Pricing of Guaranteed Stop Orders

Jonathan Josef Leicht

Andreas Rathgeber

ABSTRACT

In this article we analyse the behaviour of guaranteed stop orders on stocks in the German stock index DAX. We explain briefly how guaranteed stop orders work and then develop a jump process based on a variance gamma process to model the stock prices. We show, through simulations, that the pay-off of a guaranteed stop order is foremost governed by the volatility in the underlying stocks' overnight movements. We also demonstrate that the common linear approach to price guaranteed stop orders is too general and needs to be refined to represent the differences between stocks adequately. We show that the recent turbulence on stock markets around the world has made the guaranteed stop order interesting and that, in more tranquil periods, this order type was nearly irrelevant.

INTRODUCTION

There are many different forms of portfolio insurance, all with certain advantages and disadvantages. A common one among them is the protective put option. It is a simple tool to avoid losses below a certain barrier, but has the drawback of not providing continuous protection *(see* Bodie, Kane and Marcus, 1996). To ensure that a portfolio never drops below a certain level, dynamic insurance tools like synthetic puts are used. They have the drawback of requiring continuous trades and, therefore, cost the investor a considerable effort to use them. Additionally, Basseer [1991] has found dynamic portfolio insurance to work best in orderly markets, but it becomes impractical during periods of high volatility. Basic tools to avoid losses are risk management orders, whereby there most comon resprentative is the stop or stop loss order (SO). Unlike most other tools, the investor does not pay a premium when placing SOs, making it a widespread method. One problem with SOs as protection is that they do not guarantee a selling price equal to their barrier. If liquidity is tight and it is, therefore, difficult to find a counterparty willing to buy, the selling price may drop significantly. Many studies have shown that SOs can reinforce sudden stock price drops, a phenomenon often referred to as 'price cascades' (see Easley and O'Hara, 1991; Genotte and Leland, 1990). These price cascades occur when many investors have set stop orders with similar barriers and a price drop triggers them simultaneously. These sudden sell orders instigate other market participants to sell or even force investors to sell owing to loss limits. To counter this vulnerability to sudden stock price jumps, the guaranteed stop order (GSO) was created. It is basically a SO with the additional benefit of guaranteeing a selling price equal to the barrier. If an investor uses a GSO instead of a SO, he will still contribute to this problem, but he has the distinct advantage of not being affected by it.

In addition, a GSO is not always superior to a simple SO since the investor pays a premium in order to be insured against price jumps. Often, the stock price jumps are neglected when looking at SOs and for stocks on major firms this may be adequate in tranquil and orderly times. The recent financial crisis, however, has caused considerable turbulence and dramatically increased stock price volatility. Under these conditions we deem it necessary to take a closer look at GSOs and evaluate if their pricing is justified.

The purpose of this article is to shed some light on the value of GSOs, an instrument that has received very little scientific attention in the past. A jump process is used to account for the discontinuities of the stock market, which are the sole reason for the existence of GSOs. We will establish through stock price simulations if the method of pricing GSOs as used by CMC Markets [2005] is adequate and reflects their true value. We will analyse which factors determine the value of a GSO and how to price them accordingly.

This article is organized as follows: In the next two sections, a model for stock prices and GSOs is developed. Then the data is presented and the fitting process described. Afterwards the simulation process is explained, followed by its results. The next section provides a closer analysis, determines the major influences on GSOs, and outlines a new pricing approach. The following section then compares the previous findings to an older time period when the stock market was more tranquil. Finally, the last section concludes this article by summarizing the results and giving an outlook on possible future research on GSOs.

MODEL

In order to evaluate GSOs we must first decide on a model for the underlying stocks. The most common approach would be to assume that the stock prices follow a Brownian motion. One of the characteristics of a Brownian motion is the fact that its paths are almost surely continuous. To put it in simpler terms, a Brownian motion does not jump. On the other hand, GSOs are instruments specifically developed to protect investors against jumps in stock markets. It would, therefore, be desirable to use a stochastic process which focuses on such events. According to Cont and Tankov [2004] there are two basic categories of jump processes to choose from: jump-diffusion models and infinitely active models. The former consist of a Brownian component and rare jumps, the latter of an infinite number of jumps in each interval. In this paper, we will use a Variance Gamma Process (VGP), a process of infinite activity, as a starting point for our model. Carr, Geman, Madan and Yor [2003], Geman [2002] and Madan [2001] support that this category provides a better representation of historical stock price processes.

In the basic model *(compare* Schoutens, 2003; Cont and Tankov, 2004; Senata, 2004) the stock price S_t over time $t \geq 0$ is given by,

$$S_t := S_0 \exp(X_t) \tag{1}$$

where S_0 is the initial stock price and the exponent X_t is defined as:

$$X_t := ct + \theta G_t + \sigma W(G_t) \tag{2}$$

with the parameters c, θ, σ and the Brownian motion W. G_t is a random process independent of the Brownian Motion W(·). G_t follows a Gamma process and can be interpreted as an 'economically relevant measure of time' (Geman, Madan and Yor, 2001). The expected economically relevant time change per calendar time unit is, without loss of generality, normalized to one.

