Measuring market impact and liquidity

May 17, 2006

Galen Burghardt
galen.burghardt@newedgegroup.com

Jerry Hanweck
jerry.hanweck@newedgegroup.com

Lauren Lei
lauren.lei@newedgegroup.com

Poor execution can turn a good idea into a losing trade. Skilled traders (and their brokers) use their knowledge and experience to balance the immediacy of a transaction with the liquidity available in the market. This note reports on the work we have done to estimate the market impact of one-off trades in electronically traded futures markets.

We have invested heavily in gathering a continuous time market depth data base that allows us to observe the limit order book in nearly continuous time and to track the flow of actual trades. In the case of the limit order book, we are tracking the best five bids and best five offers. This data set is exceptionally valuable for studying market liquidity and the impact of trades of various sizes. For example, it allows us to calculate a sweep to fill measure of market impact – the effect on the price of instantaneously trading as far into the book as necessary to fill an order of a given size. That is, the cost of sweeping the book to fill an order.

Exhibit 1 shows what the resulting sweep to fill market impact profiles looked like for the E-mini S&P500 futures market during the first quarter of 2006. Information like this can be very useful to a trader. For one thing, the trader can identify the most liquid times of the trading day. For another, the trader can determine the effect of speeding up or slowing down on total transactions costs and make informed choices between tracking error and the costs of trading.

Exhibit 1
Sweep to fill cost
(E-minis S&Ps, 1/3/06 to 3/31/06)

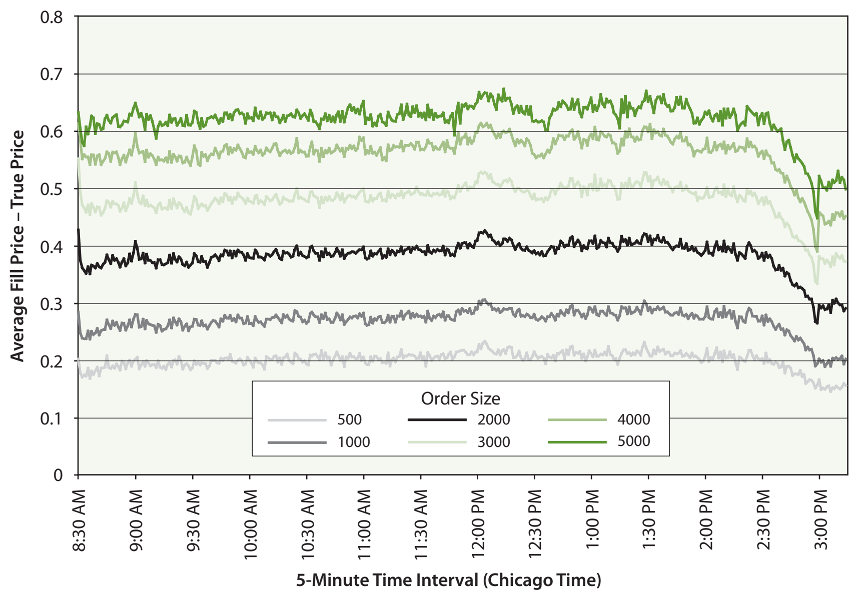

In fact, these data sets allow us to delve much more deeply into the question of market impact and liquidity and to formulate measures of market impact that take advantage of knowledge about predictable patterns of trading volume and price volatility, the risk aversion of market makers, and hidden liquidity.

AlternativeEdge series

The Alternative Edge series of research notes is a collaboration between Galen Burghardt and the Alternative Investments Group (AIG), which is led by Leslie Richman and Brian Walls. The AIG is committed to open and informed dialogue between investors and money managers. To that end, we endeavor to provide our clients with probing and objective research into questions that matter to us all. For a compilation of our research notes and information on the Alternative Edge CTA Indices, please go to newedgegroup.com.

Outline

The purpose of this note is to show the steps we have taken in using these data sets to produce useable market impact profiles for use in designing cost efficient trading programs. In the process, we

- Describe the data set we have assembled
- Show some of the things that can be learned from the limit order book
- Define sweep to fill market impact
- Describe the risks faced by a market maker
- Show how a simple risk aversion model can capture the shape of the limit order book
- Estimate the effects of hidden liquidity
- Show our model estimates for four key futures markets
- Show how we can allow for significant economic announcements

We also provide a technical appendix that shows the derivation of the risk averse market maker model that we use to describe the limit order book.

A very fat data set

To do this kind of work, we have assembled a very rich dataset for electronically traded futures that comprises price and sizes for all trades as well as the best five bids and best five asks. In assembling these data, we have recorded every instance of a change in the limit order book as well as every trade. This way, we can produce a snapshot of what the limit order book looked like at any instant. We can track additions to and withdrawals from the book, and we can infer from the status of the book the moment before each trade whether the trade occurred at the bid or the ask. From the trades themselves, we can learn about patterns of trading volume, both frequency and size, and of price volatility. Appendix A provides a list of the markets for which we are gathering these data.

These data lend themselves to insights into liquidity in several ways. For one thing, at the simplest level, they can tell you how much it would cost to fill an order of any given size immediately. Exhibit 2, for example, shows what the E-mini S&P500 book looked like at 8:40:00am on February 21st, 2006. It is worth noting that there would have been several limit order books published during the single second from 8:40:00 to 8:40:01, and that this was the first of those books.

Sweep to fill

In this case, the best bid was 1292.50 with 361 contracts and the best ask was 1292.75 with 717 contracts. Thus, you could immediately sell 361 contracts at a price of 1292.50, or you could immediately buy 717 contracts at the price of 1292.75. If you wanted to sell more than 361 contracts or buy more than 717 contracts, and to do either immediately, you would have to sell at the next best bid of 1292.25 or to buy at the next best ask of 1293.00. Since there were 509 contracts at the second best bid, you could sell as many as 870 without going deeper into the book, and if you were to sell 361 at the best bid and 509 at the next best bid, your average trade price would be 1292.35 [= (361/870)x1292.50 + (509/870)x1202.25]. If you were to buy 2,089 contracts immediately, your average price would be 1292.91 [= (717/2,089)x1292.75 + (1,372/2,089)x1293.00].

