

Bidding principles

Robert Almgren and Neil Chriss show how principal bid programme trades can be priced and evaluated as part of a trading business. By annualising the price impacts and variances of such trades, they construct an information ratio measure that can be used to set hurdles below which bids at a given discount should not be accepted

A programme trade involves the sale or purchase of a basket of stocks that is too large to be traded immediately in the market. When such trades are brokered, they take one of two forms. In an agency trade, the broker executes the trade on behalf of the client on a commission basis, and all the risk of the trade is borne by the client. In a principal trade, also called a principal basket, principal bid or risk bid, the broker directly purchases the entire basket for a fixed price, usually expressed as a discount to fair market value. By design, principal trades transfer all the risk from the client to the broker in exchange for a single price, which therefore proxies for the market price of risk for the portfolio.

Programme trading represents an increasing percentage of overall stock market volume. In 2002, programme trades averaged more than 30% of New York Stock Exchange trading volume, up from 20% in 1999 and 2000. More than half of all programme trading volume took place on the New York Stock Exchange, and, of this, between 30% and 40% was done on a principal basis (New York Stock Exchange, 2002).

We begin with two key observations about the programme business. First, a principal trade consists of two equally important components: a basket of stocks and a price. The basket is determined by the client, but the price – usually expressed as a per share discount to current market value – is agreed upon by client and broker, and the potential profitability of the trade depends on the precise price that can be secured for trading the basket. Throughout this article, we will refer to a trade as the logical unit that consists of both the basket and the price of the basket. The second observation is that since programme trading represents an investment of firm capital, the correct way to view performance measures is on an annualised basis.

This article constructs a mathematical framework for pricing and trading principal baskets, yielding two main results. First, we show that for a broad class of measures of annualised risk-adjusted return, there is a unique optimal way to trade the basket. Second, we introduce a specific measure, the information ratio of a trade, which is the ratio of annualised expected profit to annualised standard deviation of profit.

Given a proposed trade one can calculate its information ratio from known information about the basket. Because the information ratio is annualised by the expected time to complete the liquidation, it provides a way of comparing the potential profitability of trades of different sizes and levels of liquidity. This yields a powerful tool for analysing principal business.

One does not need to know all the constituents of the basket to calculate its information ratio, just the volatility and liquidity and a proposed price for the basket. In particular, one must know the effect on the prices in the basket of trading the constituents, that is, the market impact. The methods in this article work with a wide variety of market impact functions, making it possible to support a wide variety of existing models.

The information ratio is primarily a pricing tool. By specifying a hurdle rate – a minimum information ratio that every principal basket must exceed – one can determine for a given principal basket what the minimum price is that will exceed the hurdle rate. Alternatively, it may be used as an evaluation tool. If one knows the ‘going price’ of a trade, one can cal-

culate the information ratio of the trade based on that price to determine whether it is worthwhile to submit a winning bid for the business.

Efficient execution strategies. The starting point for this analysis is our previous work (Almgren & Chriss, 1999 and 2000) on the optimal liquidation of a portfolio in the presence of market volatility and transaction costs (see also Kissell & Malamut, 2002). These two factors are the main contributors to the cost of trading, or implementation shortfall (Perold, 1988), that is, the difference between the initial paper value of the portfolio and the total final proceeds of transacting. This work shows that for a given portfolio and a utility function that establishes a trade-off between expected cost and volatility risk, there is a unique optimal strategy.

A liquidation strategy is an algorithm for performing a large transaction by breaking it up into smaller pieces. Thus a strategy consists of a sequence of small trades over a period of time, which together comprise the large trade. As in our earlier work, we call the sequence of trades a ‘trading trajectory’.

Among all trading trajectories for a particular portfolio, there is a particular family that is potentially optimal, or efficient. Efficient trajectories minimise trading cost for a given level of variance. If one plots every possible trading trajectory on a plane by determining the x-co-ordinate as the variance of cost and the y-axis as expected cost, the set of attainable costs and variances is a convex and closed region, and the efficient strategies lie on its boundary. The one-parameter family of efficient trajectories form an efficient frontier, analogous to the efficient frontier in modern portfolio theory.

Strategies on the frontier may be parameterised by their speed of execution. Efficient strategies that have lower variance require more rapid trading to reduce market exposure, incurring greater transaction costs. Conversely, strategies emphasising lower transaction costs do so by trading more slowly and increasing variance.