$$E(G_t - G_{t-1}) = 1 \tag{3}$$

This normalization is contained within θ and σ. The change in the exponent X_t over one unit of calendar time, denoted by ΔX_t, can be written as:

$$\Delta X_t := X_t - X_{t-1} = c + \theta(G_t - G_{t-1}) + \sigma\sqrt{G_t - G_{t-1}}W(1) \tag{4}$$

The model, as it stands now, demands that the expected price change from one time step to the next is constant. For GSOs, however, it is important to account for the fact that the expected change between the last stock price of any given day and the next, will not be the same as between two adjacent intraday prices. This is accounted for by defining S_t as follows:

$$S_t := S_0 \exp(ID_t + ON_t) \tag{5}$$

In this definition, there are two independent stochastic processes ID_t and ON_t. ID_t controls the intraday stock movement and ON_t the overnight movement. This model has the underlying assumption that each jump is independent of previous jumps, especially that overnight jumps are independent of intraday jumps. This assumption is not entirely true. There is evidence that an over or under performance in the last few stock prices of a day is related to an over or under performance in the overnight jumps. This dependence lends itself to further study but will be ignored in this model. To be able to formally define ID_t and ON_t, we create a set T containing all points in time which coincide with the first price fixing of each day. The change in ID_t is now defined similar to ΔX_t in (4).

Formatiert: Englisch (USA)

Gelöscht: (4)

$$\Delta ID_t := \begin{cases} 0, \text{ if } t \in T \\ c + \theta G_t - G_{t-1} + \sigma_{ID}\sqrt{G_t - G_{t-1}}W(1), \text{ otherwise} \end{cases} \tag{6}$$

This definition shows that the intraday movement follows a VGP that does not affect the overnight movement. ON_t adds a normally distributed jump between the closing and the opening price but does not affect intraday movement and is defined as follows:

$$\Delta ON_t := ON_t - ON_{t-1} \begin{cases} \sim N(\mu, \sigma_{ON}^2), & \text{if } t \in T \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

GUARANTEED STOP ORDERS

As mentioned earlier, GSOs are similar to standard stop orders. The difference is that a GSO guarantees a selling price equal to the chosen barrier, while the stop order only provides the next selling price after the barrier has been broken. So basically, a GSO can be seen as an insurance that pays the difference between the barrier and the next possible selling price. We evaluate GSOs using the stock price model (5). The pay-off P_B is defined as,

Formatiert: Englisch (USA)

Gelöscht: (5)

$$P_B := |B - S_t| \tag{8}$$

assuming that t is the first point in time for the stock price S_t to be less or equal to the barrier B or, if short selling, to exceed the barrier. The model considers a time frame of one year between GSO placement and cancellation. Should the barrier not be reached in this time frame the pay-off P_B is zero. It is extremely uncommon to place (guaranteed) stop orders for a longer period. Many brokerage services do not even permit it. When placing a GSO, unlike a SO, an insurance premium has to be paid to reflect the offered protection. This fee is calculated according to CMC Markets [2005]. For stocks in the German stock index DAX the premium is set to 0.3% of the GSOs' barrier by CMC and has to be paid immediately. Stocks in other indices have different rates, but follow the same structure. The premium R is therefore defined as,

$$R := rB \tag{9}$$

with r = 0.003 in our cases. This definition is intuitive for long positions, but seems counterintuitive for short positions, where the premium rises further away from the initial stock price at which the barrier is set. This structure is probably due to the fact that the expected relative jump sizes are assumed to be constant and the absolute jumps therefore expected to increase at higher stock prices. An additional rule is that a GSO cannot be set within 5% of the current price of a stock. So, for example, if a stock is listed at 100 one could set a GSO at 95 and pay a premium of 0.285. CMC Markets allows its customers to trade in contracts for difference (CFD) instead of directly trading stocks, but the CFD prices are adjusted to the underlying stocks. Considering the short time frame interest rates are neglected and the GSO value is defined as the difference between pay-off and premium. In this paper we will focus on long positions, where the barrier is consequently below the initial stock price.

DATA

The next step is to obtain data, to which the model can be fitted. CMC Markets trades in CFDs, for which no time series were obtainable. Since these CFDs reflect stock prices nearly 1:1, we will use stock prices as a proxy. We fit our parameters to the stock prices of the 30 firms in the Deutscher Aktienindex (DAX), the most important German stock index. These stocks are also available as CFDs on which GSOs can be placed. Since we intend to simulate stock prices on a tick level, we require stock price data on the same level. The exchange in Stuttgart provides us with the necessary time series of quotes for each tick on the stock market.

In order to fit the intraday aspect of our model we collect quotes ranging from 1 April, 2009 to 14 May, 2009. With only 29 trading days this may seem to be a very brief period of time, but since we are using tick data we have an average of 1,701 data points per firm. Commerzbank AG shows the highest activity with 7,512 data points. Merck KGaA, the firm with the least activity in this period, still provides 206 data points. Therefore, we believe these 29 trading days to be sufficient for the intraday fit of our model. This time period leaves us with only 28 overnight returns, which does not suffice for an adequate fit of the model's overnight aspect. Since we assume the overnight returns to be independent of the intraday returns, we can extend the time period to gain additional data points. We use overnight returns ranging from 3 June, 2008 to 14 May, 2009. This period amounts to 239 returns, which should allow for an adequate fit. The overnight returns corresponding to dividend payment dates are removed from this series leaving one less data point for most firms. Intraday returns are not directly affected by the dividend payments, therefore none have to be removed.