Exhibit 2
Limit order book for E-mini S&P500 futures
2/21/2006 (8:40:00am)

Price	Bid			Ask		
	average price*	total	contracts	contracts	total	average price*
1293.75				1,361	5,973	1293.29
1293.50				1,089	4,612	1293.16
1293.25				1,434	3,523	1293.05
1293.00				1,372	2,089	1292.91
1292.75				717	717	1292.75
1292.50	1292.50	361	361			
1292.25	1292.35	870	509			
1292.00	1292.15	2,105	1,235			
1291.75	1292.06	2,686	581			
1291.50	1291.92	3,578	892			

* Average sweep to fill price

A practice like this is known as sweeping the book, and the resulting average price one realizes is known as a "sweep to fill" price. The immediate market impact of sweeping the book can be calculated by taking the difference between the sweep to fill price and some measure of the mid-market price. The market impact profiles shown in Exhibit 1, for example, provide an example of how this measure of market impact looked for trades of various sizes in E-mini S&P500 futures for the first quarter of 2006 on those days when there were no scheduled economic announcements of particular importance. For this illustration, the measure of mid-market we have used is true market price, which we will explain in the next section.

Profiles like these are useful guides to finding market liquidity. Notice, for example, that the market impact of a trade is higher right at the market open of 8:30am (Chicago time) than it would be even a few minutes into the trading day. You can see that market impact tends to rise as the trading day progresses and tends to be highest right around noon, when trading is slowest. It is also apparent that the market is most liquid (impact is smallest) as the market reaches the cash close time of 3:00pm and remains relatively liquid until the futures close at 3:15pm.

True market price

In the academic literature, one finds the very sensible notion that there is a true market price that lies between the best bid and the best ask. Evidence gleaned from the limit order book tends to bear this out, and those who trade or execute trades for a living will recognize the idea immediately.

Consider again the order book in Exhibit 2. There are 717 contracts offered as the best ask price of 1292.75, but only 361 contracts at the best bid price of 1292.50. From this information, traders and their brokers are likely to observe that the market is more heavily offered than bid and that the next trade is more likely to take place at the bid than at the offer.

To lend concreteness to this idea, we define order imbalance as ln(Bid size/Ask size). If the numbers of contracts at the bid and ask are equal, the ratio would be one, and this measure of order imbalance would be 0. This measure of order imbalance will be negative if the number of contracts bid is less than the number offered, and will be positive if the number of contracts bid is more than the number offered. Moreover, it will be symmetrical, so that if the number of contracts bid is twice the number offered, the measure would be $\ln(2) = 0.693$. If the number of contracts bid is half the number offered, the measure would be $\ln(0.5) = -0.693$.

Exhibit 3
Time to Next Trade Versus Order Imbalance
(E-mini S&P Futures, 2/4/06 to 2/23/06)

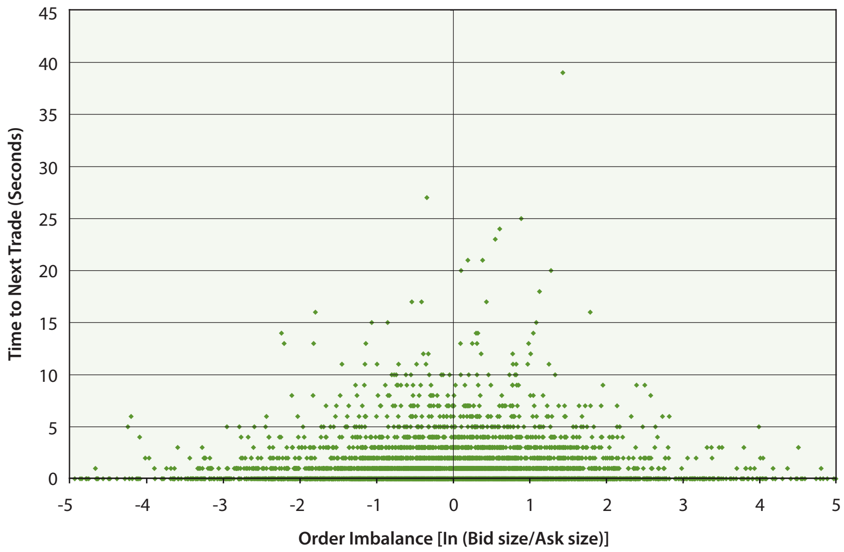

Exhibit 4
Next Trade Price Versus Order Imbalance
(E-mini S&P Futures, 2/4/06 to 2/23/06)

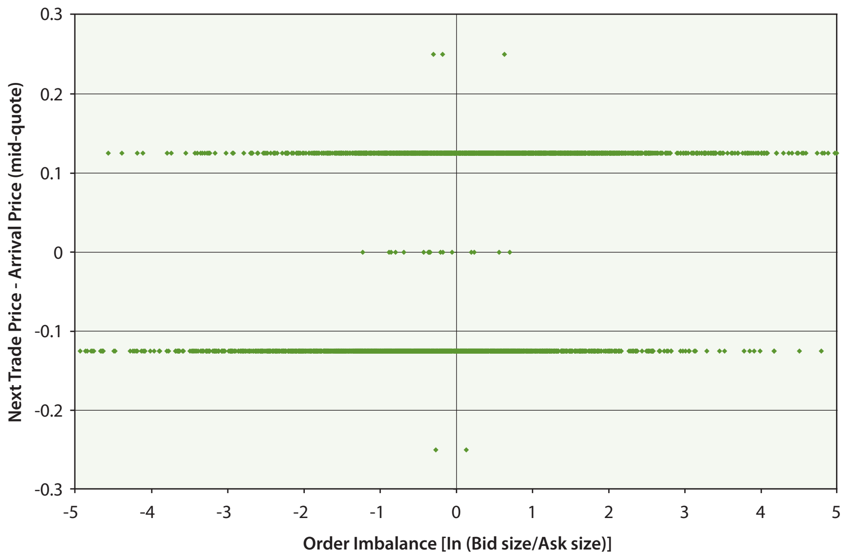

Now consider Exhibit 3, which shows the relationship between order imbalance and the time to next trade for E-mini S&Ps. When order imbalance is close to zero – that is, when the number of contracts at the bid and ask are roughly equal – the time to next trade tends to be high. In contrast, when the order imbalance is large in either direction, the time between trades tends to be small. This seems reasonable enough if a balanced order book means that the price at which people would really like to trade is equally far from both the bid price and the ask price.

