Bayesian Adaptive Trading with a Daily Cycle

Robert Almgren* and Julian Lorenz**

July 28, 2006

Abstract

Standard models of algorithmic trading neglect the presence of a daily cycle. We construct a model in which the trader uses information from observations of price evolution during the day to continuously update his estimate of other traders' target sizes and directions. He uses this information to determine an optimal trade schedule to minimize total expected cost of trading, subject to sign constraints (never buy as part of a sell program). We argue that although these strategies are determined using very simple dynamic reasoning—at each moment they assume that current conditions will last until the end of trading—they are in fact the globally optimal strategies as would be determined by dynamic programming.

*Electronic Trading Services, Banc of America Securities LLC, New York; Robert.Almgren@bofasecurities.com.

**Institute of Theoretical Computer Science, ETH Zürich; jlorenz@inf.ethz.ch.
Partially supported by UBS AG.

1 Introduction

This paper presents a model for price dynamics and optimal trading that explicitly includes the daily trading cycle and the trader's attempt to learn the targets of other market participants. This is in contrast to most current models of optimal trading strategies, that view time as an undifferentiated continuum, and that view other traders as a collection of random noise sources. This paper has two primary motivations.

The first set of motivations is academic articles by Brunnermeier and Pedersen (2005) and Carlin, Lobo, and Viswanathan (2005). In these articles, institutional trading has an explicit daily cycle, based on the assumption that at the beginning of each day, each informed market participant, or institutional investor, is exogeneously given a trade target. These participants know the targets of the other informed traders, and they must decide whether to cooperate with their peers so as not to lose value to uninformed traders, or whether to compete and take value from their peers. The novel feature of our model is that the participants do not know each others' targets, but must guess them by observing price dynamics throughout the day. We take for granted that informed participants will use all available information to compete with each other.

The second set of motivations is the popularity of execution algorithms that adapt to changes in price of the asset being traded, either by accelerating execution when the price moves in the trader's favor, or conversely. Although these optimal trade models may be derived by introducing various forms of risk aversion (Kissell and Malamut 2006; Almgren and Lorenz 2006), the most common justification for them is a belief in mean reversion or momentum of the asset price.

Our model may be interpreted as one plausible way to model price momentum. There is an underlying drift factor, caused by the net positions being executed by other institutional investors. This factor is approximately constant throughout the day because other traders execute across the entire day. Thus price increases in the early part of the day suggest that this factor is positive, which suggests that prices will continue to increase throughout the day. This is different from a short-term momentum model in which the price change across one short period of time is correlated with the price change across a preceding period; most empirical evidence shows that such correlation is weak if it exists at all. Our strategies will exploit this momentum to minimize the expected value of trading costs, somewhat in the spirit of Bertsimas and Lo (1998), except that because we focus on long-term momentum, our results can obtain higher gains.

The daily cycle is an essential feature of this model. Large institutional participants make investment decisions overnight and implement them through the following trading day. Within each day, the morning is different from the afternoon, since an intelligent trader will spend the early trading hours collecting information about the targets of other traders, and will use this information to trade in the rest of the day.

By contrast, in the market view that is implicitly assumed by most models, trade decisions are made at random times and trade programs have random durations, with no regard to the daily cycle. Thus, if one observes buy pressure from the market as a whole, one has no reason to believe that this pressure will last more than a short time. From the point of view of optimal trading, price motions are purely random.

In addition, we incorporate the very important feature of constraints on trade direction: the trader must never sell as part of a buy program, even if this would yield lower expected costs (or give an expected profit) because of anticipated negative drift in the price. This is for two reasons: First, we take the point of view of an broker/dealer executing an agency trade for a client. Second, we neglect the bid/offer spread and other fixed costs, which may greatly reduce the profitability of such reversing strategies. Our adaptive strategies simply shift buying or selling from one time period to another. This constraint is often binding, and globally affects the structure of optimal solutions. In many cases, it leads to the determination of an optimal end time for trading, and sometimes directs the strategy to stop trading completely for a finite time period in the middle of execution.

In Section 2 below, we present our model of Brownian motion with a drift whose distribution is continuously updated using Bayesian inference. In Section 3 we present optimal trading strategies which, surprisingly, can be determined by computing a "static" optimal trajectory at each moment, assuming that the best parameter estimates at that time will persist until the end of the day.