FITTING THE PARAMETERS

After obtaining the data, the next step is to decide on a method to fit the model to the data. Following Cont and Tankov (2004), there are two main approaches: the method of moments (MoM) and the maximum likelihood estimation (MLE). The MLE approach requires a known distribution function. This is the case for a standard VGP, but since we modified the process to account for overnight jumps, the distribution function is not easily obtainable. We therefore use a MoM approach following Senata (2004) who states: 'While method of moments is lacking somewhat in precision, it compensates in terms of robustness.' The overnight parameters (μ, σ_{ON}) and the intraday parameters (c, θ, κ, σ_{ID}) are fitted separately since they use different time series. The overnight parameters are fitted to the first two moments, the intraday parameters to the first four.

The fitting process for the overnight parameters is trivial since μ is simply the first and σ_A the square root of the second moment.

The first four moments give us following four equations for the intraday parameters (Senata, 2004):

$$E(X_t) = c + \theta \tag{10}$$

$$\operatorname{var}(X_t) = \sigma_{ID}^2 + \theta^2\kappa \tag{11}$$

$$s(X_t) = \frac{2\theta^3\kappa^2 + 3\sigma_{ID}^2\theta\kappa}{\left(\theta^2\kappa + \sigma_{ID}^2\right)^{\frac{3}{2}}} \tag{12}$$

$$k(X_t) = 3 + \frac{3\sigma_{ID}^4 \kappa + 12\sigma_{ID}^2 \theta^2 \kappa^2 + 6\theta^4 \kappa^3}{\left(\theta^2 \kappa + \sigma_{ID}^2\right)^2} \tag{13}$$

In (12) s represents the skewness and in (13) k the kurtosis. In a first step, we attempt to solve equations (10)–(13) using the first four moments estimated from the data. Given the complexity of these equations we are not always able to find a solution. In those cases we use equation (10) together with a second set of equations (Senata, 2004):

Gelöscht: (12)
Formatiert: Englisch (USA)
Gelöscht: (13)
Formatiert: Englisch (USA)
Gelöscht: (13)
Formatiert: Englisch (USA)
Gelöscht: (10)
Formatiert: Englisch (USA)
Formatiert: Englisch (USA)
Gelöscht: (10)

$$\operatorname{var}(X_t) \approx \sigma_{ID}^2 \tag{14}$$

$$s(X_t) \approx \frac{3\theta\kappa}{\sigma_{ID}} \tag{15}$$

$$k(X_t) \approx 3(1+\kappa) \tag{16}$$

Equations (14)–(16) assume that θ is relatively small and therefore θ^2, θ^3, θ^4 can be ignored. Since this set of equations only provides an approximation, we need to check their quality. In order to do this, we insert our approximated parameters into equations (10)–(13), computing the moments implied by our approximations. We then compare these implied moments with the moments estimated from the data and evaluated the errors. Exhibit 1 shows the fitted parameters, exhibit 2 the errors in the implied moments.

Formatiert: Englisch (USA)
Gelöscht: (14)
Formatiert: Englisch (USA)
Gelöscht: (16)
Gelöscht: (10)
Formatiert: Englisch (USA)
Gelöscht: (13)
Formatiert: Englisch (USA)

The first moment is omitted from exhibit 2 since parameter c allows for an exact fit leading to an error of zero. In this table, for example, a value of -5% means that the estimated moment is 5% below the implied moment. We can observe that in all but three cases we stay in a ±5% interval. Fresenius Medical Care AG & Co. KGaA (FME) goes up to 5.74%, Salzgitter AG (SZG) up to 8.89% and Beiersdorf AG (BEI) even to 19.91%. We remove Beiersdorf and Salzgitter from our further observation since the error is too large and simulation results with these parameters cannot be trusted. FME remains in our simulations, but we should keep in mind that the results may be erroneous.

SIMULATION PROCESS

Having fitted the parameters of our model to the data, we must now decide on a simulation process. In order to simulate the price process (5) a time grid has to be chosen. The pay-off of a GSO is determined by the first obtainable CFD price or, in our case, stock price, after the relevant barrier has been broken. Therefore, one calendar time unit is defined as one tick. A problem with this definition is that the number of ticks per day varies for each stock and day. In this model we will assume that for a given stock the number of ticks per day N_S is constant but may be different for different firms. For every stock S the average number of ticks per day is calculated and defined as N_S. To determine the value of the GSOs we run 10,000 simulations for each stock in our sample. During these simulation runs we evaluate GSOs on ten different barrier levels, ranging from 50% to 95% in 5% intervals. We simulate a period of one year, which is assumed to have 250 trading days. If a GSO is not triggered in the course of one year, it is cancelled and has a pay-off of zero. The simulation algorithm is based on Cont and Tankov [2004, p. 184] and is modified to fit our specific model.

Formatiert: Englisch (USA)
Gelöscht: (5)

SIMULATION RESULTS

In this chapter we will visualize and examine the simulation results.