Also consider Exhibit 4, which shows the relationship between the next trade price and order imbalance. On the vertical axis is the difference between the next trade price and mid-market (halfway between the best bid and ask). Most of the dots in this exhibit occur at a difference of +0.125 or -0.125, which in E-mini S&Ps represents half a tick. These cases represent instances of the price trading either at the ask or at the bid when the market was 1 tick wide. There are a few cases where the price change was a full tick (either +0.25 or -0.25) or zero. These outcomes correspond to times when the market was 2 ticks wide.

What we find in this exhibit is that trades are more likely to take place at the ask when the order imbalance is positive and more likely to take place at the bid when the order imbalance is negative. Again, just what experienced traders or brokers would expect.

Exhibit 5
Average Price Change Versus Bid/Ask Imbalance
(E-mini S&P Futures, 2/4/06 to 2/24/06)

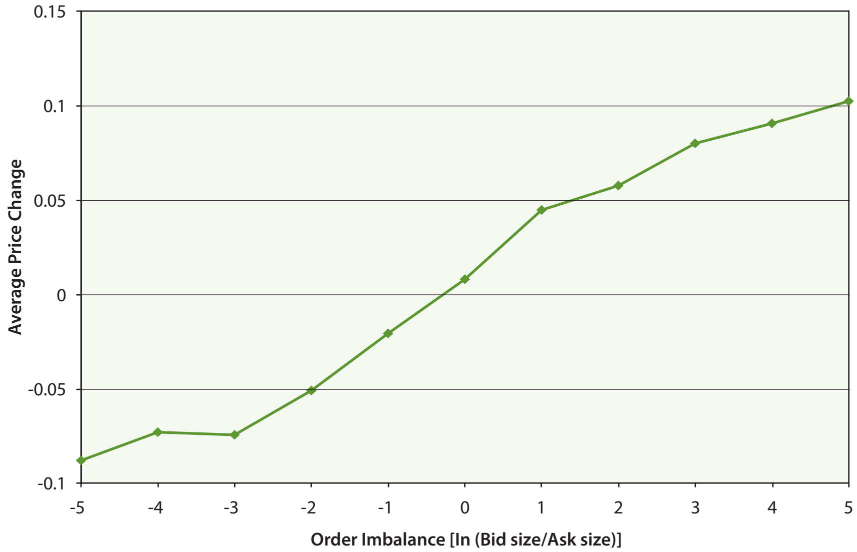

A slightly different representation of this feature of the market is shown in Exhibit 5, which shows the average price change for various order imbalance values. The effect of averaging the price changes produces a fairly well shaped curve that confirms the idea that the "true market price" can be tied reliably to imbalances in the limit order book.

Exhibit 6 provides a sample of how the idea of a true market price might be used in practice. The exhibit shows a bid/ask channel for E-mini S&Ps for the 10 seconds from 8:30:00am to 8:30:10am. It also shows a running calculation of the true market price, which was calculated as:

$$Trueprice = \left(\frac{Q_A}{Q_A + Q_B}\right)P_B + \left(\frac{Q_B}{Q_A + Q_B}\right)P_A$$

Exhibit 6
Bid, Ask, and True Market Prices
(E-mini S&P 500 Futures, 4/12/06 8:30:00 to 8:30:10am)

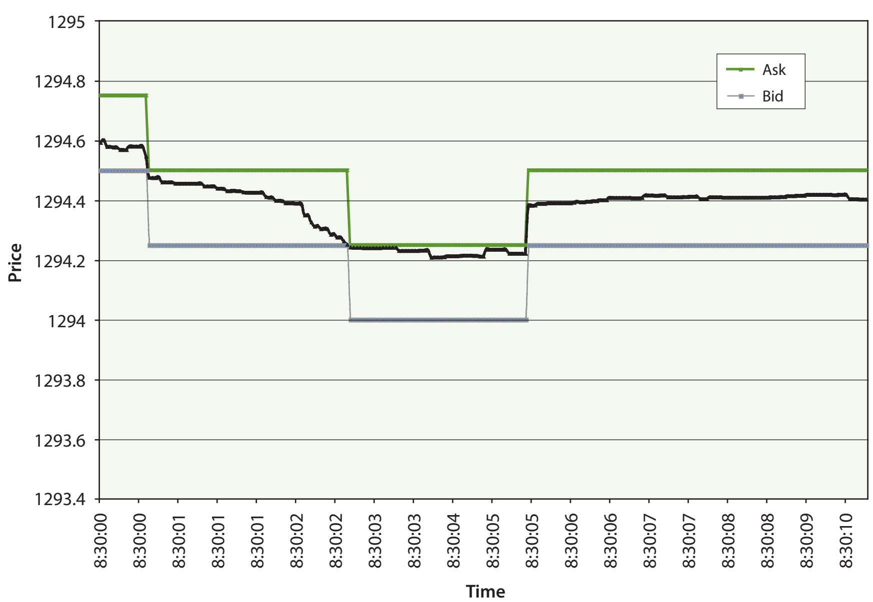

To be sure, the true price has to lie between the bid and ask prices, but there does seem to be some useful information about where the market is going based on the relationship between the true price and the surrounding bid and ask prices.

A representative market maker

For the purposes of bringing some order to the data, we will treat the great pool of traders whose business it is to provide liquidity as a single, representative, risk averse market maker. This fairly simple approach is surprisingly effective in explaining the shape of the limit order book.