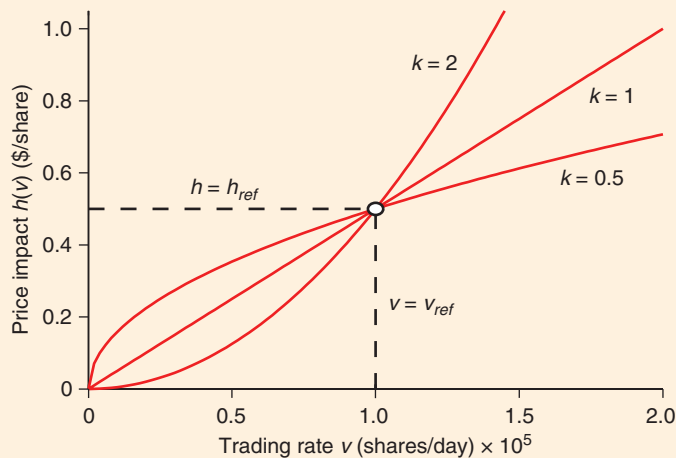
Selecting a particular optimal point on the efficient frontier requires the specification of a criterion for measuring the trade-off between risk and reward. For a single portfolio transaction, this can be done using either a utility function or a value-at-risk-type analysis (Almgren & Chriss, 2000, and Konishi & Makimoto, 2001).

Application to pricing programme trades. A programme trade requires a use of capital by a trading desk. In this article, this use of capital is evaluated employing an *ex ante* risk-adjusted return ratio – the ratio of the predicted profit to the standard deviation of predicted profit. This turns out to be the familiar information ratio, analogous to the familiar Sharpe ratio of portfolio theory.

We argue that for a principal desk engaged in ongoing business, the correct approach is to annualise cost and variance, placing them in the context of other investment opportunities. Remarkably, for each value of the discount received for trading the basket, there is then a single optimal point, independent of risk preferences. This point corresponds to a single best value for the overall information ratio.

The value of this ratio is therefore a potential tool to be used in evaluating whether to accept a certain piece of business at a certain price. In many situations, programme trading business is ‘put out to bid’. That is,

1. Non-linear impact functions



Non-linear impact function $h(v) = \eta v^k = h_{ref}(v/v_{ref})^k$. Once a reference trading rate and impact value have been chosen, the exponent k is a way to extrapolate to smaller and larger rates

the portfolio manager contacts multiple desks about a particular portfolio. Each business responds with a certain bid – the discount to fair market value required to do the trade. Programme desks often know the level required to win the business, and therefore the information ratio can be used as a hurdle or evaluation tool to decide whether or not to bid at a level to win the business.

Assumptions. To clearly describe our methodology, we will review the main assumptions that underlie the previous work on the subject.

■ **Brownian motion.** We assume that each stock in the basket follows a Brownian motion with absolute volatility σ (see equation (1) below). That is, in each time interval Δt the change in share price is normally distributed with mean zero and standard deviation $\sigma\sqrt{\Delta t}$. This is clearly different from the standard geometric Brownian motion description, but in Almgren & Chriss (2000) we show that for the purpose of short-horizon liquidations the difference is negligible.

We set the drift to zero under the assumption that over the short time horizons in which liquidations are undertaken, the scale of volatility is much greater than that of expected return. This was discussed in greater detail in Almgren & Chriss (1999). Similarly, we do not discount costs back to $t=0$. A more precise description would interpret all growth rates as excess quantities over a risk-free rate: the same conclusions would be obtained.

■ **Continuous time.** We construct our models in continuous time, even though actual trading takes place at discrete moments in time. In Almgren & Chriss (2000), the theory was worked out in discrete time and the differences are unimportant for this work. The main aspect of continuous time trading for this article is that at each moment in time there is an instantaneous rate of trading v that determines the instantaneous cost of trading due to market impact.

■ **Market impact.** Assessing the potential cost of trading requires a market impact model. Market impact is difficult to model and notoriously difficult to measure. The baseline assumption of all market impact models (Kraus & Stoll, 1972, Holthausen, Leftwich & Mayers, 1990, and Chan & Lakonishok, 1995) is that there are two different types of market impact, temporary and permanent. Temporary impact derives from paying a premium for immediacy while permanent impact refers to the component of price movements due to trading that represent a fundamental shift in the equilibrium price of the stock.

