV/NAV at the end of the previous year.

Asterisks report statistically significance of the difference between of the obtained fee levels and the benchmark category (annual crystallization) at the 10% (*), 5% (**) and 1% (***) level of significance. Significance tests based on the empirical t-distribution (bootstrap).

Table 7: The Impact of Volatility on Fee Load

Crystallization Frequency				
10% volatility		Incentive Fee	Management Fee	Total Fee Load
	Monthly	1.365%***	2.034%	3.398%***
	Quarterly	1.282%***	2.034%	3.316%***
	Semi-annual	1.214%***	2.035%	3.248%***
	Annual	1.131%	2.035%	3.166%
20% volatility				
	Monthly	2.241%***	2.027%*	4.268%***
	Quarterly	2.056%***	2.029%	4.085%***
	Semi-annual	1.895%***	2.030%	3.924%***
	Annual	1.707%	2.030%	3.737%
30% volatility				
	Monthly	3.159%***	2.018%*	5.177%***
	Quarterly	2.856%***	2.020%	4.877%***
	Semi-annual	2.606%***	2.022%	4.628%***
	Annual	2.312%	2.022%	4.335%

This table reports the average incentive fee, management fee, and total fee load for different levels of volatility, keeping the expected return constant. We reports the results for a simulation of funds' gross returns, under the assumption of normal distribution, with a mean of 0.768% per month and a standard deviation of 10%, 20%, and 30% p.a. respectively. We report the results for a 3-year investment horizon. The fee load equals the average annual fee load over the investment horizon, as a percentage of initial NAV/NAV at the end of the previous year.
Asterisks report statistically significance of the difference between of the obtained fee levels and the benchmark category (annual crystallization) at the 10% (*), 5% (**) and 1% (***) level of significance. Significance tests based on the empirical t-distribution (bootstrap).

Table 8: Trade-off between Crystallization frequency and Incentive fee

	Incentive fee (%)					
Crystallization Frequency	5	10	15	20	25	30
Monthly	2.574%	3.071%	3.599%	4.083%	4.610%	5.238%
Quarterly	2.527%	2.970%	3.456%	3.883%	4.359%	4.937%
Semi-annual	2.495%	2.908%	3.364%	3.753%	4.196%	4.731%
Annual	2.463%	2.844%	3.264%	3.622%	4.033%	4.528%

This table reports the total fee load under different combinations of the both negotiable factors, the incentive fee level and the crystallization frequency. The management fee is paid monthly and fixed at 2% p.a. The fee load is estimated by drawing random three-year sample paths from the gross CTA return data and calculating the fee load, varying the crystallization frequency and the level of the incentive fee.

Table 9: Results Robustness Checks

Robustness check		Baseline result	Result under robustness check	p-value
Backfill Bias	Monthly	4.107%	4.379%	0
	Quarterly	3.913%	4.170%	0
	Semi-Annual	3.784%	4.031%	0
	Annual	3.656%	3.893%	0
Fund Size	Monthly		3.653%	0
	Quarterly		3.487%	0
	Semi-Annual		3.373%	0
	Annual		3.258%	0
Risk-taking Behaviour	Monthly		4.107%	0.4827
	Quarterly		3.920%	0.0701
	Semi-Annual		3.792%	0.0408
	Annual		3.710%	0

This table reports the total fee load for a three-year investment horizon for the baseline case, and a set of three robustness checks.
The reported p-values test the difference in means using the empirical t-distribution (bootstrap).