Exhibit 7
Market Maker's Time Line

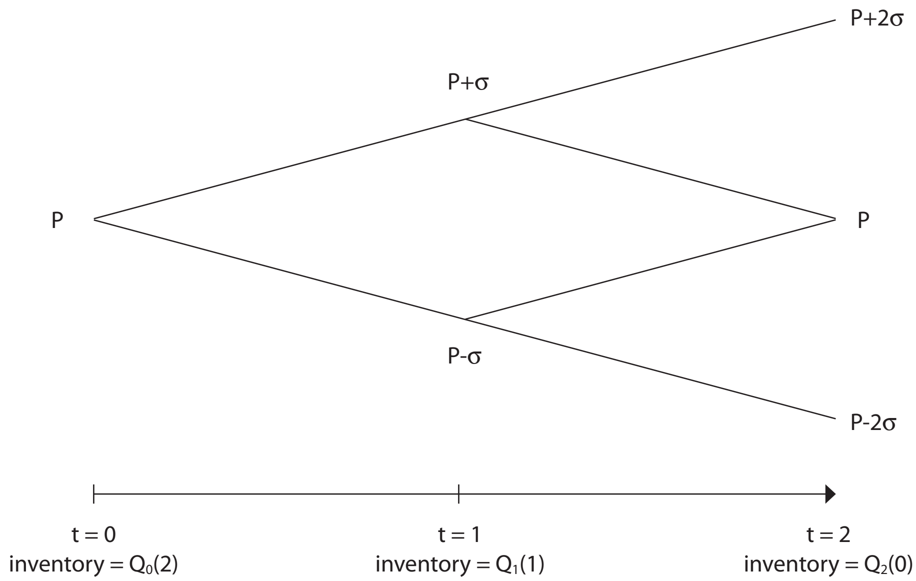

The market maker's risk

To illustrate the market maker's risk, we use a simple 2-period example that allows us to consider both the speed with which the market maker's inventory can be worked off and the price risk that comes with holding the inventory. For this example, we assume that the market maker is asked to take on an order size of 2 contracts, that the volume of trading in the market will allow the market maker to get rid of 1 contract per period, and that during each period the price can either rise or fall by an amount equal to σ. The initial price is P_0.

In this setting, as shown in Exhibit 7, the market maker's initial inventory at t_0 is 2 contracts. At t_1, he can get rid of one contract at P_1. He now has a remaining inventory of 1 contract, which he holds for one more period. Then, at t_2, he gets rid of this contract at P_2. Given this sequence of events, we can write the market maker's P/L as

$$P/L = Q_0 dP + Q_1 dP$$

If the price changes are uncorrelated from one period to the next, the variance of this P/L can be written as

$$V(P/L) = Q_0^2 dP^2 + Q_1^2 dP^2$$

and the standard deviation as

$$SD(P/L) = \sqrt{Q_0^2 dP^2 + Q_1^2 dP^2}$$

Using a price change of plus or minus σ, the market maker's total risk in our example would be

$$Totalrisk = \sqrt{2^2 \sigma^2 + 1^2 \sigma^2}$$

which can be divided by the initial order size to find risk per contract.

The market maker's risk aversion

To explain the shape of the limit order book, we can assume that the market maker is risk averse. Further, as we will show, we can do a good job of explaining the shape of actual limit order books by assuming that the market maker requires an amount α for each unit of risk taken on.

With this assumption, and using a continuous time expression for the market maker's risk, we can express the market maker's required compensation for taking on an order size of Q as

$$TotalComp(Q) = \frac{\alpha\sigma}{\sqrt{3cv}} Q^{3/2}$$

where α is the market maker's required compensation per unit of risk, σ is the standard deviation of arithmetic changes in the price, v is the rate of volume traded in the market, c is the fraction of that volume available to the market maker, and Q is the total order size. Average market maker compensation per contract would simply be

$$AverageComp(Q) = \alpha \frac{\sigma}{\sqrt{3cv}} Q^{1/2}$$

The intuition behind this expression is fairly straightforward. For one thing, an increase in risk aversion translates directly into an increase in the spread. For another, increases in price risk translate directly into an increase in what the market maker requires to take on a position. And third, while time does not seem to enter directly, it does indirectly. A quadrupling of the order size, for example, would quadruple the amount of time the market maker requires to unwind the position at whatever speed the flow of trading in the market allows. The effect on the spread, however, is proportional to the square root of the order size. Thus, a quadrupling of the order size only doubles the size of the spread. Trading volume, the square root of which appears in the denominator, works in the same way as order size but in the opposite direction. A quadrupling of useable trading volume in the market would cut the spread in half.

As a practical matter, compensation per contract in this exercise would be one half the bid/ask spread for orders of any given size since the market maker has to be willing to either buy or sell.

Exhibit 8
Sweep to Fill Costs Versus Order Size
EMINISP, 8:40am Alpha=1.94
(1/3/06 to 3/31/06)

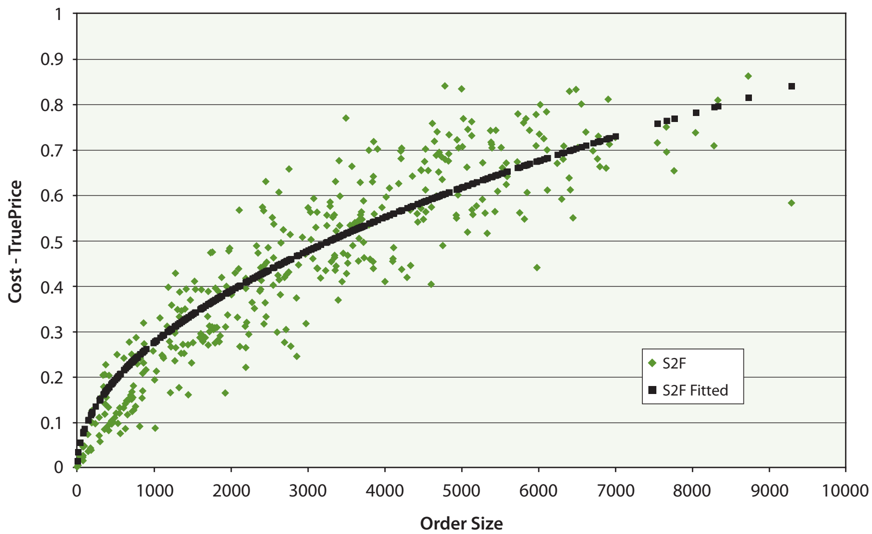

Exhibit 9
Calibrating the Model to Sweep to Fill Costs at 12:40pm
EMINISP, 12:40pm Alpha=1.85
(1/3/06 to 3/31/06)

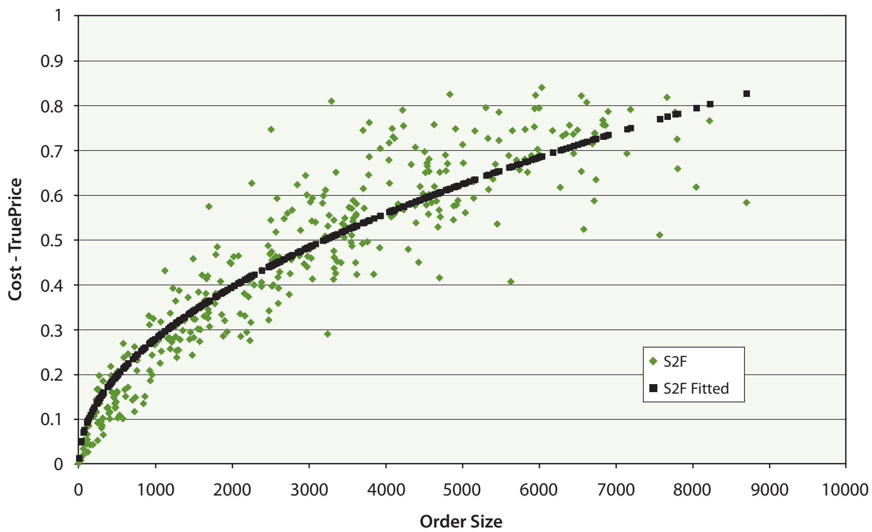

Fitting the curve to the data

The usefulness of this way of thinking about market making is illustrated in Exhibits 8 and 9, which show actual and fitted values of sweep to fill impact values for two different times of day in E-mini S&P500 trading. Exhibit 8 corresponds to 8:40am, when the market is very active. Exhibit 9 corresponds to 12:40pm, when the market is comparatively quiet.