We model temporary market impact as a specific function of trading

rate. In particular, as in Almgren (2003), we deal with non-linear impact functions of the form where market impact, as measured in adverse share price movement, is proportional to the trading rate raised to a positive power k . This subsumes a wide variety of models, including, for example, the Barra transaction cost model (Barra, 1997) and the linear transaction cost model of Almgren & Chriss (2000). It is consistent with the brief description of effective bid-ask spread as a function of block size in Perold (1988), in which convexity or concavity of the function represents different beliefs of the market about the trader.

Figure 1 shows this model, and how it may be calibrated. First, we choose a reference trading rate v_{ref} . For any particular choice of trading interval τ , v_{ref} is equivalent to a certain block size $n_{ref} = \tau v_{ref}$ traded in that time period; it may be interpreted as the market ‘depth’ in the sense of Kyle (1985) or Bondarenko (2001). For our examples, we will consider a stock that trades 1 million shares per day, and we will take v_{ref} to be 10% of that rate, or $v_{ref} = 100,000$ shares per day. For time period $\tau = 1$ hour, with 6.5 periods a day, this rate is equivalent to trading a block of approximately 15,300 shares in each hour.

Next, we choose the price impact h_{ref} , which would be incurred by steady trading at the reference rate v_{ref} . In our example, we shall assume the share price is \$50 a share, and we assume that trading $v_{ref} = 100,000$ shares per day incurs a price impact of 1%, or 50 cents a share.

To extrapolate this impact to smaller and larger trading rates or block sizes, we choose a value for the exponent k based on theoretical modelling or intuitive belief. The choice $k = 1$ corresponds to linear dependence of price impact on rate, $k > 1$ means that smaller trading rates or block sizes have a disproportionately small effect on price, while $k < 1$ means that smaller trading rates or block sizes have a relatively larger impact.

We model permanent impact as a simple linear function of trading rate (see equation (1) below). A consequence of this assumption is that the total cost incurred due to permanent impact is independent of the trading strategy (see equation (2)): the market simply responds to the total number of new shares in the block. Although permanent impact does not affect the way in which the shares are traded, it makes an important contribution to the total cost and hence to the necessary discount to fair in the price for the basket.

Optimal trading reviewed

We briefly review the notion of an optimal trading trajectory, as introduced in Grinold & Kahn (1995) and Almgren & Chriss (2000) and developed further in Almgren (2003) and Konishi & Makimoto (2001). The main idea is a two-step approach to minimising a particular risk/reward measure. First, we determine a one-parameter family of efficient trading strategies that minimise the uncertainty of trading cost for a given cost of trading or conversely minimise the cost of trading among all strategies with a given level of uncertainty. Second, we search among that family for the particular solution that optimises our chosen risk/reward measure. In the context of a principal bid, it is important to include the discount in evaluating risk and reward.

■ **Optimal liquidation trajectories.** We will develop the theory entirely for the case of a basket with a single stock. The case of a portfolio – the more important case in applications – is conceptually similar but more complex in its details. When a programme business wins a bid, it acquires a portfolio consisting of X shares of a single asset. The objective is to liquidate this portfolio ‘efficiently’, meaning that a suitable balance is found between the loss of value due to impact costs incurred by trading rapidly and the uncertainty assumed by trading slowly.

We define a trading trajectory to be a function $x(t)$ defined for $0 \leq t < \infty$, which specifies the number of shares of the asset remaining in the portfolio at time t . This is a mathematical abstraction of the notion of the ‘way’ the basket is traded. The trajectory function must have $x(0) = X$, and the optimal solutions will have $x(t) \rightarrow 0$ as $t \rightarrow \infty$ so that the basket is ultimately liquidated. In addition, optimal trajectories are differentiable and decreasing in t . In this article, we will also call $x(t)$ a trading strategy and

will use the terms trajectory and strategy interchangeably.

A key insight is that under a mean-variance optimality criterion, and only under that criterion, optimal trajectories can be determined ‘statically’, that is, at the start of trading. Information about market movements observed during the course of trading does not modify our future trading decisions. For other optimality criteria, the statically optimal trajectories are not dynamically optimal; fully optimal strategies should be determined by dynamic programming with an appropriate set of time-dependent risk preferences. Nonetheless, the statically chosen trajectories are generally considered to be adequate for practical purposes.

We have taken the continuous-time limit for analytical simplicity. In this limit, $v(t) = -\dot{x}(t)$ (dot denotes time derivative) is the instantaneous rate of selling at time t , in shares per day. We assume that the asset market price $S(t)$ follows a Brownian motion:

$$dS = \sigma dW - \gamma v dt = \sigma dW + \gamma dx \quad (1)$$

where σ is a volatility and $W(t)$ is a standard Wiener process. Note that σ is an absolute volatility, measured in dollars per share rather than per cent (per square root of time).