2 Price Model including Bayesian Update

We consider trading in a single asset whose price is $S(t)$, obeying an arithmetic random walk

$$S(t) = S_0 + \alpha t + \sigma B(t) \qquad \text{for } t \geq 0, \tag{1}$$

where $B(t)$ is a standard Brownian motion, σ an absolute volatility and α a drift. In the presence of intraday seasonality, we interpret t as a volume time relative to a historical profile.

Our interpretation is that volatility comes from the activity of the "uninformed" traders, whose average behavior can be predicted reasonably well. Mathematically, we assume that the value of σ is known precisely (for a Brownian process, σ can be estimated arbitrarily precisely from an arbitrarily short observation of the process).

We interpret the drift as coming from the activity of other institutional traders, who have made trade decisions before the market opens, and who expect to execute these trades throughout the day. If these decisions are in the aggregate weighted to buys, then this will cause positive price pressure and an upwards drift; conversely for overall net selling. We do not know the net direction of these trades but we can infer it by observing prices. We implicitly assume that these traders are using VWAP-like strategies rather than arrival price, so that their trading is not "front-loaded." This assumption is questionable; if the strategies are front-loaded then the drift coefficient would vary through the day.

Thus we assume that the drift α is constant throughout the day, but we do not know its value. At the beginning of the day, we have a prior belief

$$\alpha \sim \mathcal{N}(\bar{\alpha}, \nu^2) \qquad \text{prior belief},$$

which will be updated using price observations during the day. There are thus two sources of randomness in the problem: the continuous Brownian motion representing the uninformed traders, and the single drift coefficient representing the constant trading of the large traders.

2.1 Bayesian inference

Intuitively, as the trader observes prices from the beginning of the day onwards, he or she starts to get a feeling for the day's overall flow.

Mathematically, at time t we know the stock price trajectory $S(\tau)$ for $0 \le \tau \le t$. In fact, all of our information about the drift comes from the final value $S(t)$. Conditional on the value of α, the distribution of $S(t)$ is

$$S(t) - S_0 \sim \mathcal{N}(\,\alpha t,\, \sigma^2 t\,) \quad \text{conditional on } \alpha.$$

and after some calculation we find the unconditional distribution

$$S(t) - S_0 \sim \mathcal{N}(\,\bar\alpha t,\, (\sigma^2 + \nu^2 t)t\,) \quad \text{unconditional.}$$

We then use Bayes' rule

$$\mathrm{Prob}(\,\alpha \mid S(t)\,) = \frac{\mathrm{Prob}(\,S(t) \mid \alpha\,) \cdot \mathrm{Prob}(\alpha)}{\mathrm{Prob}(S(t))}$$

to obtain the posterior conditional distribution

$$\alpha \sim \mathcal{N}\left(\frac{\bar\alpha\sigma^2 + \nu^2(S(t) - S_0)}{\sigma^2 + \nu^2 t},\, \frac{\sigma^2}{\sigma^2 + \nu^2 t}\nu^2 \right) \quad \text{conditional on } S(t). \tag{2}$$

This represents our best estimate of the true drift α, as well as our uncertainty in this estimate, based on combination of our prior belief with price information observed to time t.

This formulation accomodates a wide variety of belief structures. If we believe our initial information is perfect, then we set $\nu = 0$ and our updated belief is always just the prior $\alpha = \bar\alpha$ with no updating. If we believe we have no reliable prior information, then we take $\nu^2 \to \infty$ and our estimate is $\alpha \sim \mathcal{N}\big((S(t) - S_0)/t, \sigma^2/t\big)$, coming entirely from the intraday observations. For $t = 0$, we will have $S(0) = S_0$, and our belief is just our prior. As $t \to \infty$, our estimate becomes $\alpha \sim \mathcal{N}\big((S - S_0)/t, 0\big)$: we have accumulated so much information that our prior belief becomes irrelevant.

2.2 Trading and price impact

The trader has an order of X shares, which begins at time $t = 0$ and must be completed by time $t = T < \infty$. For concreteness, we shall suppose $X > 0$ and interpret this as a buy order.