Using specific times of day for curve fitting is a way of controlling for trading volume and price volatility. As we will see, volume and volatility profiles exhibit a lot of regularity, so that 8:40 in the morning on one trading day will look a lot like 8:40 in the morning on another trading day.

When fitting the curves, we sought values of alpha that minimized the sum of squared errors in total trading cost. The difference between observed and fitted market impact is worth more when you are trading 5,000 contracts than when you are trading 500 contracts. So it makes sense to find estimates of the risk aversion parameter that do the best job of explaining total market impact rather than market impact per contract.

Hidden liquidity

What you see with a limit order book is not necessarily what you get. First, it shows phantom liquidity – bids and offers to which traders are not really committed and that are withdrawn either for no apparent reason or because the market begins to move in their direction. Whether these are available for sweep to fill orders is to some extent a matter of timing and fast action. Also, the limit order book does not reveal hidden liquidity – all of those potential bids and offers controlled by traders who don't want to show their hands.

Of the two, it seems that hidden liquidity is the more important consideration when analyzing market impact. As have others, we find that the apparent impact of trades tends to be smaller than sweep to fill measures of market impact would suggest.

Seconds per half tick

To get a handle on hidden liquidity, we conducted the following exercise. First, we defined spans of time that could reasonably be considered instantaneous for the purposes of measuring the flow of traffic through a market. While it certainly stretches the concept of instantaneous, it does make sense to define a span of time over which one might reasonably expect the market not to tick up or down. This kind of span depends on two things – tick size and price volatility.

For our purposes, we define this as the interval of time for which a one standard deviation price move is equal to half of a tick in each market. If price changes are normally distributed, this means that roughly 68 percent of the price changes over the interval would be within plus or minus one half of a tick.

Using realized price volatility for the six months leading up to April 11th, 2006 we translated annualized price volatility to volatility per second based on the number of trading seconds in a day (using normal open and close times) and assuming 252 trading days per year. Given volatility per second, one standard deviation over t seconds would be

$$\sigma_t = \sigma_s t^{1/2}$$

Setting this equal to one half a tick and solving for t, we get

$$t = \frac{1}{4}\left(\frac{tick}{\sigma_s}\right)^2$$

The resulting calculations for our four markets are shown in Exhibit 10.

Exhibit 10
How fast is instantaneous?

Market	Seconds per half tick
E-mini S&Ps	6
10-year Treasury notes	21
EuroStoxx	10
EuroBunds	10

Estimates of hidden liquidity

Our first step was to determine how many contracts actually traded during intervals of these lengths (5 seconds in the case of E-mini S&Ps) during the first quarter of 2006. In practice, we started at the beginning of each trading minute during regular trading hours. For each snapshot, we classified all trades that took place at a price above the true market price at the beginning of the interval as "buys" and all trades that took place at a price below the true market price as "sells." This was an imperfect distinction, but plausible. Suppose, for example, that the true market price is 1380 at 8:40am and that we observe the following E-mini S&P trades take place during the next 5 seconds.

Quantity	Price
20	1380.50
9	1379.50
17	1379.75
6	1380.25
25	1379.75

Exhibit 11
Vwap Market Impacts
(E-mini S&Ps, 1/3/06 to 3/31/06)

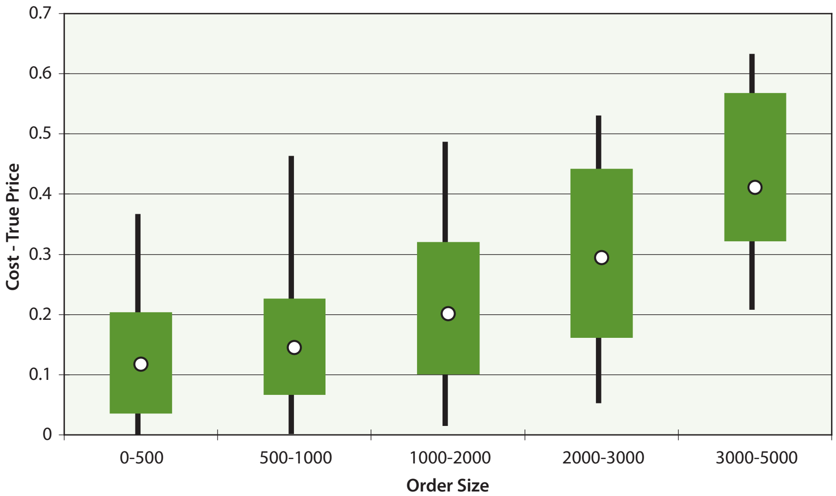

Exhibit 12
Sweep to Fill Market Impacts
(E-mini S&Ps, 1/3/06 to 3/31/06)

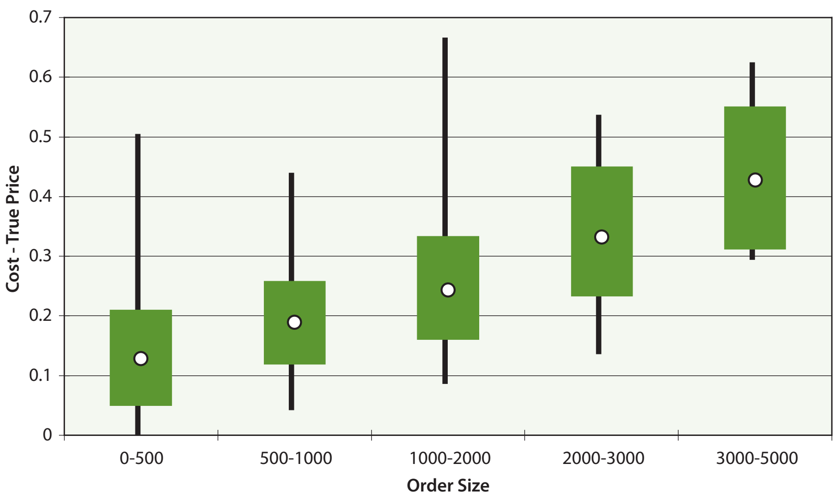

We would treat the first and fourth trades of 20 and 6 contracts respectively as "buys" and the second, third, and fifth trades or 9, 17, and 25 as "sells." The total number bought would be 26 and the total number sold would be 51.

