

National Centre of Competence in Research
Financial Valuation and Risk Management

Working Paper 319

Dynamic Hedging and Order Flow Dynamics

Carsten Murawski

First version: June 2006
Current version: June 2006

This research has been carried out within the NCCR FINRISK project on
“Credit Risk and Non-Standard Sources of Risk in Finance”

Dynamic Hedging and Order Flow Dynamics¹

Carsten Murawski

*Swiss Banking Institute, University of Zurich, Plattenstrasse 14, 8032 Zurich,
Switzerland.*

Abstract

The price process in a financial market is driven by demand and supply. Statistical analyses have shown that price “feeds back” on future demand and supply. To date, few testable models have been proposed that offer an economic explanation for this relationship. In this paper, we investigate a mechanism that might, at least to some extent, explain the positive feedback effect of prices on future order volume, namely dynamic hedging. Using order book and transaction data from the Swiss Stock Exchange and Eurex, we investigate the dependencies between theoretical hedging demand generated by option holders and actual order volume in the underlying stock market. We are able to show that the popular belief that options trading has a negative systematic impact on the underlying market in the form of a positive feedback does not seem to be warranted.

Key words: Derivative securities, dynamic hedging, feedback effects, order flow dynamics

Email address: murawski@isb.unizh.ch (Carsten Murawski).

¹ The author would like to thank Rajna Gibson for helpful discussions and relentless support throughout, and Markus Leippold for inspiring conversations on the topic. Further, the author thanks SWX SwissExchange for providing the dataset; Wolfgang Bauer for making available a preprocessed version of the data, for lots of computer code, and for numerous fruitful discussions; Luigi Vignola and his team at Zürcher Kantonalbank for a thorough introduction to the business of options and warrants trading. This research has been carried out within the project on “*Conceptual Issues in Financial Risk Management*” of the National Centre of Competence in Research “Financial Valuation and Risk Management” (NCCR FINRISK). The NCCR FINRISK is a research program supported by the Swiss National Science Foundation.

1 Introduction

This paper analyzes the relationship between dynamic hedging of options and order volume in the underlying market. We address the following research questions: (1) What are the major properties of order flow on the underlying market? (2) Is there a relationship between price change and future order volume? (3) Can we explain this relationship by dynamic hedging activities of option holders?

Starting with Mandelbrot (1963), many empirical studies of the fluctuations in the price of stocks have shown that distributions of returns have fat tails that deviate from the Gaussian distribution, including Fama (1965), Mandelbrot (1966), Officer (1972), Pagan (1996), Campbell, Lo & MacKinlay (1997), Pictet, Dacorogna & Müller (1996), Bouchaud, Gefen, Potters & Wyart (2004). This is particularly so for intraday timescales as shown by Cont, Potters & Bouchaud (1997). Bollerslev, Chou & Kroner (1992) show that these fat tails, characterized by a significant kurtosis, persist even after accounting for heteroskedasticity in the data. More recently, research has indicated that large price changes like “crashes” do not seem to be outliers of the price distribution. In fact, Gabaix, Gopikrishnan, Plerou & Stanley (2003) and Sornette (2003), among others, argue that price changes seem to be distributed according to power laws. A power law does not have a characteristic scale implying that large events occur with much higher probability than in case of the Gaussian distribution. Very often, systems with certain types of feedback mechanisms show “power-law behavior”.

The fact that significant fluctuations in prices are not necessarily related to the arrival of information, as demonstrated by Cutler, Poterba & Summers (1989), or to variations in fundamental economic variables, as investigated by Shiller (1989), has lead researcher to think that the high variability present in stock market returns may indeed correspond to certain types of positive feedback. One mechanism often suggested in this context is herding which was investigated by Cont & Bouchaud (2000) among many others.

Another mechanism that has been suggested is options trading. In fact, options markets are often proposed as the main suspect responsible for the high variability of stock prices. There is indeed some evidence that the volatility of prices in certain markets changed after the introduction of options on the assets traded in the respective market, as discussed for example by Damodaran & Lim (1991) and Bruart & Gibson (1998). It is not clear, though, whether there is a causal relationship. To date, little evidence has been produced that establishes such a causal link between options markets on the one hand and trading volume and price movements in the underlying stock markets on the other hand. Exceptions are studies by Anthony (1988) who analyzes the relationship between options and stock market trading volume and Easley, O’Hara & Srinivas (1998) who investigate options trading volume to study information dissemination in markets.

Albeit, there seems to exist an obvious link between options and underlying markets. “Producers” of options, holding a short position in a contract, typically hedge their risk by buying or selling shares in the underlying market. However, a trader’s position remains hedged, or “delta-neutral”, for only a relatively short period of time. Thus, the hedge has to be adjusted periodically. This is known as *rebalancing*. Such a hedging scheme is commonly called *dynamic hedging*. The hedging strategy is typically computed according to the Black-Scholes formula. Suppose a trader is short in a call option. When the price of the underlying falls, the trader will sell shares in the underlying, potentially causing a further fall in the stock price. This way, dynamic hedging might indeed create a positive feedback where a price change in the underlying stock leads to subsequent changes of the same sign.

In order for such a positive feedback to set in, several criteria have to be met all at once: (1) the exposure created by options has to be such that hedging demand is “consistent” with price changes; this means that the delta of the options portfolio has to be positive so that a negative price change requires the sale of shares (and vice versa for a positive price change); (2) hedging demand has to account for a considerable part of order volume in the underlying market; (3) the price impact of trading in the underlying has to be non-zero. If either of these criteria is not met, a price change is unlikely to “feed back” onto price.

If a systematic relationship between options and underlying markets in the form of positive feedback did indeed exist, it would have significant implications for market participants. They would have at least to reconsider pricing and risk management methods to take account of feedback effects. The Black-Scholes model and most of its derivatives ignore any trading costs occurring as a result of dynamic hedging. Maybe more importantly, such a link would have severe implications for public policy. It would mean that options trading, the domain of large institutional investors, creates an externality by systemically consuming liquidity and amplifying returns in the stock markets, which have a much wider range of participants. Options trading would increase liquidity costs as well as systemic risk.

The aim of this paper is to establish whether such a systematic relationship between options and underlying stock markets exists or not. Our paper contributes to the burgeoning body of empirical literature exploring “microscopic” properties of financial markets including order flow characteristics, limit order book dynamics, price impact functions, as well as price processes. Much work has been done in the last few years exploring the relationship between order flow, via the limit order book, and price as described in Biais, Hillion & Spatt (1995), Gopikrishnan, Plerou, Gabaix & Stanley (2000), and Cont (2001), Potters & Bouchaud (2003), among others. Although there are still a large number of open questions, many important discoveries have been made. These include the long-memory property and power-law distribution of order volume, the power-law tail of the distribution of limit prices, and the concave shape of the price impact function. At the moment, a large part of empirical

and theoretical research seems to focus on understanding the exact relationship between order flow and price process. Some explanations have been offered by Farmer, Gillemot, Lillo, Mike & Sen (2004), Bouchaud et al. (2004), and, in a theoretical framework, Smith, Farmer, Gillemot & Krishnamurthy (2003). However, almost no analysis to date, at least to our knowledge, has been made of the drivers of order flow.

In this paper we address this question. Thus, it extends the literature on the microscopic analysis of the *dynamics* of financial markets. We investigate the potential impact of price on future order volume, trying to “close to loop” between order volume and the price process. Our work is in the same spirit as Kambhu (1998) and Kambhu & Mosser (2001) who analyze the impact of hedging and its effect on underlying markets in the U.S. dollar swaps market. These studies find that no such link exists, at least not a systematic one. However, these analyses are based on rather “coarse” data. It has become clear recently that many properties of financial markets can only be well understood if analyzed on very short—that is, intra-day—time scales.

To our knowledge, our analysis is the first that investigates options hedging as a potential feedback mechanism using microscopic, that is, ultra-high frequency data. It is based on transaction level data on order flow and trades of the Swiss Stock Exchange (SWX), as well as daily position data from Eurex on options on stocks traded on SWX. We test a model, with several variations, that relates (theoretical) hedging demand to future order volume in the underlying market. Our model does not have any free parameter, that is, it can be tested with our data. To put it differently, “we let the data speak”. For reasons of space and due to limited computational resources we only present results for one stock, Novartis (ticker symbol NOVN).

Our null hypothesis, inspired to some extent by the results of Kambhu (1998) and Kambhu & Mosser (2001), is that there is no systematic relationship between hedging demand and future order volume. Based on various statistical tests, we do not reject our null hypothesis. This means that we do not find that hedging demand is systematically related to future order volume in the underlying market. In other words, while at times hedging demand might contribute a significant share to order volume in the underlying market, a large part of the variation seems to originate from other sources. Neither do we find a systematic relationship between hedging demand and price impact. This means that at at times when hedging demand represents a large part of order volume, price impact is not systematically higher than at other times.

Our findings are robust across many model specifications. We consider our results to be important as they cast serious doubt on the all too popular argument and widely-held belief that options trading has a systematic negative impact—in the form of a positive feedback effect—on underlying markets.

The remainder of the paper is organized as follows: We first review the relevant literature in Section 2. Then, in Section 3, we briefly describe the markets we analyze as well as our dataset. In Section 4 we describe our research methods and the model, before we present results of our analysis in Section 5. In

Section 6 we conclude.

2 Literature Review

The process of price formation in financial markets and its outcomes has been the domain of market microstructure research. While much of economics abstracts from the mechanics of trading, the microstructure literature analyzes how specific trading mechanisms affect the price formation process. Early work in this field investigates issues in relation to the stochastic nature of supply and demand. More recent work focuses on the properties of information aggregation provided by prices and markets.

Most models offered by the market microstructure literature suffer from a serious drawback, though. Their results tend to be very sensitive to one or more free parameters in the models. Therefore, these models are rather difficult to test.

In the last several years, an abundance of financial markets data has become available. In many cases, the data allows researchers to reconstruct the trading process of a given market at the most fundamental level, order by order or trade by trade. This allows researchers to develop microscopic models of financial markets. However, the sheer size of this data makes model building a complex task. To overcome this hurdle, researchers have pursued a “divide & conquer” strategy. The first step comprises the description of the data by extracting properties of the observable quantities, usually in the form of statistical regularities. Once the data is sufficiently understood, it is then described by a model. These models usually only have few, if any, free parameters so that they can be tested with the data at hand. While so far considerable progress has been made with regards to understanding the properties of markets, the second step, the development of models that describe the data, is still in its infancy.

As we know from the market microstructure literature, price formation in a given market depends on the specific trading mechanism employed in that market. In most financial markets, trading takes place continuously throughout the day and is conducted via a continuous double auction. In order to better understand the literature we now briefly describe the life cycle of an order in such a market. We also define important terms and describe the quantities that can be measured at the different stages. A more detailed description of the continuous double auction is provided in Section 3.

We call the sequence of orders submitted to the market the *order flow*. We differentiate two types of orders. *Market orders* are requests to buy or sell a given number of shares immediately at the best available price. The other type of orders is called *limit orders*. A limit order specifies a quantity that the participant wants to buy or sell, and a limit price corresponding to the

worst allowable price for the transaction. Limit orders often do not result in an immediate transaction and are stored in a queue called the *limit order book*. Buy limit orders are typically called *bids* and sell limit orders are called *offers* or *asks*.

The limit price of the best buy order in the queue is often referred to as *best bid*, whereas the limit price of the best sell order is referred to as *best ask*. The, typically non-zero, gap between the best bid and the best ask is called *spread*. Prices are continuous but have discrete quanta called *ticks*. The minimum interval that prices change on is the *tick size*.

When market orders arrive they are matched against limit orders of the opposite sign in order of first price and then arrival time. As orders are placed for a varying number of shares, matching is not necessarily one-to-one. A high density of limit orders per price is often referred to as high *liquidity* for market orders, that is, it decreases the price movement when a market order is placed. We define market liquidity or depth as the volume of shares for sale at up to k ticks from the midquote. It can be interpreted as the volume necessary to move the price by k ticks. More liquid or deeper markets can accommodate larger trades for a given price impact.

We will refer to a buy limit order whose limit price is greater than the best ask, or a sell limit order whose limit price is less than the best bid, as a *marketable limit order*. Such limit orders result in immediate transactions, with at least part of the order executed straightaway.

A limit order can also be removed from the queue by being cancelled, which can occur at any time. The best price may change as new orders arrive or old orders are cancelled.

In a continuous double auction, the relationship between traders' order placement and the short-term dynamics of asset prices, the transmission of information in security markets, the costliness of trading, and the nature of liquidity available in the market is rather complex. In the following, we describe the literature that investigates these quantities and their interaction. Our review of the findings proceeds in the order of the life cycle of a trade, as explained above. We will first look at single quantities and then at the dynamic interaction between them.