A *trading trajectory* is a function $x(t)$ with $x(0) = X$ and $x(T) = 0$, representing the number of shares remaining to buy at time t. The corresponding *trading rate* is $v(t) = -dx/dt$. We shall require that $v(t) \ge 0$ for all t, so that the program never sells as part of a buy order. Together with the endpoint constraints, this requires $0 \le x(t) \le X$, but it may also be binding in the interior of this region.

We use a linear market impact function for simplicity, although empirical work (Almgren, Thum, Hauptmann, and Li 2005) suggests a concave function. Thus the actual execution price is

$$\tilde{S}(t) = S(t) + \eta\, v(t)$$

where $\eta > 0$ is the coefficient of temporary market impact.

The *implementation shortfall* C is the total cost of executing the buy program relative to the initial value,

$$\begin{aligned} C &= \int_0^T \tilde{S}(t)\, v(t)\, dt - X S_0 \\ &= \sigma \int_0^T x(t)\, dB(t) + \eta \int_0^T v(t)^2\, dt + \alpha \int_0^T x(t)\, dt. \end{aligned} \tag{3}$$

Here α is the true drift, that determines our cost whether or not we know its value. C is a random variable, both because the price $S(t)$ is random, and because the trading strategy $v(t)$ may be adapted to S.

3 Optimal Trading Strategies

We now address the question of what trading strategies are optimal, given the above model for price evolution and market impact. In "classic" arrival price (Almgren and Chriss 2000), trajectories are determined as a tradeoff between market impact and aversion to risk caused by volatility. The trader wants to complete the trade quickly to eliminate exposure to price volatility; he or she wants to trade slowly to minimize the costs of market impact. The optimal trajectory is determined as a balance between these two effects, parameterized by a coefficient of risk aversion.

Risk-averse trading strategies can behave strangely in time (Almgren and Lorenz 2006) even in the classic framework, depending on the precise formulation of the mean-variance tradeoff. In this case, the problem is complicated by the need to account for the variance in the estimation of α. We have obtained partial solutions for the risk-averse problem but their complexity obscures the underlying structure.

To focus on the drift, which is the most important new aspect of this problem, here we neglect risk aversion; we seek to minimize only the expectation of trading cost. That is, we assume that the pressure to complete the trade rapidly comes primarily from a desire to capture the price motion expressed by the drift α, and it is this effect that must be balanced against the desire to reduce impact costs by trading slowly.

To support this description, we shall generally suppose that the original buy decision was made because the trader's belief has $\bar{\alpha} > 0$. We then expect $\alpha > 0$ in (3), and the term $\int \alpha\, x(t)\, dt$ is a positive cost. It may be that the true value has $\alpha < 0$, or that intermediate price movements cause us to form a negative estimate. Because our point of view is that of a broker/dealer executing an agency trade, we shall always require that the trade be completed by $t = T$, unless the instructions are modified.

For any deterministic trajectory $x(t)$ specified at $t = 0$, C is a Gaussian variable. Conditional on the true value of α, it has expected value

$$\mathbb{E}(C) \;=\; \eta \int_0^T v(t)^2\, dt + \alpha \int_0^T x(t)\, dt \quad \text{conditional on } \alpha. \tag{4}$$

From (2), our best estimate at time t for the value of α is

$$\alpha_*(t, S) \;=\; \frac{\bar{\alpha}\sigma^2 + \nu^2 (S - S_0)}{\sigma^2 + \nu^2 t} \tag{5}$$

where $S = S(t)$. Because the expectation (4) is linear in α, we may simply substitute the expected value α_* to see that, conditional on the information available at time t, the expected cost of the remaining program is

$$E\big(t, x(t), S, \{x(\tau)\}\big) \;=\; \eta \int_t^T v(\tau)^2\, d\tau + \alpha_*(t, S) \int_t^T x(\tau)\, d\tau.$$

On the left, t is current time, $x(t)$ is the number of shares currently remaining to buy, S is the current price, and $\{x(\tau)\}$ denotes the liquidation strategy that will be used on the remaining time $t \le \tau \le T$.

Our trading goal is to choose the remaining strategy to minimize this expected cost: determine $x(\tau)$ for $t \le \tau \le T$ so that

$$\min_{\{x(\tau)\}} E\big(t, x(t), S, \{x(\tau)\}\big). \tag{6}$$

In computing this solution, we shall assume that the drift estimate $\alpha_*(t,S)$ does not change during the interval $t \le \tau \le T$. In fact, it will change as new price information is observed. Our actual strategy will use only the initial instantaneous trade rate of this trajectory, continuously responding to price information. This is equivalent to following the strategy only for a very small time interval Δt, then recomputing. Thus our strategy is highly dynamic.