Exhibit 3 shows the GSO pay-off as simulated. Every second GSO barrier was omitted in order to keep the table at a reasonable size. Pay-offs that exceed the costs are bold and underlined. One can observe that at the highest possible barrier over half of the firms (53.6%) have a pay-off higher than the cost. This figure decreases steadily down to 10.7% at the lowest barrier.

Exhibit 4 takes a closer look at the difference between pay-off and cost of the GSOs. The numbers in the table indicate the percentage of firms with a pay-off to cost ratio in the respective interval at a certain barrier. We can observe that, at the highest barrier, half the firms have a pay-off to cost ratio between 75% and 125%. As the barrier decreases, fewer and fewer firms remain at moderate ratios. At the 85% barrier and below, over 20 out of 28 firms have ratios of either less than 50% or over 150%. This clearly shows that the linear approach to pricing GSOs does not reflect their true value. Additionally, we can observe that the pay-off decreases at a significantly higher rate than the cost and, with a barrier of 85% or lower, most are at less than 50% pay-off to cost ratio.

The pay-off of a GSO depends on two factors: the probability that the stock reaches the barrier and the amount by which it breaks the barrier. Exhibit 5 shows averages across all firms for the pay-off, the pay-off under the condition that the barrier is reached (‘pay-off when triggered’), the trigger probability and the percentage of barrier passages caused by overnight jumps. Additionally, the cost of the GSO is shown for comparison. When considering the average pay-off over the DAX firms, only barriers of 80% or higher have a positive return. This is further evidence that the linear pricing structure is not adequate for low barriers. The decline in ‘pay-off when triggered’ is due to the fact, that the expected relative jump size is constant, not the expected absolute jump size. Hence the pay-off when triggered should be proportional to the stock value and, therefore, to the barrier. The simulation results reflect this relation with the exception of the two lowest barriers. Here an increase is observable. This increase can easily be explained when considering that not all firms reach these barriers. Firms with little movement and small jumps will not reach the lower barriers while the ones with much movement and the largest jumps will. Therefore, the weights of the most active firms increase at low barriers and the average is biased upwards. As expected, the probability to reach a barrier goes down as the barrier goes down. The percentages of barrier passages that occur overnight are evidently constant at about 28% and seem to be independent of the barrier. Since the overnight movements are independent of the intraday movements in our model, they are also independent of the stock price. Hence the simulated constant percentage is in accordance with the model.

ANALYSIS

The goal of this chapter is to analyse what determines the value of a GSO and, using this information, how to adequately price GSOs on different stocks. Two important questions arise: Is the value driven by overnight movement or by intraday movement? Is it more important how likely it is to trigger the GSO or how large the price jump will be when it is triggered? Exhibit 6 illustrates the possible cases.

Exhibit 7 compares two typical ‘pay-off when triggered’ histograms. Volkswagen is a company with a very high simulated pay-off, Lufthansa with a very low simulated pay-off. It is obvious that the distribution of pay-offs is very different. Volkswagen's pay-off has much more weight in the tail than Lufthansa's. Lufthansa shows not a single simulation run with a pay-off of more than 4.5%. Volkswagen has several runs with over 10% going up to pay-offs of nearly 18%. The variance in these pay-offs is 5.01 for Volkswagen and 0.26 for Lufthansa. These results are typical for firms with high respectively low pay-off and show that the value of a GSO is strongly driven by the distribution of this conditional pay-off. This indicates that cases (II) and/or (IV) from exhibit 6 are the relevant factors.

Exhibit 8 is an excerpt from a correlation matrix at the 95% barrier that aims to answer the previous questions. It uses the simulation inputs and results of all 28 firms. For all correlations with an absolute value above 43.72% the hypothesis, that they are zero at the 99% confidence level, can be rejected. This was determined using a t-test, which, according to Zimmermann [1986], performs well even for small, non-normally distributed samples. When looking at the GSO pay-off, we see that it has a relatively high correlation with the trigger probability (57.25%), but an extremely high correlation of 99.01% with the 'pay-off when triggered'. This clearly shows that when calculating the value as the product of trigger probability and 'pay-off when triggered' the second factor is of much greater importance. Therefore, we can state the cases (I) and (III) from exhibit 6 are not of great relevance. Columns four to nine of the first line of exhibit 8 show that the overnight moments are strongly correlated to the GSO pay-off, while the intraday moments are not. This is an interesting observation since on average 60.8 intraday ticks are simulated per overnight tick. A closer look at intraday variance and pay-off reveals that larger variances lead to larger pay-offs, but there are several outliers. These outliers have higher pay-offs than implied by their intraday variances. Further examination shows that these are the firms with exceptionally high overnight variances. Hence we can conclude that intraday variance is important as long as the overnight variance is not already extremely large. The 'pay-off when triggered' and its variance are correlated to the overnight moments while the trigger probability is correlated to the overall expected value of the stock returns, which is mostly driven by the intraday expected value. This observation explains the dependence of the GSO pay-off on overnight moments, since the GSO pay-off is highly correlated to the 'pay-off when triggered'. Additionally, with a correlation of 97.00% the pay-off and its variance are clearly influenced by the same factors. This means that when the average jumps are large and the GSO pay-off therefore high, the risk of very large jumps is also high and vice versa. We can now conclude that case (II) from exhibit 6 is the most relevant factor and we are interested in firms with large jumps overnight.