Our review will focus on the more recent literature and only refer to the traditional market microstructure literature where appropriate. Comprehensive surveys of the latter include the standard reference by O'Hara (1999) as well as Madhavan (2000), Biais, Gloston & Spatt (2002), Harris (2003), and Biais, Gloston & Spatt (2005).

We begin with the empirical part of the literature. As there is no overview of findings available in the literature—at least to our knowledge—our review will be somewhat more extensive. We start by looking at order flow. Obviously, order flow drives the trading process so it is essential to understand its properties such as size distribution, (limit) price distribution, seasonality, and

so on.

One of the first studies of microscopic financial data in general and order flow in particular was conducted by Biais et al. (1995) who analyze order flow and transaction data of the Paris Bourse for three days. They find that the placement of new orders tends to be concentrated in the morning whereas cancellations tend to occur in the afternoon. Rinaldo (2003) investigates order flow on the Swiss Stock Exchange and also finds strong intraday seasonality in order placement, whereas Grammig, Heinen & Rengifo (2004), analyzing order flow on the Xetra system which includes orders for the Frankfurt Stock Exchange, do not find such patterns.

As Biais et al. (1995) further show, there is strong persistence in the time interval between orders after controlling for the time of the day. Short time intervals tend to follow short time intervals, and long time intervals tend to follow long time intervals. The most frequent events are small orders that are submitted within the bid-ask spread and are thus immediately and fully executed. The next most frequent events are new orders placed within or at the spread. This means that most activity is within or at the quotes.

Biais et al. also reveal that the probabilities of orders are conditional on the last order. This means that the probability that a given type of order occurs is larger after this event has occurred than it would be unconditionally. Biais et al. call this the *diagonal effect* and point out that it could reflect strategic order splitting, imitation, or similar, but successive, reaction of traders to the same event. We point out at this stage that the latter could be related to dynamic hedging activities of traders.

Biais et al. also find that new orders inside the quotes are relatively more frequent when the spread is large, whereas trades are relatively more frequent when the spread is tight. Such order placement offers liquidity when it is scarce and consumes it when it is plentiful. The fact that orders tend to be placed within the quotes when the spread is large generates mean reversion in the bid-ask spread. This suggests that increases in the spread are not permanent but transient changes, due to liquidity shocks.

Market buy (sell) orders tend to occur after market sell (buy) orders. This generates negative serial correlation in the changes in both the ask and the bid quotes. The finding suggests that there is negative serial correlation in the quote changes and in the spread, reflecting the dynamics in the supply of liquidity on both sides of the market.

Gopikrishnan et al. (2000) analyze transaction data for the largest 1,000 U.S. stocks for the two-year period 1994–1995. They show that the distribution of transaction volume displays a power-law decay, and that the time correlations in volume display long-range persistence. Such long-range persistence is called *long memory*. More precisely, a random process is said to have long memory if it has an autocorrelation function that is not integrable. This is the case, for example, when the autocorrelation function decays asymptotically as a power law of the form $\tau^{-\alpha}$ with $\alpha < 1$. It implies that the values from the distant past can have a significant effect on the present. Gopikrishnan et al. (2000)

also find that the long-range correlations in volume are largely due to those in the number of transactions.

Lillo & Farmer (2004) study the order flow of the London Stock exchange and establish that the signs of orders obey a long-memory process. The autocorrelation function of the series of signs of orders decays roughly as a power law with an exponent of 0.6. This means that signs of future orders are quite predictable from the signs of past orders. It would suggest strong market inefficiency. However, it appears that fluctuations in order signs are compensated for by anti-correlated fluctuations in transaction size and liquidity, which are also long-memory processes that act to make the returns whiter. In the analysis of Lillo & Farmer, positive autocorrelation coefficients are seen at statistically significant levels over lags of many thousand events, spanning many days.

Grammig et al. (2004), in their analysis of order flow on the Xetra system, find that market buy and market sell orders are almost symmetric in terms of volume.

Niederhoffer & Osborne (1966) discuss the granularity of supply and demand, defined as the difference between limit prices, for the stocks of the Dow Jones Industrial Index. They show that orders tend to cluster at particular prices.

A more detailed analysis of limit prices was conducted by Bouchaud, Mézard & Potters (2002) who analyze the order flow of the Paris Bourse. They find that the price at which new limit orders are placed is very broadly power-law distributed around the current bid-ask spread. According to their analysis, the distribution of volume at the bid and ask follows a gamma distribution.

Farmer & Zovko (2002) investigate 20 years of order flow data of the London Stock Exchange. They define the *relative limit price* as the difference between the limit price of a given order and the best price. Farmer & Zovko demonstrate that for both buy and sell orders, the unconditional cumulative distribution of relative limit prices decays roughly as a power law with an exponent of approximately -1.5. The time series of relative limit prices are characterized by an autocorrelation function that asymptotically decays with a power of -0.4.

Other studies of order flow include those of Lehmann & Modest (1994) who analyze order placement at the Tokyo Stock Exchange, and Harris & Hasbrouck (1996) who study order placement at the New York Stock Exchange.

We now turn to the order book. Biais et al. (1995), in their analysis of the order book of the Paris Bourse, establish that the bid-ask spread is typically more than twice as large as the difference between adjacent limit prices on each side of the book, while the latter are fairly constant. If the bid-ask spread is larger than twice the other spreads, they find that the slope of the book is steeper at the quotes than away from the quotes. The depth of the book is lower at the bid-ask quotes than away from the quotes, and it increases the further away the limit prices are from the best quotes. This reflects the accumulation of old, unexecuted, uncanceled orders. By definition, new marketable orders first hit the best quotes in the order book, widening the bid-ask spread, and reducing the depth at these quotes. Typically, the provision of liquidity at the

bid-ask quotes is only a small portion of the overall depth of the order book. As the spread is larger and the depth is lower at the best quotes, the price schedule is steeper at the best quotes than away from the quotes. This may reflect, so Biais et al. suggest, the correlation in the value of a given security to various bidders on the same side of the market. It may also reflect the extent of competition among bidders on the same side of the market. Furthermore, it may be related to the shading of bids compared to the underlying reservation values. In general, the price schedule is weakly concave, not departing strongly from linearity.

Biais et al. discover that the bid-ask spread as well as the relative spreads on each side of the book exhibit a U-shaped pattern along the course of the day, while the depth at the best quotes is largest during the largest two trading hours.

In addition, Biais et al. find that the widening of the spread and the decrease in depth at the best quotes are transient in that they attract new orders at or within the quotes, which restore depth, and narrow the spread.

Grammig et al. (2004), in their analysis of the Frankfurt stock market, find that larger spreads reduce the relative importance of market order trading compared to limit order submission. They classify orders according to their *aggressiveness* which relates order price and size to the prevailing price and depth of bid and ask quotes. This concept was introduced, to the best of our knowledge, by Biais et al. (1995). The depth at the best quotes stimulates the submission of aggressive limit orders on the same side of the market, as limit order traders seem to strive for price priority. Larger depth at the opposite side of the market reduces the aggressiveness of own-side limit orders. Grammig et al. find that a considerable fraction of the variation of market liquidity can be explained by two principal components, one related to liquidity and the other related to information.

Griffiths, Smith, Turnbull & White (2000) and references therein investigate the trading cost of orders. Griffiths et al. compare the cost of orders in terms of price impact of market and limit orders on the Toronto Stock Exchange. They find that the higher an order's price and quantity are set relative to those available on the limit-order book, the larger is the price impact but the lower are the opportunity costs. Conversely, passive orders have negative price impacts but high opportunity costs. These opportunity costs arise because a large percentage of the passive orders are not filled.

Goldstein & Kavajecz (2000) investigate the relationship between tick size and liquidity provision. They analyze the reduction in tick size on the NYSE from eighths to sixteenth in June 1997 and find that the change in tick size resulted in a decline in both, spreads and depths. They show that the reduction in tick size reduced the cost for small market orders, but that trading costs of large market orders increased.

The impact of orders on price has recently been the subject of several other studies. The function that relates the change of price to the volume of orders is commonly called the *price impact function*. Coppejans, Domowitz & Mad-

havan (2000) study orders and trades on the Swedish stock market and find that the price impact function is convex in order size and that it varies over time. Price impacts are higher at market opening. The price impact function does not vary with the type of the order, that is, whether it is a buy or a sell order.

Plerou, Gopikrishnan, Gabaix & Stanley (2002) find that the average price impact function displays a concave functional form that seems universal for all stocks. Moreover, they find that large price fluctuations occur when demand is small. The price impact function displays a power-law behavior for small volumes.

Lillo, Farmer & Mantegna (2002) study the average price impact of a single trade on the New York Stock Exchange for the years 1995 to 1998. They analyze data for the 1,000 most highly capitalized stocks. They find that the price impact function for those stocks collapses onto a single function. This function increases as a power law of order 0.5 for small volumes, but then increases more slowly for large volumes.

Potters & Bouchaud (2003) discover a logarithmic price impact function for NASDAQ stocks. They find weak time dependence of the response of prices, which means that price impact is quasi-permanent. This suggests that trading is interpreted by the market as new information.

Petersen & Umlauf (1994) and Knez & Ready (1996) show that for the New York Stock Exchange the most important conditioning variable for price impact is the size of an order relative to the volume at the best price.

Many of these results just mentioned are contested by a recent study by Farmer et al. (2004). In a rather interesting, and important, piece of work, they analyze large price changes on the London Stock Exchange. They establish that even for the most liquid stocks there can be substantial gaps in the order book, corresponding to a block of adjacent price levels containing no limit orders. When such a gap exists next to the best price, a new order can remove the best quote, triggering a large midpoint price change. Thus, the distribution of large price changes merely reflects the distribution of gaps in the limit order book. Put differently, large price fluctuations are driven by liquidity fluctuations, the variations in the market's ability to absorb new orders. Or, in other words, order size is not necessarily the driver of large price changes, meaning that even small marketable orders can have a large price impact, and large orders might only have a small price impact. This effect is finite in size, caused by the granularity of order flow. If market participants placed many small orders uniformly across prices, such large price fluctuations would not happen. These findings might also explain price fluctuations on longer time scales. The insights of Farmer et al. are obviously fairly important for the understanding of price dynamics in particular and market dynamics in general.

These facts do not mean, though, that fluctuations in orders are not important. However, it is not the size of orders that drives large price changes, but rather the uniformity of their coverage of price levels. Revealed supply and demand curves at any instant in time are irregular step-like functions with

long flat regions and large jumps.

Several studies investigate the impact of price impact as conditional expectation value of order imbalance, including Hasbrouck (1991), Hausman, Lo & MacKinlay (1992), Kempf & Korn (1999), Evans & Lyons (2002), and Bouchaud et al. (2004)). Gabaix et al. (2003) suggest that large price changes are due to large order imbalances being defined as the difference between the number of shares bought and the number of shares sold. Starting from the distribution of trading volume and a fit to the average price impact function, they suggest an explanation for the empirical power law distribution of stock prices. This study was criticized by Farmer & Lillo (2004) on the basis that the test presented in Gabaix et al. (2003) lacks power in the presence of correlations in the order flow and because the functional form used to describe the price impact of large orders seems to vary for large stocks. Instead, the authors suggest in Farmer et al. (2004), on the basis of an event-based analysis, that large price changes are due to the granularity of the order book, which gives rise to time-varying liquidity, as described above.

Weber & Rosenow (2004) study the price impact function for the most frequently traded stocks on NASDAQ and Island. They find that price changes, in five minute intervals, larger than five standard deviations can hardly be explained by order imbalances. Instead, they suggest that price changes are explained by unusually large slopes of the price impact function. In other words, large price changes cannot be explained by the *average* price impact function but by a dynamical version. They show that the *actual* price impact function for large events, that is, price shifts of more than five standard deviations away from the previous price, has a steeper slope than the average price impact function. This implies that in time intervals with large price movements there are less limit orders available than on average.

The actual price impact function only takes account of price changes due to marketable orders. Weber & Rosenow (2003) study the behavior of a virtual price impact function, that is, they compute the price impact of an hypothetical market order every time a new order is placed in the order book. This way, they take account of limit orders that are placed in between marketable orders but are not immediately executed. The virtual price impact of a given order volume is roughly four times stronger than the actual price impact. The average price impact typically has a concave shape. Theoretically, this would be an incentive to execute large trades as they would be less costly. The virtual price impact function computed by Weber & Rosenow, on the other hand, has a convex shape. The authors show that the correlation between returns and limit order placement can explain the difference between actual and virtual price impact function. It seems that that limit order placement is negatively correlated with returns (in contrast to the positive correlation between returns and market orders). The authors suggest that limit orders placed in response to returns provide a quantitative link between virtual and actual price impact.