The coefficient γ incorporates linear permanent impact: each share purchased drives up the long-term equilibrium price by an amount γ . We can integrate the evolution law to get:

$$S(t) = S_0 + \sigma W(t) - \gamma(X - x(t))$$

in which S_0 is the market price at the start of trading. The cumulative permanent price impact at any time is simply proportional to the number of shares traded up to that time. The total dollar cost incurred by permanent impact over the entire course of the trade programme is:

$$C_{perm} = -\int_0^\infty \gamma x(t) v(t) dt = \frac{1}{2} \gamma X^2 \quad (2)$$

which is independent of the specific trajectory x , assuming only that $x(t) \rightarrow 0$ as $t \rightarrow \infty$. Thus permanent impact has no effect on the nature of the optimal strategy, but makes a substantial contribution to its overall cost.

When the programme desk transacts at time t , the price at which the transaction takes place is given by:

$$\tilde{S}(t) = S(t) - h(v(t))$$

where the non-linear function $h(v)$ represents the temporary market impact given as a function of trading rate v . This represents the discount a trader has to offer the market to trade with him.

With these assumptions, the total cost incurred during trading due to temporary market impact is:

$$C_{temp}[x] = \int_0^\infty h(v(t)) v(t) dt = \int_0^\infty -\dot{x}(t) h(-\dot{x}(t)) dt$$


At each moment of time, we incur a cost that is the number of shares we sell in that moment, $v(t)dt$, multiplied by the per-share impact cost, $h(v(t))$. As noted above, we do not discount this cost back to $t = 0$, under the assumption that trading horizons are short enough that this can be neglected, or that all quantities are discounted at the same rate. The scalar quantity $C_{temp}[x]$ is a functional of the entire trading strategy $x(t)$.

The variance of our trading result is defined by the functional:

$$V[x] = \int_0^\infty \sigma^2 x(t)^2 dt$$

For each moment of time that we still have the asset in our portfolio, the eventual value we will obtain from the sell programme changes in proportion to the square of the number of shares we hold. The total variance is equal to the integral of variance over infinitesimal intervals because of the independent increments of the arithmetic Brownian motion.

2. Optimal trajectories



Optimal trajectories for three values of the exponent $k = 0.5, 1, 2$. Time is measured in units of the characteristic time T , and shares held are measured as fractions of the initial portfolio size. Regardless of the details, after one characteristic time, approximately 60% of the portfolio remains

Our objective is to choose the trajectory $x(t)$ to minimise both C and V , with some particular criterion to evaluate the trade-off between them. As outlined in Almgren & Chriss (2000), the most tractable approach is first to determine the family of trajectories that minimise V for a given C . To do this, we introduce a Lagrange multiplier λ , and seek to minimise the combined quantity:

$$U[x] = C_{temp}[x] + \lambda V[x] = \int_0^\infty \left[\lambda \sigma^2 x(t)^2 - \dot{x}(t) h(-\dot{x}(t)) \right] dt$$

The minimum is to be taken over all functions $x(t)$ having $x(0) = X$. It is also easy to see that optimal trajectories must have $x(t), \dot{x}(t) \rightarrow 0$ as $t \rightarrow 1$ (trajectories that do not tend to zero have infinite variance).

Standard techniques of the calculus of variations (Almgren, 2003) show that with the specific choice:

$$h(v) = \eta v^k \quad (3)$$

the unique minimising solution with the given boundary conditions is:

$$\frac{x(t)}{X} = \begin{cases} \left(1 + \frac{1-k}{1+k} \frac{t}{T}\right)^{-(1+k)/(1-k)} & \text{if } 0 < k < 1 \\ \exp\left(-\frac{t}{T}\right) & \text{if } k = 1 \\ \left(1 - \frac{k-1}{k+1} \frac{t}{T}\right)^{(k+1)/(k-1)} & \text{if } k > 1 \end{cases} \quad (4)$$