A closer look at the dependency between the simulated pay-off and the overnight expected value reveals that negative expected values have greater influence on the pay-off than positive ones. Exhibit 9 plots the simulated pay-off against the overnight expected value to visualize their relation. It is clear that highly negative expected values increase the pay-off while highly positive ones do not necessarily decrease it. The same observation can be made regarding 'pay-off when triggered'. It has to be noted, however, that only 8 out of 28 firms show negative overnight expected values, and this observation might, therefore, not be very stable.

Prior to the simulation one might have expected the kurtosis to play the most significant role instead of the variance. The importance is of the variance is due to the variance gamma distribution (VGD) used for the intraday stock log returns. Approximations (14) and (16) show that parameter κ basically defines the kurtosis and σ_{ID} the variance. In the VGD the rate of decay at the negative tail (Cont and Tankov, 2004) is governed by,

Gelöscht: (14)

Gelöscht: (16)

Formatiert: Englisch (USA)

Formatiert: Englisch (USA)

$$\lambda_{-} := \frac{\theta + \sqrt{\theta^2 + 2\frac{\sigma_{ID}^2}{\kappa}}}{\sigma_{ID}^2} \tag{17}$$

Therefore, an increase in κ or σ_{ID} leads to a smaller λ_{-}, which means a slower decay at the tail. But an increase in κ also means a less rounded peak. When combined this means that an increase in κ, which represents an increase of the kurtosis, will increase probabilities near the peak value and in the tail, but reduce intermediate values *(see also* Madan and Senata, 1990). By contrast, an increase in σ_{ID}, which represents an increase of the variance, will increase probabilities in the tail and increase symmetry, and therefore have a dominating influence.

As mentioned earlier, the pricing using the linear formula (9) does not reflect the pay-off structure of GSOs adequately. We have shown that the pay-off decreases more rapidly with decreasing barriers and the pay-off is greatly dependent on the underlying stocks' overnight volatility. An improved pricing approach should incorporate these findings. The EUREX Clearing AG factors in historical volatilities when determining margin requirements for equity options. The exact formula is not revealed, but they use, depending on the case, a 30-day or 250-day historical volatility of the underlying stock to create different margining groups. This idea could be used to refine the current grouping system used for GSOs, which only takes into account the index the underlying stock belongs to. In order to account for the pay-off decay at low barriers, we propose an exponential pricing model. A possible formula for the GSO premium for long positions could be,

$$R := a_{\sigma} \exp\left(b_{\sigma} B\right) \tag{18}$$

The stocks in the DAX are grouped according to σ_{ON}, their overnight standard deviation. Each group then has separate values for a and b, denoted as a_{σ} and b_{σ}. While this approach increases the complexity it should also strongly improve the precision of GSO pricing. In order to decide on the number of groups, the time frame of the overnight volatility, and the selection of parameters, further research will be required.

2005/2006 COMPARISON

In the previous sections, we have shown how GSOs on DAX firm stocks perform in highly volatile periods. The question arises: How do they perform in relatively tranquil periods? To answer this question we evaluate a second simulation.

We use the same model, but fit it to data from 2005 and 2006 using the same algorithms. To be precise, we use intraday data ranging from 3 April, 2006 to 15 May, 2006, which equals to 28 trading days. The overnight movement is fitted to data ranging from 14 June, 2005 to 15 May, 2006, which equals 236 trading days. We use all firms that were in the DAX in this time period and are in the original data sample. This leaves us with 23 out of 30 possible firms. Everything else stays the same as in the 2008/2009 case.

Exhibit 10 presents the averages that result from the new simulation and can be compared directly with exhibit 5. We can see that the average pay-off is far below the cost even at high barriers. This shows that, in tranquil periods, the guarantee of the GSO is not necessary. Exhibit 10 also shows that the trigger probability is very high, much higher indeed than in the original simulation. This is due to the fact that many firms were experiencing a downwards trend during this period. The contrast between high trigger probability and low average pay-off supports the finding in the previous chapter that the 'pay-off when triggered' is the decisive factor in determining the GSO value.

Exhibit 11 provides further insight into the simulated pay-offs at different barriers and can be compared directly with exhibit 4. Even at the highest barrier only two firms (8%) have a positive GSO value, and over half the firms have GSO pay-offs of less than half the cost. The only firm that has a positive GSO value down to the 75% barrier is Volkswagen, which has both a strong downwards trend and a high intraday volatility. Hence the only GSO with a strongly positive value is on a stock with a very pessimistic outlook.

We can, therefore, conclude that in tranquil periods a stop order is sufficient to protect stocks against sudden loses, and the guarantee of a GSO is not necessary. In other words, the current pricing model renders GSOs uninteresting for investors, since it does not factor in volatilities.

CONCLUSION

We have analysed the behaviour of GSOs on stocks following our modified variance gamma process. We have shown that the pricing of GSOs poorly reflects their expected pay-offs. The pricing performs best in turbulent stock markets, as we have been seeing them recently, with GSO barriers close to the stock price. Barriers further from the stock price cause a drop in the pay-off greater than the drop in the price of GSOs. A closer representation of this decrease would probably make GSOs more attractive for investors, which can only be in the interest of the issuing exchange. We have further shown that the variance of the underlying stocks' log-returns, especially its overnight returns, are the main factor in order to forecast the pay-off of a GSO and, therefore, to determine an adequate premium. Additionally, we have evaluated that in more tranquil times, e.g. 2005 and 2006, the classic stop order proves superior, since GSO premiums on average are too high. The pricing, therefore, needs to be linked to volatility in order to ensure that GSOs, as a product, remain interesting for investors.