A puzzle has been the reconciliation of the strong autocorrelation of order

flow with the uncorrelated nature of returns. Bouchaud et al. (2004) offer a potential explanation by suggesting that liquidity provision reduces the “bare” price impact, consistent with the arguments presented by Farmer et al. (2004). Weber & Rosenow (2003) show that anticorrelation of limit order flow and returns strongly reduces the price impact of market orders and is responsible for the empirically observed concave shape of price impact function.

Gopikrishnan, Plerou, Gabaix, Amaral & Stanley (2001) investigate the relation between trading activity—measured by the number of transactions—and the price change for a given stock. They find that long-ranged volatility correlations are largely due to correlations of the number of transactions in a certain time interval. These results are confirmed by a study of Rosenow (2002). He also finds that price changes are not only due to the large volume of trades but also to a large price impact of a given volume. Again, large price fluctuations seem to be caused by changes in market depth, that is, in liquidity.

We now look at the price process. The dynamics of the price result from the interplay between the order book and order flow. A major property of the price process, that the distribution of returns has fat tails, has been known for a long time, at least since the work of Mandelbrot (1963) and Mandelbrot (1966). More recent studies include Pagan (1996), Campbell et al. (1997), Pictet et al. (1996), Bouchaud et al. (2004). This is especially so for intraday timescales as shown by Cont et al. (1997). These fat tails, characterized by a significant excess kurtosis, persist even after accounting for heteroskedasticity in the data, which was demonstrated by Bollerslev et al. (1992). The heavy tails observed in these distributions correspond to large fluctuation in prices. As Cutler et al. (1989) or Shiller (1989) argue, “bursts” of volatility are difficult to explain only in terms of variations in fundamental economic variables.

Cont (2001), in a rather extensive analysis, establishes a set of stylized facts about the price variations in various markets. He finds that there is no autocorrelation of prices except on very small timescales, typically less than 20 minutes. The unconditional distribution of returns seems to display a power-law or Pareto-like tail with a finite tail index. Moreover, he finds an asymmetry between gains and losses meaning that there are large drawdowns in stock prices and indexes but not equally large upward movements.

Cont (1998) also points out that the heavy tails observed are fatter than those of a normal distribution but thinner than that of a stable Pareto-Lévy distribution. In some cases they can be represented by an exponential form, as shown by Cont et al. (1997), and in other cases by a power law with a tail index between 3 and 4, as shown by Pagan (1996).

Gabaix et al. (2003) analyze the price changes of the 30 largest stocks on the Paris Bourse for a five-year period. They also find that price changes can be characterized by a power law and show that a power-law distribution of price changes is universal across markets. They suggest that large price changes in a market are caused by the trading activities of large market participants such as mutual funds, referring to Keim & Madhavan (1997) who find that large

trades by institutional investors seem to have large price impacts.

Farmer et al. (2004) analyze several recent studies and point out that the exponent of the power law describing large price returns seems to be larger than two. This implies that the standard deviation of price returns is finite. However, it remains controversial whether a power law is always the best description of price returns. In any case, tail exponents vary from stock to stock and are not universal.

Cont (2001) finds that the distribution of returns looks more and more like a normal distribution as one increases the time scale over which returns are calculated. He calls this phenomenon *aggregational Gaussianity*. At any time, returns display a high degree of variability. Different measures of volatility have a positive autocorrelation over several days. This is usually called *volatility clustering*. Even after correcting for clustering, the residual time series still exhibit heavy tails.

In addition, Cont (2001) establishes that the autocorrelation function of absolute returns decays slowly as a function of the time lag, roughly as a power law with an exponent between 0.2 and 0.4. This is sometimes interpreted as a sign of a long memory, as described above. Most measures of volatility of an asset are negatively correlated with the returns of that asset, called the *leverage affect*.

Bouchaud et al. (2004) analyze price changes of stocks traded on the Paris Bourse. They find that prices follow roughly a random walk. However, the random walk nature of prices seems to be highly non-trivial. It originates from the interplay of market and limit order flow. Long-range correlated market orders create persistence in the price diffusion, whereas mean-reverting limit orders create “anti-persistence”. These two effects, persistence and anti-persistence, precisely compensate each other and lead to an overall approximative diffusive behavior of prices. Their results are related to studies of order flow that were described earlier in this section.

Muzy, Bacry & Delour (2001) demonstrate that the volatility of most assets shows random variations in time, with a correlation function that decays as a small power of the time lag. Their finding is quite universal across markets.

Cross correlations between stock price changes are studied by Plerou, Gopikrishnan, Rosenow, Amaral & Stanley (1999), Plerou, Gopikrishnan, Rosenow, Amaral & Stanley (2000), and Plerou, Gopikrishnan, Rosenow, Amaral, Guhr & Stanley (2002).

We now turn to dynamical aspects of the trading process. We first look at the interrelation between the order book and order flow. Biais et al. (1995) investigate how order flow reacts to the state of the order book and informational events in the market place, how it mechanically updates the state of the order book, and how it influences the subsequent evolution of the order book and trading activity in the market. They find that the conditional probability that investors place limit orders, instead of “hitting” the best quotes, is larger when the bid-ask spread is large or the order book is thin. On the other hand, investors tend to hit the best quotes when the spread is tight. This means that

investors provide liquidity when it is valuable to the marketplace and consume liquidity when it is plentiful. Traders place limit orders relatively quickly when liquidity has diminished. Biais et al. also find that after a market order has been placed, the probability that the next order will provide liquidity is high. The market's response to market orders tends to be rapid, which might reflect competition in supplying liquidity. It suggests the existence of potential liquidity supply outside the book.

In addition, Biais et al. show that the order placement is concentrated at and inside the current bid-ask quotes. A large fraction of the order placements improves upon the best bid or ask quote. Improvements in the best quote is especially pronounced when the depth at the quotes is large. On the other side, quantities in the order book are not concentrated near the best quotes. Biais et al. interpret several of their findings with regards to information effects. They find that after large sales (purchases), which consume liquidity at the best quote and thus induce a decrease in the bid (increase in the ask), there is often a new sell (buy) order placed within the quotes, which generates a decrease in the ask (increase in the bid). Biais et al. suggest that this reflects the adjustment of the market's expectation to the information content of the trade. This is consistent with models by Glosten & Milgrom (1985) and Easley & O'Hara (1987). It illustrates the positive serial correlation in changes in quotes across the two sides of the market generated by information events. Large purchases (sales) tend to occur in succession and rather quickly, a phenomenon that is in line with models by Easley & O'Hara (1987) and Easley & O'Hara (1992).

Coppejans et al. (2000) analyze the links between market liquidity, order placement behavior, and returns. They show that the variation in liquidity over time is economically and statistically significant. This finding is consistent with a theoretical model by Admati & Pfleiderer (1988).

Plerou, Gopikrishnan, Amaral, Gabaix & Stanley (2000) analyze the relationship between the number of transactions and price changes for every transaction for 1,000 stocks on the U.S. market for the two-year period from 1994–1995. They find that price movements follow a complex variant of classic diffusion, where the diffusion constant fluctuates drastically in time. In addition, they show that the power-law tails of the distribution of price changes is related to the distribution of the variance of price changes, and that long-range correlations in the absolute value of price changes are driven by the number of transactions.

Another topic of interest is the interaction of liquidity and short-horizon expected returns. Amihud & Mendelson (1986) and Amihud & Mendelson (1991) find evidence of a positive relation between asset returns and bid-ask spreads. Amihud, Mendelson & Lauterbach (1997) document large changes in asset values for stocks moving to more liquid trading systems on the Tel Aviv Stock Exchange. Brennan & Subrahmanyam (1996) and Brennan, Chordia & Subrahmanyam (1998) show that liquidity can explain the cross-sectional variation.

Coppejans et al. (2000), analyzing data of the Swedish Stock Market (OMX), find that the relationship between liquidity and returns is rather complex. They discover some kind of “self-correcting ability” of the market in the sense that volatility shocks reduce liquidity but dissipate quickly, indicating a high degree of resiliency.

Hasbrouck & Seppi (2001), Bauer (2004), and Grammig et al. (2004) examine commonality in liquidity.

Farmer & Zovko (2002) find that relative limit price levels are positively correlated with and are led by price volatility, which may potentially contribute to clustered volatility.

Feedback effects in markets due to mechanisms like herding and program trading are discussed in Grossman (1988), Gennotte & Leland (1990), Pritsker (1997), and Subrahmanyam (1991).

So far, we have looked at stock markets in isolation. We now turn to the interrelation between options and underlying stock markets. Anthony (1988) finds that trading volume on the Chicago Board Options Exchange leads trading volume in the underlying shares, with a one-day lag. Schlag & Stoll (2005) conduct a similar analysis for German DAX futures and futures options. They find that futures volume leads options volume. They also find that neither options nor futures volume predicts price changes.

Kambhu (1998) analyzes the potential impact of dynamic hedging of U.S. dollar interest rate options on liquidity in the underlying market. He finds that only in one segment, medium-term maturities, an unusually large shock to interest rates could cause the hedging of exposures to generate demand that is high relative to trading volume. In other words, underlying markets are large enough to absorb hedging demand, except for one segment and very large shocks. A serious issue of Kambhu’s analysis is his dataset. The estimate of traders’ exposure is based on information for one point in time. This means that he keeps exposure constant when he computes hedging demand for the years surrounding this date.

Kambhu & Mosser (2001) extend the above analysis by considering two different time periods, 1990–1995 and 1996–2000. They find certain indications for a positive feedback in short-term rates, and that these indications have become stronger in the second period. A likely reason is the tremendous growth of interest rate derivatives in this time period. Their analysis implies that, with the current rapid growth of derivatives markets, feedback effects might indeed become important as a potential destabilizing factor in financial markets. However, their study suffers from the same drawback as that by Kambhu (1998), that is, exposure is held constant throughout the investigation period. Last, but not least, we briefly consider several contributions to the theoretical literature of market mechanics. A survey of the literature on double auctions is provided by Friedman (1991). Early work on the trading process in financial markets was conducted by Mendelson (1982) who modeled random order placement with period clearing. Cohen, Conroy & Maier (1985) develop a model of a continuous auction, modeling limit orders, market orders, and or-

der cancellation as Poisson processes. Domowitz & Wang (1994) extend this model to arbitrary order placement and cancellation processes. All of these models are static in the sense that they treat the order process as static.

More recent models address price dynamics incorporating the feedback between order placement and price formation. The first model in this line of work was developed by Bak, Paczuski & Shubik (1997). Other models include those by Maslov (2000), Slanina (2001), Daniels, Farmer, Iori & Smith (2003), Bouchaud et al. (2002), Potters & Bouchaud (2003), and Challet & Stinchcombe (2003). A major contribution is that by Smith et al. (2003) who develop a dynamical statistical model for the continuous double auction. Their model makes predictions for basic properties of markets, such as price volatility, the depth of stored supply and demand versus price, the price impact function, and others. These predictions are based on properties of order flow and the limit order book, such as share volume of market and limit orders, typical order size, and tick size. As these quantities can be measured directly, there are no free parameters meaning that the model's predictions can be tested.

A theoretical model of large price changes based on the existence and trading behavior of large investors was developed by Gabaix et al. (2003). Hollifield, Miller & Sandås (2004) propose a model for limit order submission to study the trade-off between market and limit orders. A model for the explanation of limit versus market order placement and the resulting bid-ask spread was proposed by Cohen, Maier, Schwartz & Whitcomb (1981), whereas Parlour (1998) proposes a model relating limit order placement to the resulting price dynamics.

Several researchers have tried to explain the heavy tails of price distributions. Explanations based on statistical mechanisms so far have failed. Other explanations based on herding were cited above. Models where heavy tails emerge due to the interaction between fundamentalist and noise traders include those by Bak et al. (1997), Lux (1998), and Cont & Bouchaud (2000).

The impact of liquidity risk on the value of options has been investigated by Patie & Frey (2002), among others, and on asset prices by Acharya & Pedersen (2004).

This concludes our review of related literature. We will return to important findings during the description of our model and the discussion of our results. Certain properties should be kept in mind though. Order flow is highly auto-correlated and has long memory ranging over many days. Buy and sell order volumes tends to be symmetric. Liquidity is typically provided when it is valuable, and it is consumed when it is plentiful. Thus, a widening of the spread and a decrease in the depth at the best quotes are usually transient.

The price impact of an order can be expressed as a power law of volume. However, the exponent of the power fluctuates over time. Large price impacts are mainly due to granularity of limit prices in the order book, and not due to order volume. Negative correlation between limit order flow and returns tends

to reduce the price impact.