We would then calculate volume weighted average prices – vwaps – for the "buys" and for the "sells." In this example, the average buy vwap is 1380.44 [= (20 x 1380.50 + 6 x 1380.25) / 26], while the average sell vwap is 1379.71 [= (9 x 1379.50 + 17 x 1379.75 + 25 x 1379.75) / 51].

Given the true market price of 100, the vwap impact of buying 26 contracts would be 0.44 [= 1380.44 – 1380.00], while the vwap impact of selling 51 contracts would be 0.29 [= Absolute value (1379.71 – 1380)].

The resulting distributions of these vwap impact estimates for E-mini S&Ps are shown in Exhibit 11. For the sake of simplicity, we have grouped the trades by lot sizes with the 0 to 500 bucket containing the smallest trades and 3,000 to 5,000 containing the largest trades. The dot in each figure represents the average value of the impact. The box represents 80 percent of the observations, from 10 percent to 90

Exhibit 13
Hidden liquidity summary for E-mini S&Ps
(1/3/06 - 3/31/06)

Order size	Average impact		
	Sweep to fill	Vwap	Difference
0-500	0.129	0.117	0.011
500-1000	0.189	0.145	0.044
1000-2000	0.243	0.202	0.042
2000-3000	0.332	0.295	0.038
3000-5000	0.428	0.411	0.017

percent. The top and bottom of each "whisker" represents the maximum and minimum values of each distribution.

The second step was to determine what the sweep-to-fill impacts would have been for trades whose sizes were equal to the total buys and total sells in each interval. In the above example, we would have calculated the sweep to fill impact of buying 26 contracts given the limit order book observed at exactly 8:40:00. Similarly, we would have calculated the sweep to fill impact of selling 51 contracts given the same limit order book. The distributions of these sweep to fill impacts for various order size buckets are shown in Exhibit 12.

Comparing the two distributions, we find that actual trade prices reveal more liquidity than is apparent in the limit order book. As shown in Exhibit 13, the effect of hidden liquidity was worth slightly more than a cent for small orders and just under 2 cents for fairly large trade sizes. For intermediate sized trades, though, the presence of hidden liquidity was worth considerably more. For trades between 500 and 1,000 contracts, hidden liquidity was worth 4.4 cents, while for trades between 2,000 and 3,000 contracts, hidden liquidity was worth about 3.8 cents per contract.

Calibrating for hidden liquidity

The next to last step in this process required us to fit the market maker's required compensation curve to the two sets of data described here. To do this, we used sweep to fill and vwap impact values irrespective of the time of day for which they were calibrated. As a result, the estimates of the risk aversion parameter are of no direct use, but we use the relative values of the estimated alphas to recalibrate the curves fit for specific times of day.

Risk aversion parameter (alpha) estimates

Exhibit 14 provides a summary of our estimates of alpha for four highly active and liquid futures markets – two equity markets and two 10-year note markets. In all four cases, we chose four times of day for fitting the curve to sweep to fill impact data – a time shortly after each market's open, two times in the middle of each market's trading day, and one time shortly before the close of each market. The average unadjusted alphas for the two US markets were around 1.8, while the average unadjusted European markets were slightly more than 2.0. The hidden liquidity adjustment was greatest for 10-year Treasury note futures and smallest for EuroStoxx futures.

Exhibit 14
Risk aversion parameter (alpha) estimates
(non-event days, January through March 2006)

Market	Estimation Times				Average Alpha	Hidden Liquidity Adjustment	Adjusted Alpha
E-mini S&Ps	8:40am	10:40am	12:40pm	2:50pm	1.85	0.87	1.61
10-year Treasury notes	7:50am	9:50am	11:50am	1:50pm	1.79	0.72	1.29
EuroStoxx	2:10am	5:10am	8:40am	10:20am	2.13	0.93	1.98
EuroBunds	1:10am	4:10am	7:50am	11:50am	2.04	0.85	1.73

Volume, volatility, and market impact profiles

We have now reached the point where we can pull all of the pieces together to produce reliable market impact profiles. Our model of market impact requires three pieces of information – trading volume, price volatility, and an estimate of market makers' risk aversion. An example of how it all fits together is provided in Exhibits 15 through 18.

Exhibit 15 shows two volume profiles for E-mini S&P futures – one for a typical trading day for which there are no scheduled economic announcements of any particular importance. The other is for a trading day on which an FOMC announcement is scheduled. As is apparent, the presence of the FOMC announcement predictably changes the shape of the trading day and the amount of trading done at any given time during the day.

Exhibit 16 shows two price volatility profiles for E-mini S&P futures. Again, one is for a typical trading day, while the other is for an FOMC announcement day.

Using these profiles, we can calculate market impact profiles for orders of various sizes for the two kinds of days. For example, Exhibit 17 shows what market impact profiles would look like for order sizes ranging from 100 to 10,000 contracts at a time on a typical trading day. Exhibit 18 shows what the profiles should look like on an FOMC announcement day.

The combined effect on volume and volatility of an FOMC announcement makes the market predictably less liquid right around the announcement. The effects of the announcement on volume and volatility seem to be roughly the same orders of magnitude. But volatility matters more than volume for market impact, so the liquidity of the market falls at announcement time.