We shall argue that this strategy determined is the true optimum strategy that would be computed by a full dynamic optimization. Loosely speaking, this will be because the expected value of future updates is zero, and thus they do not change the strategy of a risk-neutral trader.

3.1 Trajectories by calculus of variations

We consider a small perturbation of the path $x(\tau) \mapsto x(\tau) + \delta x(\tau)$ for $t \le \tau \le T$. Since $x(\tau)$ is fixed at $\tau = t$ and $\tau = T$, this perturbation must have $\delta x(t) = \delta x(T) = 0$. The associated trade rate perturbation is $\delta v(\tau) = -\delta x'(\tau)$, and the perturbation in cost (assuming that $x(\tau)$ and $\delta x(\tau)$ are twice differentiable) is

$$\begin{aligned} \delta E &= \eta \int_t^T 2v(\tau)\,\delta v(\tau)\,d\tau \;+\; \alpha_* \int_t^T \delta x(\tau)\,d\tau \\ &= \int_t^T \big(-2\eta\, x''(\tau) \;+\; \alpha_* \big)\,\delta x(\tau)\,d\tau. \end{aligned} \tag{7}$$

Here $\alpha_* = \alpha_*(t, S(t))$ is the best available drift estimate using information at time t, which we assume constant for $t \le \tau \le T$. If $x(\tau)$ is an optimal solution, then there must not exist any admissible $\delta x(\tau)$ that gives $\delta E < 0$.

Unconstrained trajectories We temporarily neglect the sign constraint on $x'(\tau)$. Then $\delta x(\tau)$ may have either positive or negative values independently for each τ, and any optimizing $x(\tau)$ must satisfy the ordinary differential equation (ODE)

$$x''(\tau) \;=\; \frac{\alpha_*}{2\eta}, \qquad t \le \tau \le T. \tag{8}$$

The solution to this equation that satisfies the boundary conditions is

$$x(\tau) = \frac{T-\tau}{T-t}\,x(t) - \frac{\alpha_*}{4\eta}(\tau-t)(T-\tau), \qquad t \le \tau \le T, \tag{9}$$

and the corresponding instantaneous trade rate is

$$v(t,x) = -x'(\tau)\big|_{\tau=t} = \frac{x(t)}{T-t} + \frac{\alpha_*}{4\eta}(T-t) \tag{10}$$

as a function of time and shares remaining. This solution may violate the constraints: if α_* is large then the quadratic term in (9) may cause $x(\tau)$ to dip below zero, which would cause $v(t)$ in (10) to become negative.

This unconstrained solution is the sum of two pieces. The first piece is proportional to $x(t)$ and represents the linear (VWAP) liquidation of the current position; it is the optimal strategy to reduce expected impact costs with no risk aversion. The second piece is independent of $x(t)$ and would therefore exist even if the trader had no initial position. Just as in the solutions of Bertsimas and Lo (1998), this second piece is effectively a proprietary trading strategy superimposed on the liquidation. The magnitude of this strategy, and hence the possible gains, are determined by the ratio between the expected drift and the liquidity coefficient. Imposition of the constraint will couple these pieces together.

Constrained trajectories If the constraint becomes binding, then it is no longer obvious that the integration by parts procedure used to derive (7) is valid. For example, if a trajectory that crosses the axis $x = 0$ is simply clipped to satisfy $x \ge 0$, then the derivative will be discontinuous. A more refined use of the calculus of variations gives the additional condition

$v(\tau)$ must be continuous (though not necessarily differentiable).

Thus when solutions meet the constraint, they must do so smoothly. Solutions are obtained by combining the ODE (8) in regions of smoothness, with this "smooth pasting" condition at the boundary points.

The result may be summarized as follows. There is a critical drift value α_c such that

- If $|\alpha_*| \le \alpha_c$, then the constraint is not binding. The solution is the one given in (9) and (10).