We now turn to the description of the markets we investigate and of our dataset.

3 Description of Markets and the Dataset

As mentioned in the Introduction, we develop a model of the relationship between options and underlying stock markets that we then test with data of the Swiss Stock Exchange (SWX) and Eurex, the German-Swiss derivatives exchange. We thus deem it important to describe certain characteristics of the two markets as well as the dataset we use. This is the aim of this section.

3.1 Markets

The cash market in our analysis is the SWX, whereas options are traded on Eurex, the derivatives exchange. We consider only Swiss stocks and options written on them. Options on Swiss stocks are traded almost exclusively on Eurex.

3.1.1 The Cash Market

We first describe the underlying cash market. The following borrows heavily from Bauer (2004) who provides a much more extensive description of the SWX market. We thus refer the reader to this paper and the references contained therein for details.

SWX is a pure order-driven market. Listed securities are permanently traded, with voluntary market making. We will consider the main segment of SWX—the blue chips—which together make up the Swiss Market Index (SMI). These securities are traded on virt-x, an electronic trading system based in London. The market model employed by virt-x is a limit order book model with price-time priority, that is, orders are executed in the order of price (first priority) and then submission time (second priority).

Trading hours are from 9:00 to 17:00 on every business day. Trading takes place continuously. There are rules for trading halts, comparable to the circuit breakers at the New York Stock Exchange, to avoid extreme market movements. If the price of a security would shift by 2% (25%) or more from the previous trade's or closing price, then trading is interrupted for 15 minutes (five minutes for stocks with prices of less than CHF10). No orders are executed during the last 10 minutes of the trading day. Instead, they are matched in a double auction that takes place at 17:30. This auction determines the official closing

prices for the day. From 17:30 until 22:00 and continuing from 6:00 to 9:00, the so-called pre-opening takes place. During the pre-opening, orders, both bids and offers, may be entered or deleted in the order book but there is no matching, that is, no trades are executed. A theoretical opening price is computed and displayed continuously to give guidance to traders. At the opening at 9:00, the opening price is determined and orders are matched according to the matching rules of SWX. In order to establish the opening price (or upon resumption of trading after an interruption), the highest-execution principle is employed. This means that the price is fixed such that the highest possible turnover is achieved. If the potential opening price would shift by 2% (25% for stocks with a price of less than CHF10) or more from the reference price, which is essentially equal to the previous traded price, opening is delayed by 15 minutes (5 minutes).

All trades with a size of less than CHF200,000 must be executed on the SWX trading system within official trading hours, whereas larger orders may be traded off-exchange. Such off-exchange transactions must be reported within 30 minutes of their execution.

Tick size for the various securities varies with stock price. Accordingly, large changes in stock price will induce a change in tick size, which in turn will change the spread. This is a particularity of SWX and is not the case in other markets like the New York Stock Exchange where tick size remains fixed.

The two main types of orders are limit and market orders. Besides, there are three other types of orders, hidden-size orders, accept orders, and fill-or-kill orders.

We now turn to a description of the continuous double auction with a special emphasis on the quantities that can be observed and measured.

The “life” of a trade begins with the submission of an order. We call the sequence of orders submitted to the market the *order flow*. There are two major types of orders. *Market orders* are requests to buy or sell a given number of shares immediately at the best available price. The other type of orders is called *limit orders*. A limit order specifies a quantity that the participant wants to buy or sell, and a limit price corresponding to the worst allowable price for the transaction. Limit orders do not necessarily result in an immediate transaction and are stored in a queue called the *limit order book*. Buy limit orders are typically called *bids* and sell limit orders are called *offers* or *asks*. There are separate queues for buy and sell orders. Order flow can be observed and we can extract, among others, submission time, type of order, quantity, limit price, delete time as well as the reason for deletion.

The limit price of the best buy order in the queue is often referred to as *best bid*, which we denote by $b(t)$, whereas the limit price of the best sell order is referred to as *best ask*, denoted by $a(t)$. The, typically non-zero, gap between the best bid and the best ask is called *spread*, $s(t)$, that is, $s(t) = a(t) - b(t)$. Prices are continuous but have discrete quanta called *ticks*. The minimum in-

terval that prices change on is the *tick size*, and we denote it by dp .

When market orders arrive they are matched against limit orders of the opposite sign in order of first price and then arrival time. As orders are placed for a varying number of shares, matching is not necessarily one-to-one. A high density of limit orders per price results in high *liquidity* for market orders, that is, it decreases the price movement when a market order is placed. We define market liquidity or *depth* as the volume of shares for sale at up to k ticks from the midquote. It can be interpreted as the volume necessary to move the price by k ticks. More liquid or deeper markets can accommodate larger trades for a given price impact.

Let $n(p, t)$ be the stored density of limit order volume at price p , which we will call the *depth profile* of the limit order book at any given time t . The total stored limit order volume at price level p is $n(p, t)dp$. For unit order size the shift in the best ask $a(t)$ produced by a buy market order is given by solving the equation

$$\omega = \sum_{p=a(t)}^{p'} n(p, t)dp \quad (1)$$

for p' . The shift in the best ask $p' - a(t)$ is the instantaneous *price impact* for buy market orders. A similar statement applies for sell market orders.

We will refer to a buy limit order whose limit price is greater than the best ask, or a sell limit order whose limit price is less than the best bid, as a *marketable limit order*. Such limit orders result in immediate transactions, with at least part of the order executed instantly.

A limit order can also be removed from the queue by being cancelled, which can occur at any time. The best price may change as new orders arrive or old orders are cancelled.

The order book can be completely reconstructed from order and trade data. The properties described above such as spread and depth profile can thus be observed at any time.

Once two orders are matched, a trade is executed “producing” a price. On most exchanges, including SWX, orders can be matched outside the order book. However, these trades have to be reported to the exchange so that they are also observable. We call trades executed in the trading system *on-book* trades and trades executed outside the trading system *off-book* trades. Trade data allows to measure time, volume, price, and trade type, among others.

For markets operating a fully electronic trading system, the life cycle of an order can be traced from submission to deletion, and all relevant data created inbetween can be observed and measured.

3.1.2 Eurex

Next, we describe the options market under consideration and the data that can be observed. At the time of writing, Eurex is the largest derivatives ex-

change in the world. It operates a fully electronic trading system similar to the one employed by SWX. Trading is also order-driven but with market makers who continuously provide quotes for the products traded.

Eurex offers options on all shares contained in the Swiss Market Index (SMI) as described in Eurex (2005). For the options we consider, one option contract is based on either 10 or 100 shares. Contract months for the options we consider are the next 60 months; more precisely, the three nearest calendar months, the three following quarterly months of the March, June, September, and December cycle thereafter, and the four following semi-annual months of the December cycle thereafter, and the following annual months of the December cycle thereafter.

For every contract month, contracts with different exercise (strike) prices are available. The last trading day of options, before expiry, is the third Friday of each contract month, if this is an exchange trading day; otherwise, it is the exchange trading day immediately preceding that Friday.

At the end of every trading day, the exchange calculates a settlement price for every contract. This price is used for margining. The calculations are based on various option pricing models. The volatility used is the implied volatility derived from daily closing prices of the respective contract, as described in Eurex (2003). The settlement price, which is also reported, can be used to “back out” the implied volatility for the contract.

Eurex provides historical data on trading including all trades executed on the exchange as well as daily market statistics. The latter contain, among others, the open interest for every single contract. Based on this information the aggregate amount of options outstanding on a given underlying can be computed.

3.2 The Dataset

We now provide a description of our dataset. As before, we first look at the data of the underlying market (SWX), and subsequently discuss the data of the options market (Eurex). We consider data for the period from January to December 2002.

3.2.1 The Underlying Market

Our description of the dataset is organized along the life cycle of an order. This means that we will start with order flow, then turn to the order book, and finally discuss transaction data. Both, order and transaction data are on event level, that is, the data comprises every order submitted to SWX and every transaction executed in the respective period.

Table 1
Descriptive statistics on order volume

This table presents descriptive statistics on order volume for Novartis (NOVN) for the period 1 January to 31 December 2002. Order volume is highly skewed and fat-tailed.

	Buy Orders	Sell Orders
Number of orders	8,806,700	9,087,630
Minimum	1	1
Maximum	50,000,300	250,000,000
1. Quartile	780	650
3. Quartile	5,989	5228
Mean	4,683	4,712
Median	2,000	2,000
Sum	$4.124 \cdot 10^9$	$4.282 \cdot 10^9$
Standard deviation	56,405	26,262
Skewness	812	949
Kurtosis	706,688	903,637

3.2.1.1 Order Flow We consider all orders submitted to SWX between 1 January and 31 December 2002. The dataset comprises 1,645,246 limit orders and 144,187 market orders, a total of 1,789,433 orders. 50.0% of the limit orders are orders to buy, whereas the respective fraction of market orders is 40.3%. The share of marketable limit orders is 21.65%, meaning that 27.9% of all orders submitted (and 50.4% of those orders traded against) lead to a trade immediately.

44.8% of all limit orders were not traded against at all, 32.7% were involved in one trade, and 27.4% were fully matched in a trade. Of the market orders, 2.2% did not result in a single trade, 78.2% were involved in one trade, and 15.2% were fully executed in one trade.

46.5% of limit orders were deleted by the respective trader and 4.5% expired. The respective fractions of market orders are 2.2% and 0% (the number of market orders that were manually deleted equals the number of market orders that did not result in a trade, see above). Table 1 shows descriptive statistics of order volume, separated into buy and sell orders. The mean volume per order was 4,683 shares for buy orders and 4,712 shares for sell orders. The distribution is highly skewed, though, and has fat tails, as skewness and kurtosis show. The average volume per day in the period considered was 16.6 million shares for buy orders and 17.2 million shares for sell orders, or approximately 0.06% of the free float at the time.

Fig. 1. Empirical cumulative distribution function of order volume (log scale)

The figure below shows the empirical cumulative distribution functions of buy and sell order volume, respectively, for Novartis (NOVN) on log scale. The dashed line shows the cumulative distribution function of a normal distribution. A close observation of the plots shows that the empirical distribution function has heavier tails than that of the normal distribution.

Fig. 2. Correlelograms for buy and sell order volume

Below we show autocorrelation and cross-correlation plots of buy and sell order volume. Order volume is aggregated over five-minute intervals to make the data comparable.

Figure 1 displays the empirical cumulative distribution function of buy and sell orders for the entire period in log scale. The dotted line shows the cumulative distribution function of the normal distribution. The empirical cumulative distribution function is close to the cumulative distribution of the normal distribution. However, it has fatter tails. In Figure 2 we show autocorrelation and cross-correlation plots of buy and sell order volume. Data is reported in “event time”. In order to make buy and sell volume series comparable, we aggregate the data over five-minute intervals. As detected previously by several researchers, order volume is highly autocorrelated. It also seems to be cross-correlated.

Table 2
Estimation of VAR(1) model of order volumes

The table below shows the results of a vector autoregression model of order volume. The underlying dataset comprises NOVN order volume for January 2002 aggregated over five-minute intervals on log scale. We differentiate four types of orders: limit buy, limit sell, market buy, and market sell orders.

Bold face indicates a p -value of 0.01 or lower for the respective coefficient. The numbers in brackets below the coefficients show the standard errors.

	V_t^{lb}	V_t^{ls}	V_t^{mb}	V_t^{ms}
Valid cases	1,510	1,510	1,510	1,510
Degrees of freedom	1,505	1,505	1,505	1,505
Residual SS	1,116.11	949.51	4,005.11	3,568.01
R^2	0.296	0.307	0.094	0.083
\bar{R}^2	0.294	0.305	0.092	0.080
	V_t^{lb}	V_t^{ls}	V_t^{mb}	V_t^{ms}
Intercept	4.357723 (0.277663)	4.718117 (0.256103)	1.807999 (0.525983)	3.430203 (0.496452)
V_{t-1}^{lb}	0.286509 (0.030367)	0.165572 (0.028009)	0.070798 (0.057525)	0.038039 (0.054295)
V_{t-1}^{ls}	0.290312 (0.032861)	0.377091 (0.030309)	0.255217 (0.062248)	0.131567 (0.058754)
V_{t-1}^{mb}	0.053768 (0.013720)	0.046243 (0.012655)	0.189198 (0.025990)	0.051701 (0.024531)
V_{t-1}^{ms}	-0.000235 (0.014626)	0.014722 (0.01349)	0.046266 (0.027706)	0.221662 (0.026151)

SS: sum of squares

\bar{R}^2 : adjusted R^2

Fig. 3. Probability density of order volume aggregated over different time intervals

The figure below shows density plots of order volume aggregated over different time intervals. Buy order volume is shown in the left column and sell order volume in the right column. The top row shows volume aggregated over five-minuted intervals, whereas the bottom row shows volume aggregated over one-hour intervals. The dashed line shows the probability density of the normal distribution.