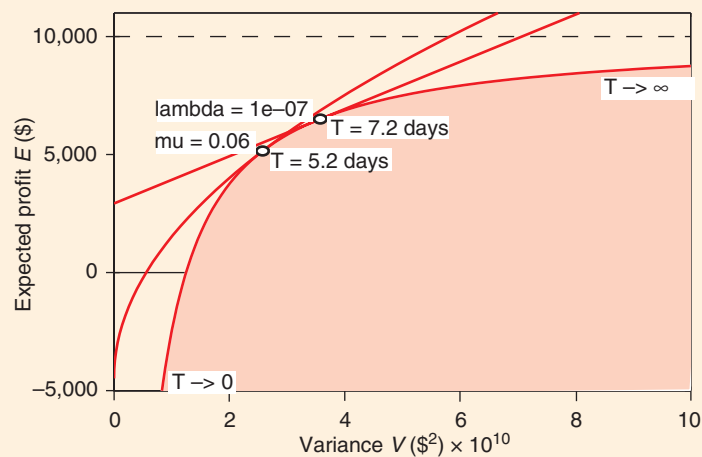
where T , known as the ‘characteristic time’, is given by:

$$T = \left(\frac{k\eta}{\lambda\sigma^2}\right)^{1/(k+1)} X^{(k-1)/(k+1)}$$

These solutions tend to zero as $t \rightarrow \infty$. In fact, for $k > 1$, the portfolio holdings reach zero at a finite time T_{max} prior to $t = \infty$ (Almgren, 2003). See figure 2.

The characteristic time attached to a strategy x is a reasonable estimate of the time in which the liquidation process is ‘substantially’ completed and represents a reasonable description of the rate of trading: the initial

3. Efficient frontier



Efficient frontier including discount, for a single principal bid. Points far down along the vertical axis correspond to very rapid trading: variance is minimised at high expense. Points far to the right correspond to very slow trading: the full discount is kept as profit but high variance is incurred. The tangent line and tangent curve indicate optimal points for two different risk/reward functions: a mean-variance criterion $E - \lambda V$ and the value-at-risk measure $E - \mu\sqrt{V}$

A. Parameter values for example case

Initial stock price	$S_0 = \$50$ a share
Basket size	$X = 100,000$ shares
Annual volatility 32%	$\sigma = 1$ (\$/share)/ $\sqrt{\text{day}}$
Daily volume	$v_{\text{day}} = 1$ million shares
Reference trading rate	$v_{\text{ref}} = 10\%$ $v_{\text{day}} = 10^5$ shares a day
Reference price impact	$h_{\text{ref}} = 1\%$ $S_0 = 50$ cents per share
Exponent	$k = 1$
Impact coefficient	$\eta = h_{\text{ref}}/v_{\text{ref}}^k = 5 \times 10^{-6}$ (\$/share)/(share/day)
Discount	$D = 10$ cents per share

rate is $-\dot{x}(t=0) = X/T$, and the rate slows as the portfolio is partially liquidated. This time depends on the initial portfolio size X and the Lagrange multiplier (risk-aversion coefficient) λ , as well as on market parameters.

As noted, trading trajectories are static in the sense that information observed during the trading period does not affect the trajectory. This should not be confused with merely stating that the trading trajectory is optimal from the point of view of $t=0$. Rather, if at some time $t_1 > 0$ the programme trader re-evaluates his remaining holdings as a fresh liquidation with trade-off parameter λ and with a portfolio $X_1 = x(t_1)$ (the portfolio contents at time t_1), the optimal trajectory thus produced would be identical to the remaining trajectory on the original liquidation. More precisely, if the new optimal trajectory is $x_1(t)$, then $x_1(t) = x(t)$ for $t \geq t_1$.

For a fixed starting portfolio size, the execution time T is a one-to-one function of our artificial parameter λ , and hence (for given values of η and k) we may consider the trajectories to be parameterised by T rather than λ . That is, once an approximate time has been chosen over which the programme is to be executed, the solution (4) is the optimal trading strategy on that time scale, depending only on the exponent k .

The efficient frontier including discount. We now stop thinking of the problem of liquidating a basket as disconnected from the price of the basket. Our aim is to eventually calculate the information ratio of liquidating a basket incorporating the value of the discount to fair value

received in the transaction. Thus, we assume that the trader will receive a discount of D dollars per share for the basket in a principal trade and explicitly calculate the expected cash received in trading out the basket and its variance. For example, if the programme desk were able to dump the entire portfolio into the market with no market impact, it would earn a profit of DX dollars.