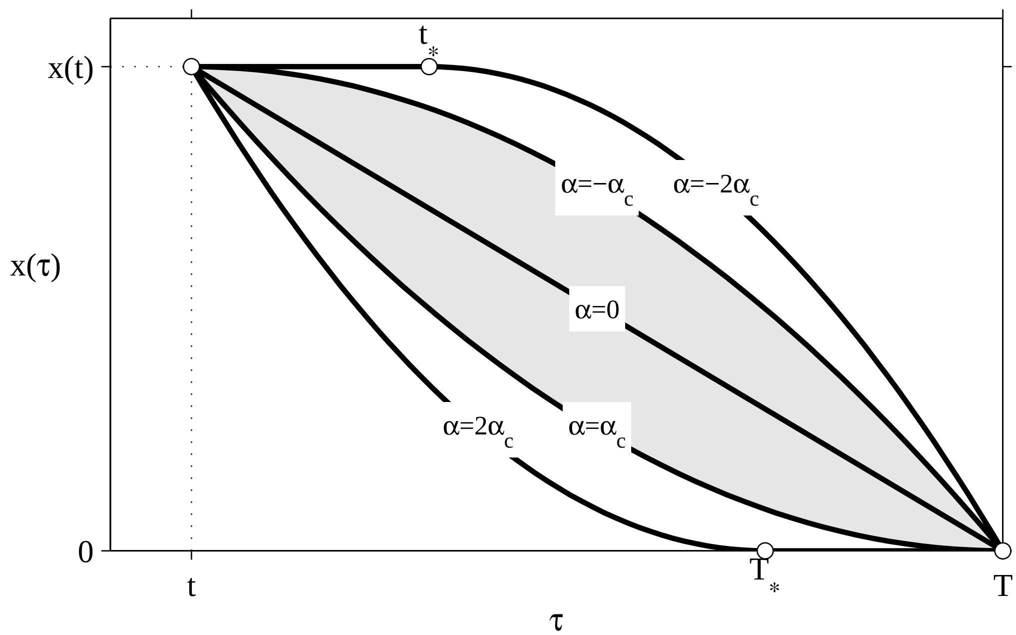

Figure 1: Constrained solutions $x(\tau)$, starting at time t with shares $x(t)$ and drift estimate α. For $\alpha > 0$, the trajectories go below the linear profile to reduce expected purchase cost. For $|\alpha| \le \alpha_c$, the constraint is not binding (shaded region). At $\alpha = \alpha_c$ the solutions become tangent to the line $x = 0$ at $\tau = T$, and for larger values they hit $x = 0$ with zero slope at $\tau = T_* < T$. For $\alpha < -\alpha_c$, trading does not begin until $\tau = t_* > t$.

- If $\alpha_* > \alpha_c$, then the solution is the one of (9,10), with a shortened end time $T_* < T$ determined by

$$T_* - t \;=\; \sqrt{\frac{4\eta\, x(t)}{\alpha_*}}.$$

  This value is determined so that $x'(T_*) = x(T_*) = 0$. The threshhold value α_c is the value of α_* for which $T_* = T$:

$$\alpha_c\big(\, x(t),\, T - t\,\big) \;=\; \frac{4\eta\, x(t)}{(T-t)^2}.$$

- If $\alpha_* < -\alpha_c$, then the solution is the one of (9,10), except that trading does not begin until a starting time $t_* > t$ determined by

$$T - t_* = \sqrt{\frac{4\eta\, x(t)}{-\alpha_*}}.$$

This value is determined so that $x'(t_*) = 0$ and $x(t_*) = x(t)$. The threshhold value α_c is the value of $-\alpha_*$ for which $t_* = t$.

Figure 1 illustrates these solutions.

Then the overall trade rate formula may be summarized as

$$v(t,x,S) = \begin{cases} 0, & \alpha_* < -\alpha_c \\ \dfrac{x}{T-t} + \dfrac{\alpha_*}{4\eta}(T-t), & |\alpha_*| < \alpha_c \\ \dfrac{x}{T_*-t} + \dfrac{\alpha_*}{4\eta}(T_*-t) = \sqrt{\dfrac{\alpha_* x}{\eta}}, & \alpha_* > \alpha_c \end{cases} \qquad (11)$$

where $\alpha_* = \alpha_*(t, S(t))$ by (5). This is our Bayesian adaptive strategy: it is a specific formula for the instantaneous trade rate as a function of price, time, and shares remaning.