In Table 2 we show the results of a vector autoregression. We categorize orders according to buy and sell as well as limit and market. We denote by V_t^{lb} , V_t^{ls} , V_t^{mb} , and V_t^{ms} the log of the volume of limit buy, limit sell, market buy, and market sell orders, respectively (unless noted otherwise, when we speak of volume we will now always mean the logarithm of volume). Our response variables are

$$V_t = \begin{pmatrix} V_t^{lb} \\ V_t^{ls} \\ V_t^{mb} \\ V_t^{ms} \end{pmatrix}.$$

We regress each of the explanatory variables V_{t-1}^{lb} , V_{t-1}^{ls} , V_{t-1}^{mb} , and V_{t-1}^{ms} . The model satisfies the definition of a seemingly unrelated regressions model, as described for example in [Gourieroux & Jasiak \(2001\)](#). The results show that there is “causality” within and across limit orders, but less so for market orders.

[Ranaldo \(2003\)](#) analyzes intraday patterns of trading volume on SWX. A large number of buy limit orders are submitted right after the market opening. Limit order submission decreases around noon. Sell limit orders are lagged by about one hour. Trading volume in the afternoon, between 14:00 and 15:30, is influenced by trading activity and disclosures in the U.S. markets. During this time of the day, market depth is reduced and the bid-ask spread widens. Another interesting property of order volume is the fact that its distribution changes when orders are aggregated over time. Figure 3 shows the volume of buy and sell orders aggregated over five-minute and one-hour intervals, respectively. It seems that the larger the aggregation interval, the closer the distribution approaches a normal distribution. Note, though, that the empirical distribution of volume retains its fat tails even when volume is aggregated over time. In fact, order volume in our sample aggregated over one-hour intervals is gamma-distributed.

Important properties of order flow to keep in mind are (1) a high percentage (50.4%) of all orders are marketable orders, meaning that they are traded against immediately, (2) there is considerable autocorrelation across all types of orders and cross-correlation between limit buy and limit sell orders, (3) order volume aggregated in one-hour intervals is gamma-distributed.

Table 3
 Estimation of the price impact function

This table presents the results of the estimation of the price impact function. We regress the price change caused by a trade, Δp , on transaction volume, V^{tr} , on logarithmic scale, that is, $\Delta p = \beta V^{tr}$. Both, price change and transaction volume are normalized.

	Δp
Degrees of freedom	25,078
Residual SE	0.7785
R^2	0.3939
\bar{R}^2	0.3939
F statistic	$1.63 \cdot 10^4$
p -value	0.000
	Δp
V^{tr}	0.627603 (0.004916)

\bar{R}^2 : adjusted R^2

3.2.1.2 Order Book and Price Impact We now turn to the order book and the price impact of trades. We noted previously that the effect of an order on price depends on the state of the order book at the time when the order is submitted. The effect of orders on price is captured by the price impact function, a mapping from volume to price change. We investigate all on-order book trades in January 2002, a total of 25,931 data points, to find the relationship between price change and transaction volume. We recall that the price impact function is defined as

$$I_t(Q) = \mathbf{E}[\Delta p \mid Q, t], \quad (2)$$

where Q denotes the logarithm of the volume of the trade, and Δp denotes the (log) return. We find that $I_t(Q)$ is a concave function of volume that can be fitted by a power law. The results of the estimation are reproduced in Table 3. We find the relation $\Delta p = Q^{0.62}$. We note that the exponent of 0.62 is between 0.5, found by Plerou, Gopikrishnan, Gabaix & Stanley (2002), Gabaix et al. (2003), and Potters & Bouchaud (2003), and 0.76, found by Weber & Rosenow (2003). It should also be noted that the exponent varies considerably. For the period we consider, the standard deviation is approximately 0.07.

The exponent does not display any autocorrelation nor does it seem to be (systematically) cross-correlated with transaction volume across periods. There is strong evidence, as demonstrated above, that there is an increasing

relationship between volume and price changes, albeit concave. In other words, price changes are higher the larger the transaction volume. We have thus established a link between order flow and price changes. Now we look at the other direction, namely the relationship between price changes and (future) order flow.

We conduct a vector autoregression of order one where we link changes in price to order volume. As expected from previous analyses, volume is autocorrelated. Buy volume is also cross-correlated with price change. In other words, it seems that, at least in case of buy orders, price change leads order volume.

Table 4
Estimation of VAR(1) model of order volume and price change

The table below shows the results of a vector autoregression model of order volume and price change. The underlying dataset comprises NOVN transaction volume for the period from January to December 2002 aggregated over one-hour intervals. Both price and order volume are on log scale. We differentiate between buy order volume and sell order volume. Bold face indicates a p -value of 0.01 or lower for the respective coefficient. The numbers in brackets below the coefficients show the standard errors.

	V_t^b	V_t^s	Δp_t
Valid cases	2,524	2,524	2,524
Degrees of freedom	2,521	2,521	2,521
Residual SS	5,167.26	4,216.77	0.079255
R^2	0.182	0.090	0.010
\bar{R}^2	0.181	0.090	0.009
	V_t^b	V_t^s	Δp_t
Intercept	0	0	0
V_{t-1}^b	0.516248	0.243844	0.000381
	(0.002579)	(0.023279)	(0.000101)
V_{t-1}^s	0.475409	0.758457	-0.000378
	(0.025407)	(0.022952)	(0.000100)
Δp_{t-1}	11.668416	3.772844	0.035895
	(5.420676)	(4.896814)	(0.021229)

SS: sum of squares

\bar{R}^2 : adjusted R^2

Fig. 4. Open interest in options on NOVN

The figure below shows aggregate open interest for call and put options on NOVN for the period from January to December 2002. The sudden drops represent contract expiries. These drops are most significant at the end of every quarter. The solid line shows open interest in call options whereas the dashed line shows open interest in put options.

Such a (statistical) relationship has been documented previously, for example by *Gourieroux & Jasiak (2001)*. It is the starting point for the analysis we present in the following where we suggest an economic explanation for the link. More precisely, we investigate whether order flow can be, at least partly, explained by dynamic hedging. Before we present our analysis in detail we briefly describe the data of the options market that we use.

3.2.2 Eurex

The Eurex dataset contains daily statistics on options contracts including the open interest per contract as well as the settlement price. The latter allows us to back out implied volatilities. Figure 4 displays aggregate open interest in options on NOVN for the period from January to December 2002. Average open interest in call options was 516,087 million contracts, and 447,630 contract for put options. The average total open interest was 963,717 contracts or 96.3 million shares, compared to an average daily trading volume of 7.0 million shares in the period we consider.

4 Research Methods and Model

Having described our dataset we now turn to the description of our research methods as well as our model.

4.1 Research Methods

As we mentioned in the Introduction already, we would like to take a rather “agnostic” approach in our analysis. This means that we would like to make as few assumptions as necessary with regards to behavior of traders, utility functions, beliefs, etc. Most importantly, we want our model to be testable. Most models in the traditional market microstructure literature depend significantly on the choice of certain free parameters, that is, parameters which cannot be

observed. These include the share of noise traders vs. informed traders, the distribution of information, the shape of traders' utility functions, among others. The serious drawback of many of these models is that they cannot be tested.

Our approach is different. Our model only contains parameters that can be observed and measured. Therefore, the predictions of our model can be validated with data, and the model can be falsified. The choice as well as the specification of our model is in this spirit.

The working (null) hypothesis for our analysis is that order volume in the underlying market cannot be predicted by hedging demand. The latter can be computed with the data on open interest. If this hypothesis is false, we will have evidence that there does exist a systematic relationship between price change and order volume, via dynamic hedging, representing a potential feedback mechanism.

The observations which we will use to test the predictions of our model are provided by the SWX and Eurex data. The model and its description is described in the remainder of this section. We then discuss the results of our tests of the model in the subsequent section.

4.2 Model

The purpose of our model is to explain future order volume in the stock market by hedging demand arising from options trading. "Producers" of options, holding a short position in the contract, typically hedge their risk by buying or selling shares in the underlying market. As prices of the underlying assets tend to fluctuate, a trader's position typically only remains hedged for a relatively short period of time. Thus, the hedge has to be adjusted periodically. In other words, the trader will continuously trade in the underlying market creating order flow.

In our analysis of dynamic hedging, we make a number of assumptions. First, we assume that one side in the options market (the long side or "customers") do not dynamically hedge their options positions because doing so would negate the investment or hedging objective that motivated the purchase of the option. Thus, we need consider only dealers' ("producers" or the short side) hedging demands. Second, we assume that producers in the market compute their hedging demand based on the Black-Scholes framework. This assumption is motivated by conversations with traders in the market according to which most market participants do indeed use the Black-Scholes framework for pricing and hedging of this type of options instead of one of its refinements. Third, we assume that all producers in the market have the same exposure profile. We compute hedging demand for a global portfolio containing all the open interest in the market on a given day. This approach might over-estimate true hedging demand as some producers might hedge their exposure internally, for example,

with other lines of business. On other hand, it might also underestimate true hedging demand if, for example, true exposures are significantly different from each other. Considering these effects, we think that our approach provides a good approximation to true hedging demand. Fourth, we assume that producers hedge their positions completely in every time period. More precisely, we assume that they employ the Black-Scholes formula to compute the delta of their position and hold the corresponding number of shares in the underlying in order to be “delta-neutral”, that is, fully hedged.

Our computation of hedging demand is based on what we call the *global position*. This is a hypothetical position containing the aggregated open interest in the various option contracts on a given stock outstanding at a given time. We construct the global position based on (daily) data on the open interest in the various contracts. We do not adjust the open interest during the day as average trading volume is very small, less than 1% of aggregate open interest. We then compute the delta of the global position. At this point, we briefly recapitulate the Black-Scholes methodology for pricing and hedging of options as developed by Black & Scholes (1973) and Merton (1973).

The Black-Scholes formula for the price of a European call option on a non-dividend paying stock is given by

$$c_t = S_t N(d_1) - X e^{-r(T-t)} N(d_2),$$

where

$$d_1 = \frac{\ln\left(\frac{S_t}{X}\right) + \left(r + \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}},$$

and

$$d_2 = d_1 - \sigma\sqrt{T-t},$$

and t is the current time, S_t is the stock price at time t , X is the exercise price of the option, r is the risk-free interest rate, σ is the volatility of the stock price, and T is the expiry date of the option. $N(x)$ denotes the cumulative probability distribution function for a variable that is normally distributed with a mean of zero and unit standard deviation. The price of a European put option is given by

$$c_t = X e^{-r(T-t)} N(-d_2) - S_t N(-d_1).$$

The delta of an option is defined as the rate of change of the option price with respect to the price of the underlying asset. Formally, delta is the partial derivative of the call price with respect to the underlying stock price, that is,

$$\Delta = \frac{\partial c}{\partial S}.$$

Because delta changes over time, an investor’s position remains delta-hedged (or delta-neutral) for only a relatively short period of time. Thus, the hedge has to be adjusted periodically. This is known as *rebalancing* or *delta hedging*.

For a European call option on a non-dividend paying stock, it can be shown from the Black-Scholes formula that

$$\delta^C = N(d_1).$$

Using delta hedging for a short position in a European call option, therefore, involves keeping a long position of $N(d_1)$ shares at any given time. For a European put-option on a non-dividend paying stock, the delta is given by

$$\delta^P = N(d_1) - 1.$$

The delta of an options portfolio dependent on a single asset, with price S , is

$$\Delta = \frac{\partial \Pi}{\partial S},$$

where Π is the value of the portfolio.

The delta of the portfolio can be calculated from the deltas of the individual options in the portfolio, as described for example in Hull (2002). If a portfolio consists of an amount, ω_i , of option i with $i = 1 \leq i \leq n$, the delta of the portfolio is given by

$$\Delta = \sum_{i=1}^n \omega_i \delta_i, \quad (3)$$

where δ_i is the delta of the i th option. This formula can be used to calculate the position in the underlying asset necessary to carry out delta hedging. When this position is implemented, the delta of the portfolio, including the new position in the asset, is zero so that the portfolio is delta neutral. All equations above for the value of options give the value of a single option. In order to compute the Δ of the global position in terms of number of shares, we have to multiply the δ of the various contracts by the contract size.