In general, because of market impact, the total expected profit of the transaction would be $E = DX - C_{\text{perm}} - C_{\text{temp}}$, and the variance will be as stated above. We can explicitly calculate the expected profit and its variance as functions of the execution time T :

$$E(T) = \left(D - \frac{1}{2}\gamma X \right) X - \frac{k+1}{3k+1} \eta \left(\frac{X}{T} \right)^k X$$

$$V(T) = \frac{k+1}{3k+1} \sigma^2 T X^2$$

Shortly we will use $E(T)$ and $V(T)$ to construct the information ratio of a trade, but for now we examine some of their properties. Clearly, what we see is that the expected profit of the trade is the total discount DX reduced by the permanent impact and by a temporary impact amount. The effect of permanent impact is to reduce the expected profit per share, as reflected in the size of the discount D , by an amount proportional to the portfolio size, while temporary impact is proportional to a per share reduction in expected profit of $(X/T)^k$.

It is worth noting the dependence of E and V on T . Short liquidation times $T \rightarrow 0$ correspond to $E \rightarrow -\infty$ and $V \rightarrow 0$: all profit is dissipated in impact costs but no market risk is incurred. Long times $T \rightarrow \infty$ correspond to $E \rightarrow (D - \frac{1}{2}\gamma X)X$ and $V \rightarrow \infty$: temporary impact costs are avoided completely by holding the portfolio essentially forever, but at the expense of any certainty of profit. (In principle, if the portfolio were held forever without trading then permanent impact costs would also be avoided: $T = \infty$ is not the same as the limit $T \rightarrow \infty$.) Assuming that the discount is at least enough to compensate for permanent impact, there is an intermediate point at which $E = 0$ (the zero profit trade): impact costs are exactly compensated by the discount on average, but risk is taken to achieve this.

This 'efficient frontier' is illustrated in figure 3 for typical parameter values taken from Almgren & Chriss (2000) and Almgren (2003) and summarised in table A. The discount is taken to be $D = 10$ cents per share on a portfolio of $X = 100,000$ shares, so the maximum possible profit is $DX = \$10,000$. We take linear impact for simplicity.

Our previous work (Almgren & Chriss, 2000) focused on various ways to draw tangent lines to this frontier, in order to maximise either a mean-variance criterion $E + \lambda V$ or a VAR measure $E + \mu\sqrt{V}$ for a single trade in isolation. Now we want to consider this trade as part of an ongoing business.

Performance measures

Here, we define the information ratio of a principal trade assuming a given discount of D dollars per share for the basket. First, we note that the above analysis did not take account of the fact that different baskets will have different optimal liquidation times. If principal bids are to be considered in the context of an ongoing business in relation to multiple investment opportunities, then we must view the expected profit of a trade in units that are comparable across different optimal times. We do this directly by annualising the expected return by the expected amount of time it takes to liquidate substantially all of the basket, as determined by the characteristic time of the trade T . If a positive profit can be made, that is, if $E > 0$, then the trader prefers a shorter liquidation time to a longer time, other things being equal.

Annualisation. To annualise the expected profit and its variance, we assume that the full invested capital becomes available for reinvestment after one characteristic time T . In fact, as can be seen from the trajectories in figure 2, liquidation is a continuous process: some capital is available im-

mediately, and recovery of the full capital formally requires an infinite length of time. Nonetheless, T is a reasonable average value, and is the simplest way to compare different trajectories.

Assuming that T is measured in years, annualising is simply a matter of dividing by T . The expectation and variance per year of trading are:

$$\frac{E(T)}{T} = \frac{(D - \frac{1}{2}\gamma X)X}{T} - \frac{k+1}{3k+1} \eta \left(\frac{X}{T}\right)^{k+1}$$

$$\frac{V(T)}{T} = \frac{k+1}{3k+1} \sigma^2 X^2$$

The annualised expectation is composed of two terms. The first term is the average rate at which you accept the discount payment D , reduced by the cost of permanent impact. Since this is a fixed amount per portfolio, it is increased by trading rapidly. The second term is the impact costs incurred by permanent trading at a constant rate X/T , with a numerical correction coefficient to account for the non-linear shape of the trajectory.

Note that the annualised variance is independent of liquidation time. This can be interpreted as saying that in the course of repeated execution, you are always invested in the market by the same amount on average. This has an important implication that is one of the key results of this article. It says that if you re-cast the efficient frontier in terms of annualised expectation and annualised variance, it collapses to a single point.