Since $dx = -v\,dt$, this gives an ordinary differential equation for the trajectory $x(t)$, with a stochastic element due to the presence of $S(t)$. (It is not a stochastic differential equation since dB appears only in dS, not in dx. Thus $x(t)$ will have a first time derivative, but not a second derivative.)

3.2 Examples

Figures 2 and 3 show examples of the strategies computed by this method. To produce these pictures, we began with a prior belief for α having mean $\bar{\alpha} = 0.7$ and standard deviation $\nu = 1$. For each trajectory, we generated a random value of α from this distribution, and then generated a price path $S(t)$ for $0 \le t \le 1$ with volatility $\sigma = 1.5$. For example, on a stock whose price is \$100 per share, these would correspond to 1.5% daily volatility, and an initial drift estimate of +70 basis points with a substantial degree of uncertainty.

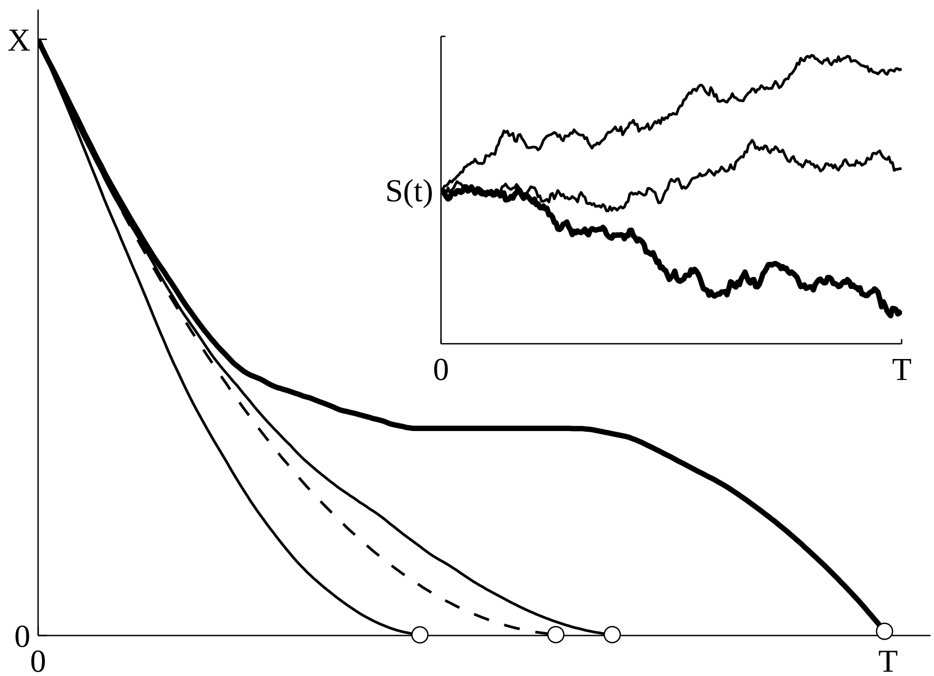

Figure 2: Adaptive trajectories. The inset shows price processes, and the main figure shows trade trajectories $x(t)$. The dashed line is the static trajectory using the prior belief for the drift value. In this example, we have selected realizations with very high drift to highlight the solution behavior, including temporary interruption of trading when the drift estimate becomes more negative than the critical value (bold line).

We set the impact coefficient $\eta = 0.07$ and the initial shares $X = 1$, meaning that liquidating the holdings using VWAP across one day will incur realized price impact of 7 basis points.

Then for each sample path, we evaluate the share holdings $x(t)$ using the Bayesian update strategy (11) and plot the trajectories. For comparison, we also show the optimal static trajectory using only the initial estimate of the drift.

In Figure 2, to illustrate the features of the solution, we show a rather extreme collection of paths, having very high realized drifts. In Figure 3 we show a completely representative selection.