Our computation of Δ , the delta of the global position, is based on a complex algorithm. This algorithm extracts the data relevant for a given period from the datasets. It then calculates implied volatilities on the basis of the previous day's settlement price of the options using standard numerical estimation techniques. To compute the relevant interest rate, the algorithm fits a yield curve to the Swiss interbank market rates prevailing at the time. Based on these parameters, the algorithm then computes the delta of the global position at time t , Δ_t .

The hedging demand in terms of shares for a given period t is given by

$$H_t = \Delta_t - \Delta_{t-1}. \quad (4)$$

Intuitively, this means that hedging demand in a given period is equal to the difference between the delta of the current and the previous period. In other words, the delta at the end of period $t - 1$, Δ_{t-1} , changes to Δ_t at the end of period t . The trader now needs to trade $\Delta_t - \Delta_{t-1}$ shares in the next period, $t + 1$, to adjust his hedge. This quantity is given by H_t .

Fig. 5. Order and hedging volume for NOVN

The figure below shows order volume, V^b and V^s , as well as (theoretical) hedging volume, H^b and H^s , for the period from January to December 2002. Order volume is aggregated over one-hour intervals. Hedging volume is based on rebalancing in one-hour intervals.

Fig. 6. Density plots for order and hedging volume for NOVN

The figure below displays density plots of order volume, V^b and V^s , as well as (theoretical) hedging volume, H^b and H^s , for the period from January to December 2002. Order volume is aggregated over one-hour intervals. Hedging volume is based on rebalancing in one-hour intervals.

Fig. 7. Goodness-of-fit plot

The figure below displays kernel density estimates (solid lines) and density functions of a fitted gamma distribution (dashed lines) for buy and sell order volume. The Kolmogorov-Smirnov test for goodness of fit indicates that the gamma distribution describes these two quantities well.

Let V_t denote aggregate trading volume in the stock market for the period between $t - 1$ and t . Our aim is to explain V_t with hedging volume. This means that we would like to relate V_t to H_{t-1} .

We thus propose a model of the form

$$V_t = \beta_0 + \beta_1 H_{t-1}. \quad (5)$$

Our goal is to estimate the parameters β_0 and β_1 . In the stock market there are two order flows, one for buy orders and one for sell orders. We denote buy order flow by V^b and sell order flow by V^s . Whenever H is positive, we will relate it to V^b , and vice versa whenever it is negative.

Figure 5 displays buy and sell order volume as well as hedging volume for the period from January to December 2002. In Figure 6 we display density plots for these quantities. Summary statistics are presented in in Table 5.

Observation of the density plots and summary statistics suggests that the distribution of order volume can be described well by a gamma distribution. We thus fit a gamma distribution to both buy and sell order volume. The scale parameters are 534.21 and 569.87 for buy and sell order volume, respectively, and the rate parameters are 37.71 and 40.35, respectively. In Figure 7 we show estimated density functions and the density function of the fitted gamma distribution. We apply the Kolmogorov-Smirnov test with the null hypothesis that V^b and V^s are gamma distributed to validate the choice of the distri-

Table 5
Descriptive statistics on order and hedging volume

This table presents descriptive statistics on order volume, V^b and V^s , as well as on hedging demand, H^b and H^s , for the period from January to December 2002. Order volume is aggregated over one-hour intervals. Hedging volume is based on rebalancing in one-hour intervals.

	V^b	V^s	H^b	H^s
Number of time periods	1,140	1,144	1,140	1,144
Minimum	12.25	12.33	7.65	9.24
Maximum	17.76	15.78	15.85	16.04
1. Quartile	13.76	13.69	12.79	12.78
3. Quartile	14.57	14.54	14.09	14.09
Mean	14.17	14.12	13.44	13.44
Median	14.15	14.12	13.48	13.46
Sum	16,150	16,155	15,327	15,371
Standard deviation	0.61	0.59	0.91	0.91
Skewness	0.28	0.00	-0.26	-0.05
Kurtosis	1.08	-0.33	0.77	-0.19

bution. p -values significantly larger than 0.1 for both V^b and V^s confirm the goodness of the fits.

All parameters that enter our model, the predictors and the responses can be measured based on our dataset. In order to keep our model tractable and economically meaningful we use standard linear estimation techniques as described below.

A preliminary analysis of the data shows that hedging demand fluctuates much more than order volume. A “visual fit” of Equation 5 suggests a relationship between order volume and hedging demand of the form $V = 7.5 + 0.5H$. This would mean that hedging demand can indeed explain a significant part of order volume.

We now turn to a more formal analysis. The distribution of V^b and V^s , the response variables, obviously influences the choice of the statistical test we will employ. As the two quantities are gamma distributed we choose the Generalized Linear Model (GLM) framework, as described in Dobson (2002) for example, with the identity function as the link function.

We regress the following equation:

$$V_t = g(\beta_0 + \beta_1 H_{t-1}) + u_t, \quad (6)$$

where u_t denotes the error term in period t and g is the inverse of the link function, as described in more detail in Section 5. The results of the regressions as well as a discussion of them are the topics of the next section.

5 Results

As described in the previous section, we test the null hypothesis that hedging demand is not systematically related to future order volume in the underlying market. We first present the results of our tests and then discuss their robustness with regards to different model specifications.

5.1 Tests

In our tests, we analyze systematic effects of hedging demand, our predictor variables, on future order volume, our dependent variables. As the dependent variables are gamma distributed, we select the Generalized Linear Model (GLM) for our analysis. GLM generalizes the standard general linear model in two respects: first, whereas in the general linear model the dependent variables are expected to follow the normal distribution, in GLM the distribution of the dependent variables can be (explicitly) non-normal; secondly, the dependent variable values are predicted from a linear combination of predictor variables, which are “connected” to the dependent variable via a link function. In the general linear model, this link function is simply the identity function. When performing the tests, we explicitly specify that the dependent variables are gamma distributed, as indicated by our analyses described above. As link function we initially choose the identity function but vary this function later when checking for robustness.

More formally, we test two equations,

$$V_t^b = g(\beta_0 + \beta_1 H_{t-1}^b) + u_t, \quad (7)$$

and

$$V_t^s = g(\beta_0 + \beta_1 H_{t-1}^s) + u_t, \quad (8)$$

where u_t denotes the error term, and g is the inverse of the link function. Equation 7 relates positive hedging demand to buy order volume, and Equation 8 relates negative hedging demand to sell order volume in the respective periods.

We perform two regressions for each of the two equations. In Tests 1a and 2a we explicitly set β_0 , the intercept term, to zero. In Tests 1b and 2b we allow the intercept term to be different from zero. Table 6 presents the results of the regression. In all four tests, the p -values for every parameter are below

Table 6
Results of GLM regression

The table below presents the results of four regressions using the Generalized Linear Model for gamma distributed dependent variables with the identity function as link function. We test the following four equations,

$$\begin{aligned} \text{(Test 1a)} \quad V_t^b &= \beta_1 H_{t-1}^b + u_t, \\ \text{(Test 1b)} \quad V_t^b &= \beta_0 + \beta_1 H_{t-1}^b + u_t, \end{aligned}$$

where

$$\begin{aligned} V_t^b &: \text{Buy order volume in period } t, \\ H_t^b &: \text{Hedging demand in period } t, \\ u_t &: \text{Error term in period } t, \end{aligned}$$

as well as

$$\begin{aligned} \text{(Test 2a)} \quad V_t^s &= \beta_1 H_{t-1}^s + u_t, \\ \text{(Test 2b)} \quad V_t^s &= \beta_0 + \beta_1 H_{t-1}^s + u_t, \end{aligned}$$

where

$$\begin{aligned} V_t^s &: \text{Sell order volume in period } t, \\ H_t^s &: \text{Hedging demand in period } t, \\ u_t &: \text{Error term in period } t, \end{aligned}$$

Bold face indicates p -values below 0.01. Standard errors are shown in brackets.

	Test 1a	Test 1b	Test 2a	Test 2b
Intercept	0	12.31339	0	12.59247
		(0.26322)		(0.25510)
H_{t-1}^b	1.05815	0.13790		
	(0.00236)	(0.01956)		
H_{t-1}^s			1.05536	0.11381
			(0.00231)	(0.01896)
Residual SS	1,210.4	401.0	1,223.8	387.1
AIC	3,346.9	2,074.4	3,340.6	2,014.2

SS: sum of squares

0.001. Based on these results, we have to reject our null hypothesis that the regression coefficients are zero.

We perform several tests to check for the validity of the test results. We employ the Shapiro-Wilk test for the null hypothesis that residuals are normally distributed. We obtain p -values of more than 0.1 for each of the four tests, implying that we cannot reject the null hypothesis. The KPSS test indicates that residuals are stationary. A Box test on the residuals indicates that residuals of Tests 1a and 2a are independent; however, we have to reject independence in case of Tests 1b and 2b.

Given the values of the Akaike Information Criterion (AIC) and the residual sums of squares for the four tests, we conclude that Tests 1b and 2b have a higher “explanatory power” than Tests 1a and 2a, and therefore discard Tests 1a and 2a.

The result of the Box test, that residuals in Tests 1b and 2b are not independent, was to be expected given the values of the coefficients β_0 and β_1 . The values of β_1 in the two tests are very small meaning that the “contribution” of H is very low. Given that order flow, represented by V , is autocorrelated, as we describe in Section 3, we would expect the residuals in Tests 1b and 2b to be autocorrelated, too.

In order to account for autocorrelation among the residuals, we apply the Generalized Least Squares (GLS) model to Equations 7 and 8. We fit an AR(1) model to the residuals of Test 1b and 2b and find that the autocorrelation coefficients are highest for the first lag and much smaller for all remaining lags. We now design two additional tests, Tests 1c and 2c, where we specify an autocorrelation function for the residuals of order 1.

Given that the dependent variables in our equations are gamma distributed, GLS might not be consistent. However, we saw in Section 3 that order flow is approximately normally distributed. To justify the use of the GLS model, we run it on Tests 1b and 2b but without specifying a correlation structure for the residuals. The differences in the coefficient estimates are less than 0.02 and 0.003 for β_0 and β_1 , respectively. We conclude that the slight deviation of the distribution of the dependent variables from the normal distribution has a negligible effect on the coefficient estimates. Table 7 displays the results of Tests 1c and 2c. As expected, the values of β_1 are not different from zero in both tests. This means that the explanatory power of H is negligible. In other words, we do not find a systematic relationship between hedging demand and future order volume.

As we mentioned previously, various research has demonstrated that (large) price changes in a market are not necessarily related to large volumes. In other words, large price changes are often caused by small volumes. This fact indicates that the magnitude of price change is determined by the interplay between volume and price impact. The latter depends on the state of the order book, in particular, the granularity of limit prices. We discussed this extensively in Section 2. In Section 1, we pointed out that a positive feedback of hedging on price requires both, that hedging demand “drives” order volume

Table 7
Results of GLS regression

The table below presents the results of two regressions using the General Least Squares model. We test two equations,

$$\text{(Test 1c)} \quad V_t^b = \beta_0 + \beta_1 H_{t-1}^b + u_t,$$

where

V_t^b : Buy order volume in period t ,

H_t^b : Hedging demand in period t ,

u_t : Error term in period t ,

as well as

$$\text{(Test 2c)} \quad V_t^s = \beta_0 + \beta_1 H_{t-1}^s + u_t,$$

where

V_t^s : Sell order volume in period t ,

H_t^s : Hedging demand in period t ,

u_t : Error term in period t ,

Bold face indicates p -values below 0.01. Standard errors are shown in brackets.

	Test 1c	Test 2c
Intercept	13.725081	14.107111
	(0.223578)	(0.220710)
H_{t-1}^b	0.032793	
	(0.016477)	
H_{t-1}^s		0.001035
		(0.0162694)
Residual SS	422.7	399.2
AIC	1,819.1	1,801.5

SS: sum of squares

in the underlying market, and that price impact is non-negligible. We do not find a systematic relationship between hedging demand and order volume as discussed earlier in the section. This does not mean, though, that hedging demand at times makes up a large part of or “drives” order volume in the underlying market. If price impact in such periods is large, hedging demand might, in such periods, create positive feedback.

Fig. 8. Hedging demand and price impact

The figures below show the ratio of hedging demand H in a given period t and order volume in period $t + 1$ (solid line) as well as the exponent of the price impact function (dashed line). The figure on the right shows the respective quantities for buy volume, and the figure on the left shows the quantities for sell volume. All quantities are normalized.