The 'annualised' efficient frontier is shown in figure 4. The shaded region in figure 3 has collapsed into a half-infinite vertical line and the frontier itself has collapsed to a single point. This is a direct consequence of the annualised variance being independent of trading time. The striking consequence of this is that any measure of risk-adjusted profitability as constructed from a tangent line or curve, for example, as shown in figure 3, regardless of functional form or parameter values, will pass through the highest point on this line. This means in particular that there is a unique best way to trade for any reasonable risk-adjusted return measure, regardless of any particular risk/reward preferences, and it is found simply by finding the value of T that maximises $E(T)/T$. This gives:

$$T_{opt} = \left(\frac{(k+1)^2}{3k+1} \right)^{1/k} \frac{\eta^{1/k} X}{(D - \frac{1}{2}\gamma X)^{1/k}}$$

To emphasise our point once more, T_{opt} is the parameter representing the optimal trading strategy across a broad spectrum of possible risk-adjusted return measures. It is independent of risk/reward preferences, but depends on the discount D . Our next aim is to define and evaluate a particular risk-adjusted return measure.

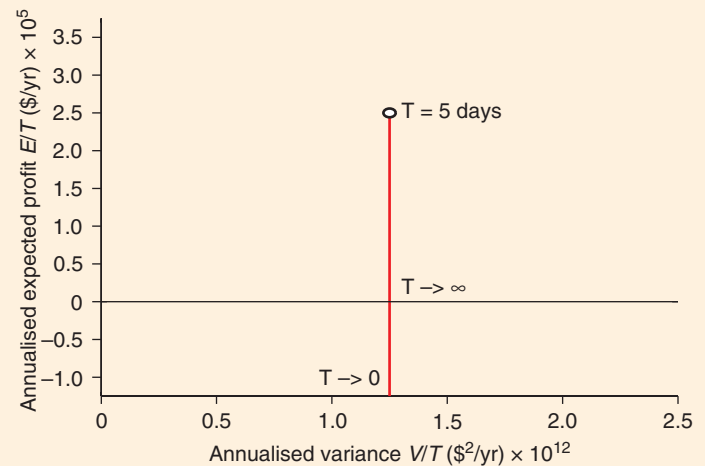
Definition of the information ratio. We define and calculate the information ratio of a trade, including the effect of the discount D received in the principal transaction. For a given characteristic time T , the information ratio with respect to T represents the annualised risk-adjusted expected profit that may be achieved by trading along the trajectory with parameter T :

$$I(T) = \frac{E(T)/T}{\sqrt{V(T)}/T}$$

Note that $I(T)$ is the risk-adjusted return for the basket implicitly assuming a discount D received for the trade and a trading time parameter T . The question is, for which T is $I(T)$ maximal? The answer to this is the information ratio of the trade and is clearly given by T_{opt} . We can substitute this into the formula for $I(T)$ to give the following formula:

$$I_{max} = \frac{(3k+1)^{(k+2)/2k} (D - \frac{1}{2}\gamma X)^{(k+1)/k}}{(k+1)^{(3k+4)/2k} X \eta^{1/k} \sigma} \quad (5)$$

4. Annualised efficient frontier



'Efficient frontier' incorporating the price of the basket, for annualised principal bids. The variance per year is independent of the trading time. Thus the accessible region collapses to a single line. Expected profit per year has a single internal maximum indicated by the disc. This represents the optimal trading strategy for repeated transactions, regardless of the risk/reward trade-off criterion

B. Specific optimal solutions

k	T	I_{max}
1/2	$0.810 \times \frac{\eta^2 X}{D^2}$	$1.063 \times \frac{D^3}{X \sigma \eta^2}$
1	$\frac{\eta X}{D}$	$0.707 \times \frac{D^2}{X \sigma \eta}$
2	$1.134 \times \frac{\eta^{1/2} X}{D^{1/2}}$	$0.449 \times \frac{D^{3/2}}{X \sigma \sqrt{\eta}}$

Optimal trading time and maximum information ratio for three values of the market impact exponent k . We denote $\tilde{D} = D - \frac{1}{2}\gamma X$, the discount reduced by anticipated permanent impact costs

Since the numerator and denominator in the definition of I are both proportional to portfolio size, it would be equivalent to consider E and \sqrt{V} as percentage return and risk. Thus this quantity allows comparisons across baskets and other investment opportunities of arbitrary size. It has units of $year^{-1/2}$, and thus should be compared only with other annualised measures.

Applications to bidding. We can now state the two main applications of the information ratio as the answers to two questions: for a given level of discount D that we can demand for this trade, what is the information ratio of the basket, and for a given information ratio hurdle, what is the minimum discount we must demand in order to clear this hurdle?