We find it encouraging that these theoretical profiles make sense. For one thing, the impact profiles for a typical trading day look like the sweep to fill profiles shown at the outset in Exhibit 1. The market tends to become less liquid as the trading day progresses until shortly before the close. The theoretical

Exhibit 15
Trading Volume for E-mini S&P Futures
(Thousands of contracts per interval, 3/23/05 to 3/23/06)

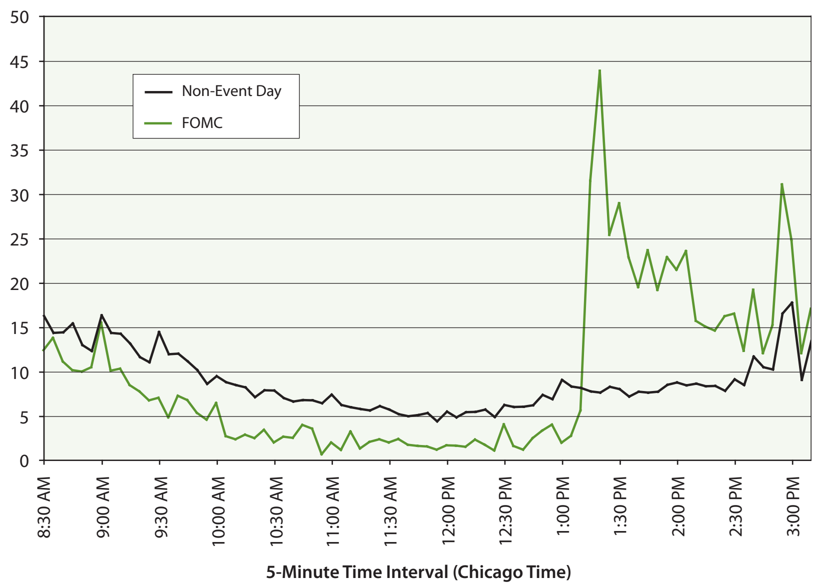

Exhibit 16
Price Volatility for E-mini S&P Futures
(Points per interval, 3/23/05 to 3/23/06)

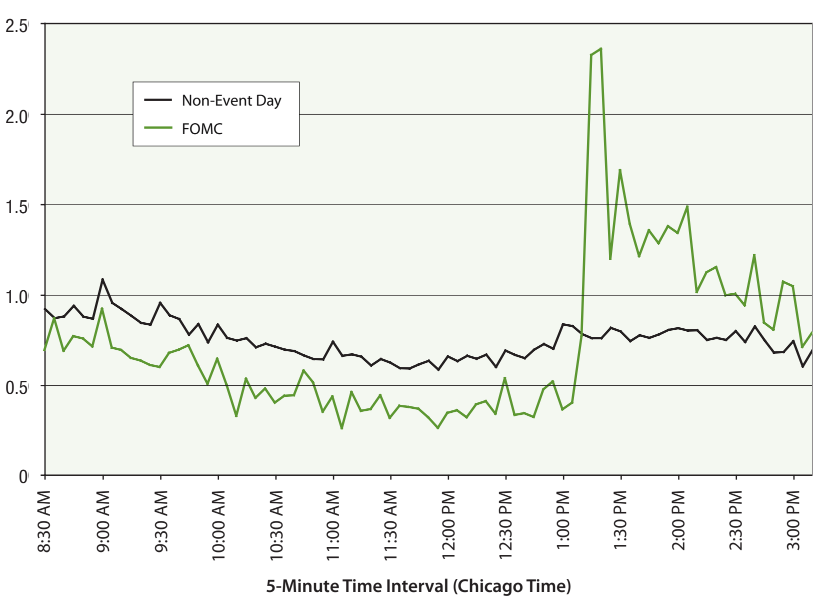

Exhibit 17
Market Impact for E-mini S&P Futures on a Typical Day
(Projected spread to "true" price)

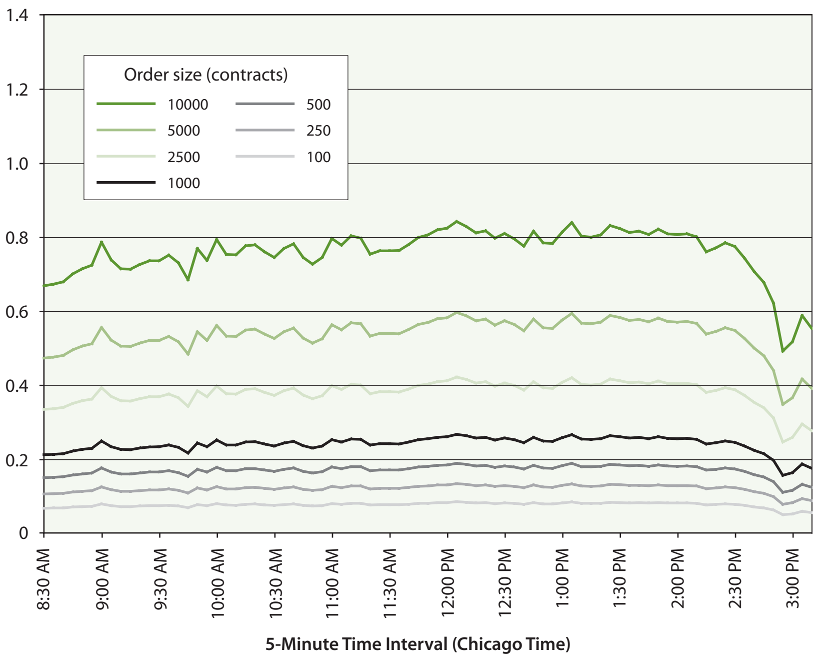

Exhibit 18
Market Impact for E-mini S&P Futures on an FOMC Announcement Day
(Projected spread to "true" price)

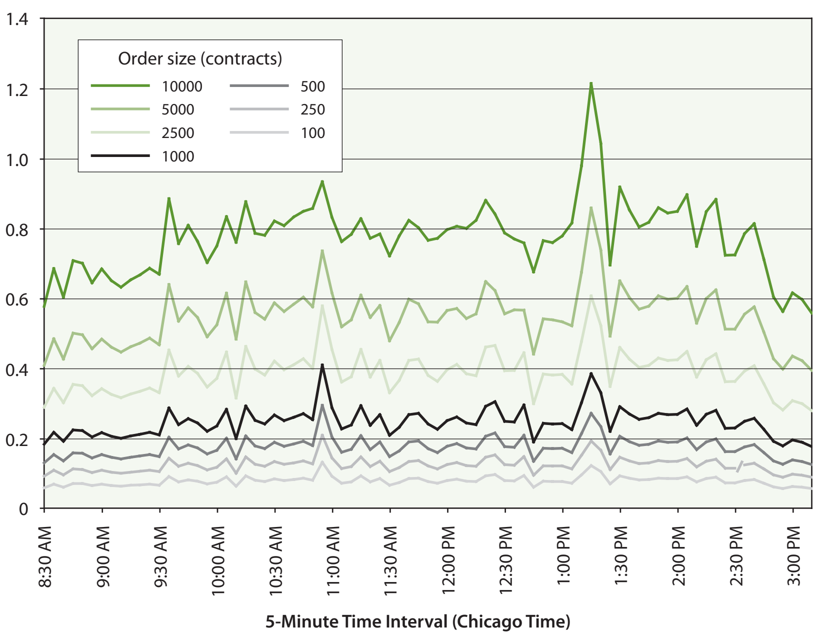

impact estimates are lowest at the cash close – another feature of the observed sweep to fill profiles. Also, the theoretical impact profiles, which have been reckoned using estimates of alpha adjusted for hidden liquidity, are slightly lower than the sweep to fill profiles.