Future research may reveal a pricing approach more precise than the linear one, maybe through σ-dependant exponential behaviour, as proposed earlier. Additionally, it can be interesting to see how GSOs behave under a more flexible model allowing for variable ticks per day or dependencies between returns, especially overnight returns dependent on intraday movement. Finally, one could take a closer look at the value of GSOs when short selling.

REFERENCES

Basseer, P. “Reducing Market Exposure with Portfolio Insurance.” Risk Management, 38 (1991), pp.36-45.

Bodie, Z., A. Kane, and A. Marcus. Investments, 3rd ed. n. p.: Irwin, 1996.

Carr, P., H. Geman, D. Madan, and M. Yor. “The Fine Structure of Asset Returns: An Empirical Investigation.” Journal of Business, 75 (2002), pp. 305-332.

CMC Markets. Guaranteed Stop Loss Order, London 2005.

Cont, R., and P. Tankov. Financial Modelling With Jump Processes, London: Chapman & Hall/CRC, 2004.

Easley, D., and M. O’Hara. “Order Form and Information in Securities Markets.” The Journal of Finance, 46 (1991), pp. 905-927.

Geman, H. “Pure Jump Lévy Processes for Asset Price Modeling.” Journal of Banking and Finance, 26 (2002), pp. 1297-1316.

Geman, H., D. Madan, and M. Yor. “Time Changes for Lévy Processes.” Math. Finance, 11 (2001), pp. 79-96.

Gennote, G., and H. Leland. “Market Liquidity, Hedging and Crashes.” American Economic Revies, 80 (1990), pp. 999-1021.

Madan, D. “Financial Modeling with Discontinuous Price Processes.” In O. Barndorff-Nielson, T. Mikosch, and S. Resnik, eds., Lévy Processes – Theory and Applications, Boston, MA: Birkhäuser, 2001.

Madan,D., and E. Senata. “The Variance Gamma (V.G.) Model for Share Market Returns.” Journal of Business, 63 (1990), pp. 511-524.

Schoutens, W. Lévy Processes in Finance, Chichester: John Wiley & Sons Ltd, 2003.

Seneta, E. “Fitting the Variance-Gamma Model to Financial Data.” Journal of Applied Probability, 41 (2004), pp. 177-187.

Zimmermann, D. “Tests of Significance of Correlation Coefficients in the Absence of Bivariate Normal Populations.” Journal of Experimental Education, 54 (1986), pp.223-227.

Exhibit 1

Fitted Parameters

firm	Intraday				ON	
	c	θ	σ	κ	μ	σ
	x 10^{-5}	x 10^{-5}	x 10^{-3}	x 1	x 10^{-3}	x 10^{-2}
ADS	-71.18	85.32	6.55	1.44	-0.20	1.63
ALV	-0.10	6.96	3.54	1.83	1.81	2.48
BAS	-7.71	14.72	3.13	3.57	-2.90	4.87
BAY	-13.12	18.26	2.41	3.15	2.34	1.63
BEI	33.10	-32.27	4.69	25.97	-0.05	1.44
BMW	9.05	-6.87	6.76	1.02	-0.91	2.20
CBK	2.00	-2.32	3.46	2.23	5.08	3.80
DAI	-18.86	22.96	3.32	2.10	3.22	2.33
DB1	-8.30	20.13	5.71	3.04	1.15	2.39
DBK	-5.17	8.51	2.77	3.11	2.70	2.86
DPW	-18.65	32.43	4.09	4.99	1.30	2.08
DTE	6.72	-7.27	1.61	15.27	1.82	1.33
EOAN	-6.41	11.61	2.43	2.78	-0.06	2.11
FME	-104.65	105.46	4.22	0.92	1.65	1.10
FRE3	-5.35	65.96	7.01	1.03	1.02	1.54
HEN3	-28.76	58.83	4.81	1.33	-0.33	1.72
HNR1	109.16	-129.13	11.84	2.33	4.35	2.58
LHA	6.73	0.95	3.17	1.95	-0.08	1.71
LIN	-43.03	59.45	5.51	2.82	0.71	1.49
MAN	16.13	15.17	6.58	2.28	0.71	2.23
MEO	-67.54	176.83	9.93	1.05	1.43	1.83
MRK	-106.73	69.10	6.75	2.07	1.64	1.48
MUV2	-80.75	96.92	5.88	1.63	1.14	1.77
RWE	-16.39	23.98	2.71	1.23	0.78	1.61
SAP	-25.53	51.06	4.01	1.21	0.88	1.92
SDF	13.26	2.64	4.43	1.40	-6.99	9.63
SIE	-16.35	21.72	3.69	1.35	3.07	2.09
SZG	-81.19	107.19	7.27	2.99	0.87	1.61
TKA	-3.60	16.35	4.85	1.97	1.56	2.30
VOW	37.41	-23.95	8.69	2.28	-2.19	5.29