We briefly investigate whether there is a systematic relationship between the share of hedging demand of total order volume on the one hand and price impact on the other hand. Our null hypothesis is that there is no such systematic relationship. To test our null hypothesis, we regress the exponent of the price impact function, as discussed in Section 3, to the ratio H/V . For this analysis, we normalize both quantities. A high exponent of the price impact function means that volume causes comparably larger price changes. A high ratio H_{t-1}/V_t means that hedging demand drives order volume in the underlying market. If there is a systematic (positive) relationship between these two quantities, then this will indicate that at times when hedging demand dominates order volume, price impact is high, meaning that hedging might create a positive feedback. In Figure 8, we show the ratio of hedging demand H in a given period t and order volume in period $t + 1$ as well as the exponent of the price impact function for January 2002². Both quantities are normalized. One clearly sees that while at times there is a close relationship between the ratio H/V , this is not the case for values more than one standard deviation away from the mean. We now test this relationship more formally. We denote the exponent of the price impact function by γ , that is $\Delta p = Q^\gamma$. We test the following equation:

$$\gamma_t = \frac{H_{t-1}}{V_t} + u_t, \quad (9)$$

where u_t denotes the error term. We test this equation separately for buy and sell order volume. Table 8 presents the results of the tests. None of the coefficients is significant. This implies that we do not have to reject our null hypothesis that there is systematic impact between the ration of hedge demand and order volume, and price impact.

Let us recapitulate. As we described in Section 1, hedging volume might trigger a positive feedback effect on price if three prerequisites are met all at once: (1) the delta of the options portfolio is positive; (2) hedging demand drives order volume in the underlying market; and (3) the price impact of orders is non-negligible. For the time period we investigate, is fluctuates and is positive from time to time. We do not find evidence for that fact that hedging demand is systematically related to order volume in the underlying market, for example, that it represents a constant share of order flow, or that is responsible for

² Computation of the price impact function is very resource-intensive. Due to limit computing resources, we only present values for January 2002.

Table 8
Results of LM regression

The table below presents the results of two regressions using the standard linear model. We test two equations,

$$\text{(Test 1d)} \quad \gamma_t = \beta_0 + \beta_1 \frac{H_{t-1}^b}{V_t^b} + u_t,$$

where

- γ_t : Exponent of price impact function in period t ,
- V_t^b : Buy order volume in period t ,
- H_t^b : Hedging demand in period t ,
- u_t : Error term in period t ,

as well as

$$\text{(Test 2d)} \quad \gamma_t = \beta_0 + \beta_1 \frac{H_{t-1}^s}{V_t^s} + u_t,$$

where

- γ_t : Exponent of price impact function in period t ,
- V_t^s : Sell order volume in period t ,
- H_t^s : Hedging demand in period t ,
- u_t : Error term in period t ,

Bold face indicates p -values below 0.01. Standard errors are shown in brackets.

	Test 1d	Test 2d
Intercept	0.00000 (0.10050)	0.00000 (0.01050)
H_{t-1}^b/V_t^b	0.006612 (0.01010)	
H_{t-1}^s		0.08699 (0.01056)
Residual SS	1.00500	1.00200
R^2	0.00004	0.00757

SS: sum of squares

the fluctuations in order volume. Finally, at times when hedging demand does contribute a large part of order volume in the underlying market, we do not find that price impact is comparatively high in such periods. This leads us to the conclusion that there does not seem to be a systematic influence of hedging demand on price in the form of a positive feedback.

5.2 Robustness of results

We perform a large variety of tests to confirm that our results are robust across model specifications. We now describe those tests that we consider most important.

One dimension in our robustness tests is the length of a period, and thus the frequency of rebalancing. In Tests 1a–c and 2a–c, the time interval is one hour. We chose alternative intervals including 30 minutes, 2 hours, 3 hours, and half-days. The results do not change significantly.

Another dimension in our robustness tests is the stock price, S_t , based on which hedging demand in period t is calculated. In Tests 1a–c and 2a–c, we select the last reported price in the respective time interval. To account for possibly considerable changes of the stock price during time intervals, we selected alternative prices including mean and median price in the respective interval and more complicated specifications. However, results remain essentially the same.

We also use different transformations of dependent and explanatory variables—instead of the logarithmic transformation in 1a–c and 2a–c—including a linear transformation and a higher-order polynomial transformation. Again, we find the same results.

Finally, we use different link functions in the GLM model, in particular, the inverse function and the log. As before, we do not find significantly different results.

From the results of these robustness tests we conclude that our findings are robust with regards to model specifications.

6 Conclusion

In this paper, we analyze to what extent hedging demand arising from options trading can explain future order volume in the underlying stock market. In other words, we investigate whether there is a systematic relationship between hedging demand and order flow.

As we pointed out in Sections 1 and 2, it is often argued that options have potentially detrimental effects on the respective underlying markets. Dynamic hedging, that is, the continuous rebalancing of options portfolios, might in-

deed create positive feedback effects in the underlying market. Price changes in the underlying market require an adjustment of the hedging portfolio by buying or selling shares in the underlying. If the delta of the options portfolio is positive, an increase in the price of the underlying market will require the purchase of additional shares and vice versa for a decrease in the share price. If there is a non-negligible price impact of the resulting trades, then indeed rebalancing will be likely to create a positive feedback effect.

For such a feedback effect to actually set in, hedging demand needs to make up a large part of order volume. In addition, the price impact of trading at these times needs to be significantly different from zero. If either of these two requirements is not fulfilled, then rebalancing is unlikely to create a feedback effect. If it is not hedging demand that “drives” order volume, then price changes will mostly originate from other sources. Secondly, if the price impact of trading in the periods when hedging drives order volume is zero, then, again, a positive feedback effect is unlikely to occur. We investigate in this paper whether order volume in the underlying market can be predicted by hedging demand. If there is no systematic relationship between the two quantities, then rebalancing will be unlikely to drive price changes in the underlying market.

Such a relationship between derivatives and underlying markets was previously investigated by Kambhu (1998) and Kambhu & Mosser (2001) for the U.S. dollar swaps market. They find that no such relationship exists, except for one segment of the market. Typically, liquidity in the underlying market is much larger than hedging demand and changes in prices, that is, the yield curve, originate from other sources. However, their data is rather “coarse”. While they use daily data of interest rates, they keep open interest in the derivatives constant over the period they look at, which is many years long. We saw in Section 3 that open interest changes significantly throughout the year. Furthermore, in a considerable number of recent studies, reviewed in Section 2, it has been shown that many phenomena in financial markets can only be explained by the analysis of very “fine” data, that is, intra-day data in event time or intervals of only a few minutes.

We investigate the relationship between hedging demand and order volume in the underlying market for the Swiss stock market. As such, our analysis contributes to the growing literature on “microscopic” analyses of the dynamics of financial markets. Whereas this literature so far has focused on the impact of order flow, via order book, on price, little analysis has been done on the origins of order flow. Furthermore, we are not aware of any analysis on such a microscopic level that investigates the relationship between price changes and future order flow. While it has been known that price change does have an impact on future volume, the respective analyses are based on daily data. Furthermore, at least to our knowledge, little work has been done that investigates hedging as a mechanism for such a relationship, with the exception of Kambhu (1998) and Kambhu & Mosser (2001).

There are reasons why, in our opinion, no such analyses have been conducted. First, the literature on “microscopic” empirical analyses of financial markets

so far has focused on the description of single quantities like order volume, granularity of limit order prices, price impact, and price change. To date, little work has been done on the dynamics of these quantities and their interaction. A second likely reason, related to the first, is the (non-)availability of appropriate data. To conduct a proper analysis of hedging demand and order flow, one needs intra-day data of order flow in the underlying market and at least daily data on open interest in the options market. So far, such data has been difficult to obtain.

Our analysis is based on event-level data of order flow and transactions on the Swiss Stock Exchange, as well as daily data on open interest on Eurex, the (German-)Swiss derivatives exchange. Based on this data, we are able to compute hedging demand on an intra-day basis and compare it to order flow in the respective time periods. Our null hypothesis, partly influenced by the analyses of Kambhu (1998) and Kambhu & Mosser (2001), is that hedging demand is not systematically related to order volume in the underlying market. In other words, we hypothesize that future order volume cannot be predicted by current hedging demand (alone). On the basis of a broad range of tests, we do not need to reject our null hypothesis. This means that we do not find evidence for a systematic relationship between hedging demand and future order volume in the underlying market. Our results are robust with regards to a many changes in our model specification.

We do not find evidence for a systematic relationship between hedging demand and price impact, either. This means that at times when hedging demand represents a large part of order volume, price impact is not systematically higher. We consider our results to be important for at least two reasons. First, if there was a systematic relationship between hedging demand and order volume, then hedging activities would be likely to create positive feedback effects on price changes. This would have serious implications for options pricing and risk management policies. Secondly, the presence of such a systematic relationship would imply the existence of externalities of options trading on the underlying market. This would have serious implications for public policy as options trading, the domain a large institutional traders, would increase the costs of trading in the underlying market, where a much broader range of investors are active, in the form of liquidity costs.

Based on our analysis, we are lead to assume that no such systematic relationship exists. Our results casts serious doubt on the common argument and widely-held belief that options trading has a systematic negative impact—in the form of a positive feedback effect—on underlying markets.

References

- Acharya, V. V. & Pedersen, L. H. (2004), Asset pricing with liquidity risk. Working paper, Stern School of Business, New York University.
- Admati, A. & Pfleiderer, P. (1988), ‘A theory of intraday trading patterns: Volume and price variability’, *Review of Financial Studies* **1**(1), 3–40.
- Amihud, Y. & Mendelson, H. (1986), ‘Asset pricing and the bid-ask spread: An empirical investigation’, *Journal of Financial Economics* **17**(2), 223–249.
- Amihud, Y. & Mendelson, H. (1991), ‘Volatility, efficiency, and trading: Evidence from the Japanese stock market’, *Journal of Finance* **46**(5), 1765–1789.
- Amihud, Y., Mendelson, H. & Lauterbach, B. (1997), ‘Market microstructure and securities values: Evidence from the Tel Aviv Stock Exchange’, *Journal of Financial Economics* **45**(3), 365–390.
- Anthony, J. (1988), ‘The interrelation of stock and option market trading-volume data’, *Journal of Finance* **43**(4), 949–961.
- Bak, P., Paczuski, M. & Shubik, M. (1997), ‘Price variations in a stock market with many agents’, *Physica A* **246**(3–4), 430–453.
- Bauer, W. (2004), Commonality in liquidity in pure order-driven markets. Working paper, University of Zurich.
- Biais, B., Gloston, L. & Spatt, C. (2002), The microstructure of stock markets. Unpublished working paper, Université de Toulouse.
- Biais, B., Gloston, L. & Spatt, C. (2005), ‘Market microstructure: A survey of microfoundations, empirical results, and policy implications’, *Journal of Financial Markets* **8**(2), 217–264.
- Biais, B., Hillion, P. & Spatt, C. (1995), ‘An empirical analysis of the limit order book and the order flow in the Paris Bourse’, *Journal of Finance* **50**(5), 1655–1689.
- Black, F. & Scholes, M. (1973), ‘The pricing of options and corporate liabilities’, *Journal of Political Economy* **81**, 637–659.
- Bollerslev, T., Chou, R. Y. & Kroner, K. K. (1992), ‘ARCH modeling in finance: A review of the theory and empirical evidence’, *Journal of Econometrics* **51**(1–2), 5–59.
- Bouchaud, J.-P., Gefen, Y., Potters, M. & Wyart, M. (2004), ‘Fluctuations and response in financial markets: the subtle nature of ‘random’ price changes’, *Quantitative Finance* **4**(2), 176–190.
- Bouchaud, J.-P., Mézard, M. & Potters, M. (2002), ‘Statistical properties of stock order books: empirical results and models’, *Quantitative Finance* **2**(4), 251–256.