The answer to the first question is simply I_{max} , the information ratio of the basket assuming a discount of D . Put a different way, equation (5) tells us that given a discount D that a desk can demand for a trade, it will yield an information ratio of I_{max} and this information ratio cannot be exceeded. This can then be used to determine if the trade clears a particular hurdle rate or not.

The second question may be answered simply by inverting (5) to yield:

$$D_{\min} = \frac{1}{2} \gamma X + \left(\frac{(k+1)^{(3k+4)/2k}}{(3k+1)^{(k+2)/2k}} X \eta^{1/k} \sigma I_{\min} \right)^{k/(k+1)} \quad (6)$$

The formula for D_{\min} gives the maximum that a desk can bid for a given basket while still clearing the minimum information ratio threshold of I_{\min} .

In table B, we present the specific forms of these expressions for a few particular and important values of k . Although the analytical expressions above are complex, for a specific choice of k they reduce to simple numerical coefficients. What is particularly noteworthy is the relationship between price of the principal basket (as embodied in the discount to fair and both the information ratio and optimal time for liquidation. It depends on the market impact function assumed, and can be quite sensitive to small movements. For example, for the Barra model $k = \frac{1}{2}$, the maximum information ratio increases as the cube of the discount, after allowing for permanent impact, and the optimal time decreases as the square of the discount.

One interpretation of these relationships is that small changes in the price of a principal bid, expressed as a cents-per-share discount to fair value, can have a significant impact on both the risk-adjusted profitability of the trade and the time it takes to liquidate the trade. For instance, a basket that commands a 2.5 cents-per-share discount to fair is more than twice as profitable on a risk-adjusted returns basis versus one that commands a 2 cents-per-share discount.

Conclusion

This article introduces a framework for analysing and pricing principal bids by means of a risk-adjusted return measure applied to the trading principal baskets called the information ratio. The information ratio may be used to compare baskets of different sizes, volatility and liquidity, traded across widely varying time horizons. For a given set of portfolio parameters, there is a unique way to trade that maximises the information ratio. This offers programme traders and risk managers two analytical tools. First, it provides a measure of risk-adjusted profitability for a principal trade at a given price point; this can be used to evaluate a given bid for a principal basket in comparison with other investment opportunities. Second, it may be used to construct information ratio hurdles, in the spirit of return on capital hurdles for trades, which a given principal trade must clear. Traders may then set their price for the basket to achieve this hurdle. ■

Robert Almgren is a professor in the mathematics and computer science departments at the University of Toronto. Neil Chriss is managing director of quantitative strategies at SAC Capital Management in New York. E-mail: neil.chriss@sac.com, almgren@math.toronto.edu

REFERENCES

Almgren R, 2003
Optimal execution with nonlinear impact functions and trading-enhanced risk
Applied Mathematical Finance 10, pages 1–18

Almgren R and N Chriss, 1999
Value under liquidation
Risk December, pages 61–63

Almgren R and N Chriss, 2000
Optimal execution of portfolio transactions
Journal of Risk 3(2), pages 5–39

Barra, 1997
Market impact model handbook

Bondarenko O, 2001
Competing market makers, liquidity provision, and bid-ask spreads
Journal of Financial Markets 4(3), pages 269–308

Chan L and J Lakonishok, 1995
The behavior of stock prices around institutional trades
Journal of Finance 50, pages 1,147–1,174

Grinold R and R Kahn, 1995
Active portfolio management
Probus Publishing, chapter 13, pages 291–314

Holthausen R, R Leftwich and D Mayers, 1990
Large-block transactions, the speed of response, and temporary and permanent stock-price effects
Journal of Financial Economics 26, pages 71–95

Kissell R and R Malamut, 2002
Optimal trading strategies
Unpublished draft

Konishi H and N Makimoto, 2001
Optimal slice of a block trade
Journal of Risk 3(4), pages 33–51

Kraus A and H Stoll, 1972
Price impacts of block trading on the New York Stock Exchange
Journal of Finance 27, pages 569–588

Kyle A, 1985
Continuous auctions and insider trading
Econometrica 53, pages 1,315–1,336

New York Stock Exchange, 2002
Press releases

Perold A, 1988
The implementation shortfall: paper versus reality
Journal of Portfolio Management 14(3), pages 4–9

The Basel Committee's recent CP3 paper will have a direct impact on your business
Read how it will affect you at www.BaselAlert.com/CP3

CP3 Alert in association with
Deloitte & Touche & **BaselAlert.com**