Exhibit 2

Errors in Implied Moments

firm	error		
	variance	skewness	kurtosis
ADS	-2.45%	1.99%	-2.79%
BAS	-0.79%	0.65%	-1.21%
BAY	-1.81%	1.48%	-2.68%
BEI	-12.28%	9.07%	-19.91%
DPW	-3.14%	2.53%	-5.00%
DTE	-3.12%	2.52%	-5.59%
FME	-5.74%	4.51%	-5.06%
FRE3	-0.91%	0.75%	-0.91%
HEN3	-1.99%	1.63%	-2.21%
LIN	-3.28%	2.65%	-4.62%
MEO	-3.32%	2.67%	-3.24%
RWE	-0.96%	0.80%	-1.05%
SAP	-1.96%	1.60%	-2.09%
SIE	-0.47%	0.39%	-0.54%
SZG	-6.51%	5.08%	-8.89%

Notes: Firms that do not appear in this table have an error below 0.01%.

Exhibit 3

Simulated GSO Payoff

firm	barrier				
	95%	85%	75%	65%	55%
ADS	**0.435**	0.246	0.120	0.050	0.020
ALV	**0.295**	0.088	0.025	0.008	0.001
BAS	**1.286**	**0.860**	**0.577**	**0.344**	**0.199**
BAY	0.151	0.023	0.002	0.001	0.000
BMW	**0.597**	**0.463**	**0.343**	**0.228**	0.140
CBK	**0.418**	**0.332**	**0.230**	0.144	0.084
DAI	0.233	0.081	0.022	0.005	0.001
DB1	**0.456**	0.199	0.080	0.029	0.008
DBK	**0.289**	0.107	0.033	0.008	0.003
DPW	0.273	0.061	0.011	0.001	0.000
DTE	**0.297**	0.166	0.089	0.039	0.013
EOAN	0.247	0.068	0.018	0.004	0.000
FME	0.193	0.064	0.016	0.004	0.001
FRE3	0.171	0.017	0.001	0.000	0.000
HEN3	0.224	0.041	0.006	0.001	0.000
HNR1	**0.985**	**0.662**	**0.436**	**0.261**	0.136
LHA	0.225	0.063	0.016	0.004	0.001
LIN	0.282	0.088	0.021	0.005	0.001
MAN	**0.322**	0.081	0.017	0.002	0.000
MEO	0.241	0.028	0.002	0.000	0.000
MRK	**0.520**	**0.407**	**0.285**	0.170	0.081
MUV2	**0.302**	0.104	0.032	0.008	0.002
RWE	0.157	0.030	0.004	0.000	0.000
SAP	0.172	0.022	0.002	0.000	0.000
SDF	**3.637**	**2.792**	**2.185**	**1.618**	**1.167**
SIE	0.246	0.072	0.017	0.003	0.000
TKA	**0.295**	0.096	0.024	0.007	0.001
VOW	**1.953**	**1.549**	**1.276**	**0.922**	**0.656**

Exhibit 4

GSO Payoff Buckets

payoff / cost	barrier									
	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
< 50%	0%	39%	64%	64%	71%	75%	75%	75%	79%	82%
50% - 75%	18%	25%	4%	11%	4%	0%	4%	7%	4%	7%
75% - 100%	29%	4%	7%	0%	0%	4%	4%	0%	7%	4%
100% - 125%	21%	7%	0%	4%	4%	4%	4%	7%	4%	0%
125% - 150%	4%	0%	4%	4%	4%	4%	4%	4%	0%	0%
> 150%	29%	25%	21%	18%	18%	14%	11%	7%	7%	7%

Exhibit 5

Averages over all Simulations

average	Barrier									
	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
cost	0.285	0.270	0.255	0.240	0.225	0.210	0.195	0.180	0.165	0.150
payoff	0.532	0.402	0.315	0.256	0.210	0.171	0.138	0.112	0.090	0.071
payoff when triggered	0.730	0.686	0.646	0.608	0.563	0.544	0.507	0.434	0.474	0.447
trigger probability	63.97%	45.45%	33.91%	26.42%	21.13%	17.23%	14.16%	11.56%	9.47%	7.66%
Overnight probability	28.09%	28.30%	28.79%	28.37%	27.13%	26.57%	28.70%	28.33%	29.36%	28.84%

Notes: Averages were calculated in two steps. Step 1: For each firm: average over all simulations. Step 2: Average over each firms average. In the case of 'payoff when triggered' and 'overnight probability', only those firms where considered that reached the respective barrier in at least one simulation run.