- Brennan, M. J., Chordia, T. & Subrahmanyam, A. (1998), ‘Alternative factor specifications, security characteristics, and the cross-section of expected stock returns’, *Journal of Financial Economics* **49**(3), 345–373.
- Brennan, M. J. & Subrahmanyam, A. (1996), ‘Market microstructure and asset pricing: On the compensation for illiquidity in stock returns’, *Journal of Financial Economics* **41**(3), 441–464.
- Bruart, M. & Gibson, R. (1998), ‘The effects of newly listed derivatives in a thin stock market’, *Review of Derivatives Research* **2**(1), 59–86.
- Campbell, J. Y., Lo, A. H. & MacKinlay, A. C. (1997), *The Econometrics of Financial Markets*, Princeton University Press.
- Challet, D. & Stinchcombe, R. (2003), ‘Limit order market analysis and modelling: on a universal cause for over-diffusive prices’, *Physica A* **324**(1–2), 141–145.
- Cohen, K. J., Conroy, R. M. & Maier, S. F. (1985), Order flow and the quality of the market, in Y. Amihud, T. Lo & R. Schwartz, eds, ‘Market Making and the Changing Structure of the Securities Industry’, Lexington Books.
- Cohen, K. J., Maier, S. F., Schwartz, R. A. & Whitcomb, D. K. (1981), ‘Transaction costs, order placement strategy, and existence of the bid-ask spread’, *Journal of Political Economy* **89**(2), 287–305.
- Cont, R. (1998), ‘A Langevin approach to stock market fluctuations and crashes’, *European Physics Journal B* **6**(4), 543–550.
- Cont, R. (2001), ‘Empirical properties of asset returns: stylized facts and statistical issues’, *Quantitative Finance* **1**(2), 223–236.
- Cont, R. & Bouchaud, J.-P. (2000), ‘Herd behavior and aggregate fluctuations in financial markets’, *Macroeconomic Dynamics* **4**(2), 170–196.
- Cont, R., Potters, M. & Bouchaud, J.-P. (1997), Scaling in stock market data: Stable laws and beyond. eprint cond-mat/9705087.
- Coppejans, M., Domowitz, I. & Madhavan, A. (2000), Liquidity in an automated auction. Working paper, Duke University.
- Cutler, D., Poterba, J. M. & Summers, L. H. (1989), ‘What moves stock prices?’, *Journal of Portfolio Management* **15**(3), 4–12.
- Damodaran, A. & Lim, J. (1991), ‘The effects of option listing on the underlying stocks’ return process’, *Journal of Banking and Finance* **15**(3), 647–664.
- Daniels, M. G., Farmer, J. D., Iori, G. & Smith, E. (2003), ‘Quantitative model of price diffusion and market friction based on trading as a mechanistic random process’, *Physics Review Letters* **90**(10), 108102/1–4.
- Dobson, A. J. (2002), *An Introduction to Generalized Linear Models*, Chapman Hall.
- Domowitz, I. & Wang, J. (1994), ‘Auctions as algorithms: Computerized trade execution and price discovery’, *Journal of Economic Dynamics and Control* **18**(1), 29–60.

- Easley, D. & O'Hara, M. (1987), 'Price, trade size, and information in security price adjustment', *Journal of Financial Economics* **19**(1), 69–90.
- Easley, D., O'Hara, M. & Srinivas, P. S. (1998), 'Option volume and stock prices: Evidence on where informed traders trade', *Journal of Finance* **53**(2), 431–465.
- Easley, D. & O'Hara, M. (1992), 'Time and the process of security price adjustment', *Journal of Finance* **47**(2), 576–605.
- Eurex (2003), 'Risk based margining'. Eurex Frankfurt.
- Eurex (2005), 'Eurex-products'. Eurex Frankfurt.
- Evans, M. D. D. & Lyons, R. K. (2002), 'Order flow and exchange rate dynamics', *Journal of Political Economy* **110**(1), 170–180.
- Fama, E. F. (1965), 'The behavior of stock-market prices', *Journal of Business* **38**(1), 34–105.
- Farmer, J. D., Gillemot, L., Lillo, F., Mike, S. & Sen, A. (2004), 'What really causes large price changes?', *Quantitative Finance* **4**(4), 383–397.
- Farmer, J. D. & Lillo, F. (2004), 'On the origin of power-law tails in price fluctuations', *Quantitative Finance* **4**, C7–C11.
- Farmer, J. D. & Zovko, I. (2002), 'The power of patience: A behavioural regularity in limit-order placement', *Quantitative Finance* **2**(5), 387–392.
- Friedman, D. (1991), The double auction institution: A survey, *in* D. Friedman & J. Rust, eds, 'The Double Auction Market: Institutions, Theory, and Evidence', SFI Studies in the Sciences of Complexity, Addison-Wesley, Reading, pp. 3–25.
- Gabaix, X., Gopikrishnan, P., Plerou, V. & Stanley, H. E. (2003), 'A theory of power-law distributions in financial market fluctuations', *Nature* **423**, 267–270.
- Genotte, G. & Leland, H. (1990), 'Market liquidity, hedging, and crashes', *American Economic Review* **80**(5), 999–1021.
- Glosten, L. & Milgrom, P. (1985), 'Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders', *Journal of Financial Economics* **14**(1), 71–100.
- Goldstein, M. A. & Kavajecz, K. A. (2000), 'Eighths, sixteenths, and market depth: changes in tick size and liquidity provision on the nyse', *Journal of Financial Economics* **56**, 125–149.
- Gopikrishnan, P., Plerou, V., Gabaix, X., Amaral, L. A. N. & Stanley, H. E. (2001), 'Price fluctuations and market activity', *Physica A* **299**, 137–143.
- Gopikrishnan, P., Plerou, V., Gabaix, X. & Stanley, H. E. (2000), 'Statistical properties of share volume traded in financial markets', *Physical Review E* **62**(4), R4493.
- Gourieroux, C. & Jasiak, J. (2001), *Financial Econometrics: Problems, Models, and Methods*, Princeton University Press.

- Grammig, J., Heinen, A. & Rengifo, E. (2004), Trading activity and liquidity supply in a pure limit order book market: An empirical analysis using a multivariate count data model. Working paper, Catholic University of Louvain.
- Griffiths, M. D., Smith, B. F., Turnbull, D. A. S. & White, R. W. (2000), 'The costs and determinants of order aggressiveness', *Journal of Financial Economics* **56**, 65–88.
- Grossman, S. J. (1988), 'An analysis of the implications for stock and futures price volatility of program trading and dynamic hedging strategies', *Journal of Business* **61**(3), 275–298.
- Harris, L. E. (2003), *Trading and Exchanges*, Oxford University Press.
- Harris, L. & Hasbrouck, J. (1996), 'Market vs. limit orders: The SuperDOT evidence on order submission', *Journal of Financial and Quantitative Analysis* **31**, 213–231.
- Hasbrouck, J. (1991), 'Measuring the information content of stock prices', *Journal of Finance* **46**(1), 179–207.
- Hasbrouck, J. & Seppi, D. (2001), 'Common factors in prices, order flows and liquidity', *Journal of Financial Economics* **59**(3), 383–411.
- Hausman, J. A., Lo, A. W. & MacKinlay, A. C. (1992), 'An ordered probit analysis of transaction stock prices', *Journal of Financial Economics* **31**(3), 319–379.
- Hollifield, B., Miller, R. A. & Sandås, P. (2004), 'Empirical analysis of limit order markets', *Review of Financial Studies* **71**, 1027–1063.
- Hull, J. C. (2002), *Options, Futures, and Other Derivatives*, 5 edn, Prentice Hall.
- Kambhu, J. E. (1998), 'Dealer's hedging of interest rate options in the U.S. dollar fixed-income market', Federal Reserve Bank of New York Economic Policy Review.
- Kambhu, J. E. & Mosser, P. C. (2001), 'The effect of interest rate options hedging on term-structure dynamics', Federal Reserve Bank of New York Economic Policy Review.
- Keim, D. & Madhavan, A. (1997), 'Transaction costs and investment style: An inter-exchange analysis of institutional equity trades', *Journal of Financial Economics* **46**, 265–292.
- Kempf, A. & Korn, O. (1999), 'Market depth and order size', *Journal of Financial Markets* **2**, 29–48.
- Knez, P. J. & Ready, M. J. (1996), 'Estimating the profits from trading strategies', *Review of Financial Studies* **9**(4), 1121–1163.
- Lehmann, B. M. & Modest, D. M. (1994), 'Market structure and liquidity on the Tokyo Stock Exchange', *Journal of Finance* **49**(3), 951–984.

- Lillo, F. & Farmer, J. D. (2004), 'The long memory of the efficient market', *Studies in Nonlinear Dynamics and Econometrics* **8**(3). <http://www.bepress.com/snnde/vol8/iss3/art1>.
- Lillo, F., Farmer, J. D. & Mantegna, R. N. (2002), Single curve collapse of the price impact function of the New York Stock Exchange. Working Paper, Santa Fe Institute.
- Lux, T. (1998), 'The socio-economic dynamics of speculative markets: Interacting agents, chaos, and the fat tails of return distributions', *Journal of Economic Behavior and Organization* **33**(2), 143–165.
- Madhavan, A. (2000), 'Market microstructure: A survey', *Journal of Financial Markets* **3**, 205–258.
- Mandelbrot, B. (1963), 'The variation of certain speculative prices', *Journal of Business* **36**(4), 394–419.
- Mandelbrot, B. (1966), 'Forecasts of future prices, unbiased markets, and "martingale" models', *Journal of Business* **39**(1), 242–255.
- Maslov, S. (2000), 'Simple model of a limit order-driven market', *Physica A* **278**(3–4), 571–578.
- Mendelson, H. (1982), 'Market behavior in a clearing house', *Econometrica* **50**(6), 1505–1524.
- Merton, R. (1973), 'The theory of rational option pricing', *Bell Journal of Economics and Management Science* **4**, 141–183.
- Muzy, J. F., Bacry, E. & Delour, J. (2001), 'Modeling financial time series using multifractal random walks', *Physica A* **299**, 84–92.
- Niederhoffer, V. & Osborne, M. F. M. (1966), 'Market making and reversal on the stock exchange', *Journal of the American Statistical Association* **61**(316), 897–916.
- Officer, R. R. (1972), 'The distribution of stock returns', *Journal of the American Statistical Association* **67**(340), 807–812.
- O'Hara, M. (1999), *Market microstructure theory*, Blackwell Publishers.
- Pagan, A. R. (1996), 'The econometrics of financial markets', *Journal of Empirical Finance* **3**, 15–102.
- Parlour, C. A. (1998), 'Price dynamics in limit order markets', *Review of Financial Studies* **11**(4), 789–816.
- Patie, P. & Frey, R. (2002), Risk management for derivatives in illiquid markets: A simulation study, in K. Sandmann & P. Schönbucher, eds, 'Advances in Finance and Stochastics', Springer, pp. 137–159.
- Petersen, M. A. & Umlauf, S. (1994), Quote revisions and the quoted size: An empirical analysis. Working paper, Kellogg School of Management, Northwestern University.

- Pictet, O. V., Dacorogna, M. M. & Müller, U. A. (1996), Heavy tails in high-frequency financial data. Working paper, Olsen and Associates.
- Plerou, V., Gopikrishnan, P., Amaral, L. A. N., Gabaix, X. & Stanley, H. E. (2000), ‘Economic fluctuations and anomalous diffusion’, *Physical Review E* **62**(3), R3023.
- Plerou, V., Gopikrishnan, P., Gabaix, X. & Stanley, H. E. (2002), ‘Quantifying stock-price response to demand fluctuations’, *Physical Review E* **66**, 027104.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L. A. N., Guhr, T. & Stanley, H. E. (2002), ‘Random matrix approach to cross correlations in financial data’, *Physical Review E* **65**, 066126.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L. A. N. & Stanley, H. E. (1999), ‘Universal and nonuniversal properties of cross correlations in financial time series’, *Physical Review Letters* **83**(7), 1471–1474.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L. A. N. & Stanley, H. E. (2000), ‘Econophysics: financial time series from a physics point of view’, *Physica A* **279**, 443–456.
- Potters, M. & Bouchaud, J.-P. (2003), ‘More statistical properties of order books and price impact’, *Physica A* **324**, 133–140.
- Pritsker, M. (1997), Liquidity risk and positive feedback. Working paper, Federal Reserve Board, Washington.
- Rinaldo, A. (2003), ‘Order aggressiveness in limit order book markets’, *Journal of Financial Markets* **7**(1), 53–74.
- Rosenow, B. (2002), ‘Fluctuations and market friction in financial trading’, *International Journal of Modern Physics* **13**(3), 419–425.
- Schlag, C. & Stoll, H. (2005), ‘Price impacts of options volume’, *Journal of Financial Markets* **8**, 69–87.
- Shiller, R. J. (1989), Comovements in stock prices and comovements in dividends. Working paper, National Bureau of Economic Research.
- Slanina, F. (2001), ‘Mean-field approximation for a limit order driven market model’, *Physical Review E* **64**, 056136.
- Smith, E., Farmer, J. D., Gillemot, L. & Krishnamurthy, S. (2003), Statistical theory of the continuous double auction. Working Paper, Santa Fe Institute, Santa Fe.
- Sornette, D. (2003), *Why Stock Market Crash*, Princeton University Press, Princeton, NJ.
- Subrahmanyam, A. (1991), ‘Risk aversion, market liquidity, and price efficiency’, *Review of Financial Studies* **4**(3), 417–441.
- Weber, P. & Rosenow, B. (2003), Order book approach to price impact. eprint cond-mat/0311457.

Weber, P. & Rosenow, B. (2004), Large stock price changes: volume or liquidity?
eprint cond-mat/0401132.