Where do we go from here?

At this point, we have enough confidence in this approach to measuring market impact that we will extend the analysis to cover all electronically traded markets that are active enough to provide us with enough data. Also, we plan to extend the analysis to cover sequential trades so that we can say something useful about very large trades.

Our market impact estimates will be useful in at least three trading applications. These include optimal trading strategies – that is, those that are designed to minimize market impact or produce the best tradeoff between market impact and tracking error. A second is the design of tactical execution rules for working orders. Our insights into liquidity and the possibility of using the limit order book to determine the likelihood of being filled at the bid or the offer will be a natural fit for such tools. A third is in the analysis of benchmarks, tracking error, and execution costs.

Appendix A: Market Data Sets

Contract		*Outrights*	*Spreads*
Equities	Emini SP	12/14/05	12/14/05
	Emini Nasdaq	12/14/05	12/14/05
	Emini Russell	12/14/05	12/14/05
	Canada S&P60	5/4/06	
	DAX	7/5/05	11/16/05
	Eurostoxx 50	7/12/05	11/16/05
	CAC40	1/10/06	3/23/06
	AEX	2/16/06	
	FTSE100	7/28/05	4/25/06
	IBEX35	5/5/06	
	Mini Dow	3/28/06	12/14/05
	SPI 200	5/4/06	
	TOPIX Index	8/30/05	
	Nikkei 225	5/3/06	
	Hang Seng		
	MSCI-Taiwan		
	MSCI-SG		
Bonds	US 30 year	12/14/05	3/29/06
	US 10 year	12/14/05	12/14/05
	US 5 year	12/14/05	3/29/06
	US 2 year	12/14/05	3/29/06
	Canada 10 year	5/4/06	
	Bund	7/12/05	9/20/05
	Bobl	7/28/05	
	Schatz	7/28/05	
	Liffe Gilt	7/27/05	4/25/06
	Australian 10 year	5/4/06	
	CONF	5/5/06	11/16/05
	JGB 10 year	8/30/05	
Money Market	Eurodollar	12/14/05	12/14/05
	Liffe Euribor	7/28/05	
	Fed Funds	5/5/06	
	3 Month Canadian BAs		
	Euroyen		
	Short Sterling		
Currencies	British Pound	12/14/05	
	Euro	12/14/05	
	Japanese Yen	12/14/05	
Energy	Brent Crude (ICE)	4/3/06	
	Gas Oil (ICE)	4/3/06	
	WTCL (ICE)	4/14/06	

Appendix B: Derivation of the Market Depth Model

Consider an order of size *Q* arrives at the market maker at time *a*. We assume the market maker can liquidate the position at the mid price in proportion to the average volume over time. Specifically, the market maker will liquidate the entire position by time *T*, where *T* solves the equation:

$$c\int_a^T dV(t) = Q$$

where $dV(t)$ is the instantaneous rate of market volume traded at time *t*, and *c* is a constant of proportionality representing the relative speed at which the market maker can liquidate inventory: $c = 1$ implies the market maker liquidates at the average rate of trading in the market. Following this notation, let $Q(t)$ be the quantity remaining to be liquidated at time *t*:

$$Q(t) = Q(a) - c\int_a^t dV(s) = c\int_t^T dV(s) = c\left[V(T) - V(t)\right]$$

Let $p(t)$ represent the mid price at time *t*. Suppose the mid price evolves according to an arithmetic Brownian motion:

$$p(t) = p(a) + \int_a^t \sigma(s)dZ(s)$$

where $\sigma(s)$ denotes the instantaneous price volatility at time *s*. Then the gain (or loss) experienced by the market maker in liquidating according to the above schedule is:

$$\begin{aligned}\Pi(T) &= c\int_a^T \left[p(t) - p(a)\right]dV(t) \\ &= c\int_a^T \left(\int_a^t \sigma(s)dZ(s)\right)dV(t) \\ &= \int_a^T \left(c\int_s^T dV(t)\right)\sigma(s)dZ(s) \\ &= \int_a^T Q(s)\sigma(s)dZ(s).\end{aligned}$$

The variance of the market maker's profit is therefore:

$$\begin{aligned}\operatorname{var}\left[\Pi(T)\right] &= \operatorname{var}\left[\int_a^T Q(s)\sigma(s)dZ(s)\right] \\ &= \int_a^T \operatorname{var}\left[Q(s)\sigma(s)dZ(s)\right] \\ &= \int_a^T Q^2(s)\sigma^2(s)\operatorname{var}\left[dZ(s)\right] \\ &= \int_a^T Q^2(s)\sigma^2(s)ds.\end{aligned}$$

Suppose the market maker demands a constant price of risk, α. Then the spread to mid that the market maker will charge for quantity Q is:

$$\text{spread} = \alpha\sqrt{\operatorname{var}\left[\Pi(T)\right]}\Big/Q = \alpha\sqrt{\int_a^T Q^2(s)\sigma^2(s)ds}\Big/Q$$

*Alternative**Edge** series*

- ***How important are daily return data?***
- ***Understanding drawdowns***
- ***Every drought ends in a rainstorm***
- ***Measuring market impact and liquidity***
- ***Superstars versus teamwork***
- ***What you should expect from trend following***
- ***There are known unknowns***
- ***Costs of active management***
- ***Correlations and holding periods:***
  The research basis for the AlternativeEdge Short-Term Traders Index

Please feel free to contact Galen Burghardt (312-762-1140, galen.burghardt@newedgegroup.com) or Ryan Duncan (312-762-1122, ryan.duncan@newedgegroup.com) with questions about the AlternativeEdge series of research notes.