Exhibit 6

Factors Determining GSO Value

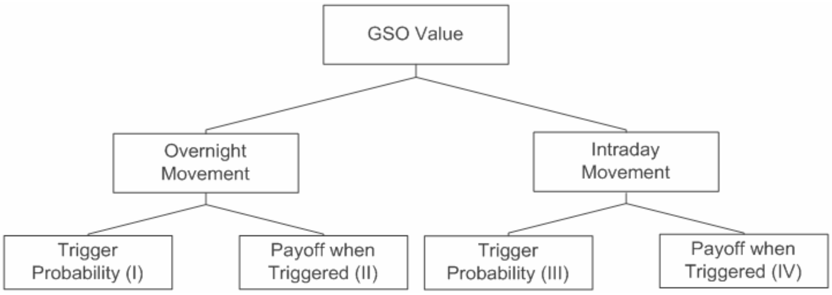

Exhibit 7

Payoff Histogramm for Volkswagen and Lufthansa at 95 % Barrier

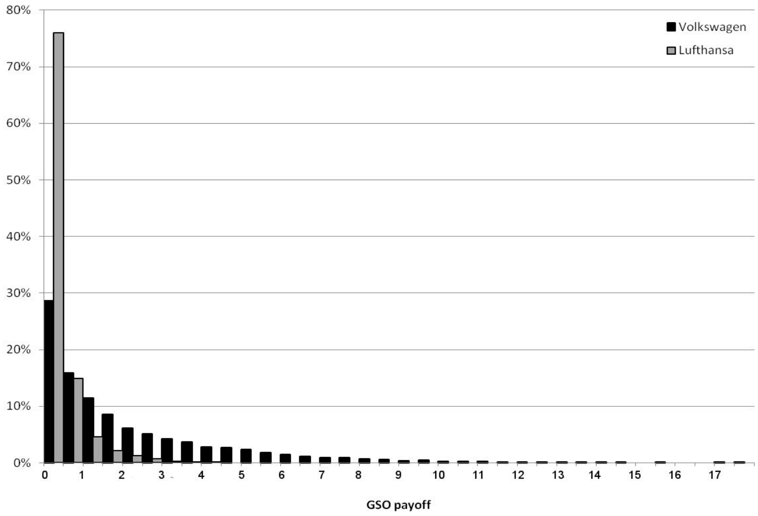

Exhibit 8

Correlations

	trigger probability	payoff when triggered	variance of payoff when triggered	moments of the log returns: intraday exp. value	intraday variance	intraday skewness	intraday kurtosis	overnight exp. value	overnight variance	overall exp. value
simulated payoff	57.25%	99.01%	95.79%	-8.95%	17.00%	-18.70%	-6.32%	-73.79%	95.40%	-35.88%
trigger prob		46.23%	39.65%	-59.68%	18.44%	-41.69%	19.94%	-23.25%	42.85%	-69.78%
payoff w.t.			97.00%	2.31%	18.10%	-13.48%	-9.74%	-76.06%	96.11%	-26.10%
variance of p.w.t.				0.54%	0.17%	-10.82%	-8.25%	-75.73%	99.25%	-20.43%

Exhibit 9

Payoff - Overnight Expected Value

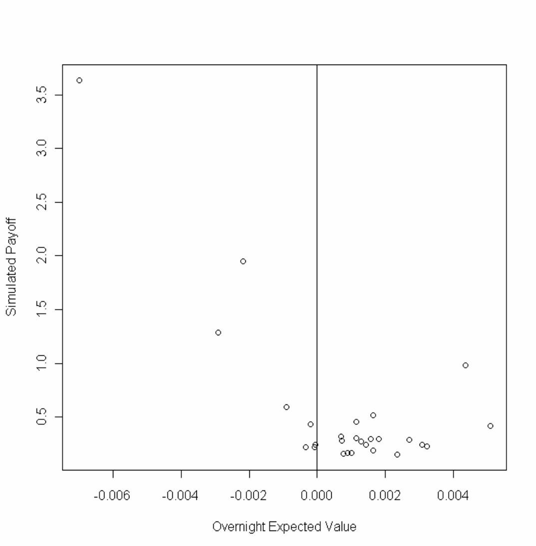

Exhibit 10

Averages over all Simulations (05/06 Data)

average	barrier									
	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
cost	0.285	0.270	0.255	0.240	0.225	0.210	0.195	0.180	0.165	0.150
payoff	0.149	0.127	0.111	0.095	0.081	0.066	0.052	0.040	0.029	0.020
payoff when triggered	0.188	0.177	0.171	0.163	0.139	0.146	0.121	0.139	0.106	0.102
trigger probability	78.06%	69.21%	63.22%	57.44%	51.24%	44.54%	37.78%	31.20%	25.12%	19.73%
Overnight probability	9.67%	9.48%	9.89%	9.93%	12.62%	10.98%	8.06%	8.43%	8.19%	9.46%

Notes: Averages were calculated in two steps. Step 1: For each firm: average over all simulations. Step 2: Average over each firm's average. In the case of 'payoff when triggered' and 'overnight probability', only those firms where considered that reached the respective barrier in at least one simulation run.

Exhibit 11

GSO Payoff Buckets (05/06 Data)

payoff / cost	barrier									
	95%	90%	85%	80%	75%	70%	65%	60%	55%	50%
< 50%	52%	57%	74%	78%	83%	83%	87%	87%	91%	91%
50% - 75%	30%	26%	9%	4%	4%	4%	4%	4%	0%	4%
75% - 100%	9%	9%	9%	9%	9%	9%	4%	4%	9%	4%
100% - 125%	4%	4%	4%	4%	0%	0%	0%	4%	0%	0%
125% - 150%	0%	0%	0%	0%	0%	4%	4%	0%	0%	0%
> 150%	4%	4%	4%	4%	4%	0%	0%	0%	0%	0%