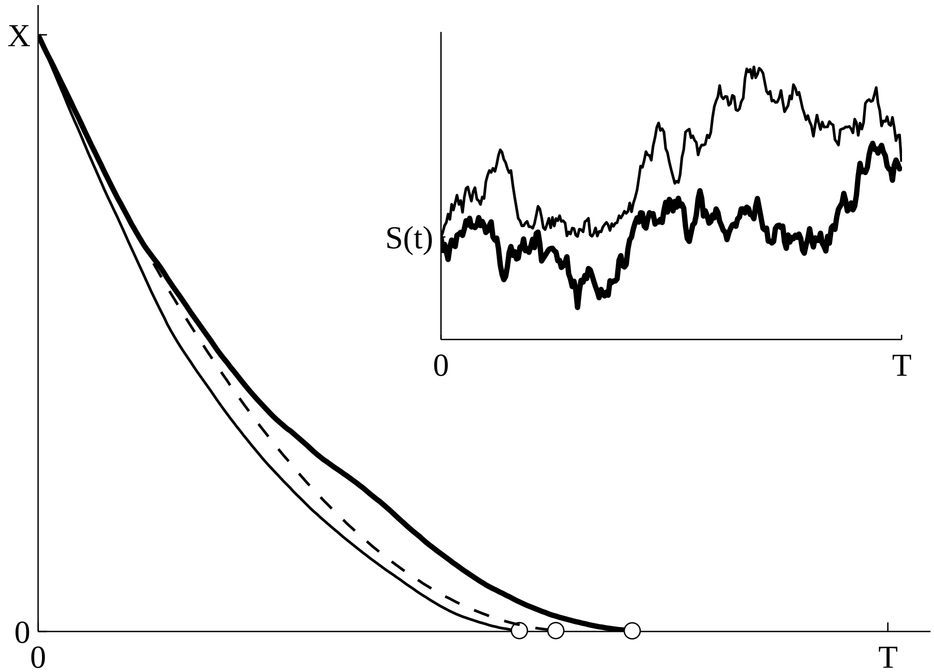

Figure 3: Same layout as Figure 2, but with more realistic sample paths. In both, the price initially trends downward (more strongly for the light path), causing the trader to estimate a drift that is smaller than his prior belief and to slow down his trading relative to the static solution.

3.3 Optimality of the Bayesian adaptive strategy

The Bayes adaptive strategy (11) is “locally” optimal in the sense that at any intermediate time we use all the new information available and recompute the trajectory for the remainder as though we would use the same estimate until the end of trading. Since we do expect to update our estimate, it is not obvious that this is the true optimal strategy.

Using the methods of stochastic optimal control we can formulate a Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE) for the value function of the corresponding dynamic program (see Appendix A).

For the unconstrained case, we can solve this PDE analytically. The corresponding optimal strategy is computed by differentiating the value function, and this agrees precisely with (10). This computation verifies that the local solution is the dynamic optimal solution, for the unconstrained case.

Furthermore, we may explicitly determine the gains due to adaptivity. The value function for the dynamic strategy, at $t = 0$ with initial shares X, may be written

$$E_{\text{dyn}} = E_{\text{stat}} - \Delta$$

where

$$E_{\text{stat}} = \frac{X^2\eta}{T} + \frac{X\bar{\alpha}T}{2} - \frac{T^3\bar{\alpha}^2}{48\eta}$$

is the expected cost of the non-adaptive strategy determined at $t = 0$ using the prior expected drift $\bar{\alpha}$. The additional term,

$$\Delta = \frac{\sigma^2 T^2}{48\eta}\int_0^1 \frac{(1-\delta)^3}{(\delta+\rho)^2}\,d\delta, \qquad \rho = \frac{\sigma^2}{\nu^2 T}, \quad \delta = \frac{t}{T},$$

is the reduction in expected cost obtained by using the Bayesian adaptive strategy (note that $\Delta > 0$).

The gain Δ is independent of initial portfolio size X and thus, as discussed above, it represents the gains from a proprietary trading strategy superimposed on the risk-neutral liquidation profile. It can be seen that $\Delta \sim \mathcal{O}(T^4)$ when T is small and $\mathcal{O}(T^2)$ when T is large, so the adaptivity adds very little value if applied to short-term correlation. This accounts for the small gains obtained by Bertsimas and Lo (1998), as discussed by Almgren and Chriss (2000).

For the constrained case, the HJB equation has complicated boundary conditions, and we are unable to determine an explicit analytic solution. However, we do not believe that imposition of the constraint should change the relation between the static and the dynamic solutions. Thus we believe that our Bayesian dynamic solution is the dynamic optimal solution in the constrained case as well.

4 Conclusion

We have presented a simple model for momentum in price motion based on daily trading cycles, and derived optimal risk-neutral adaptive trading strategies. The momentum is understood to arise from the correlated trade targets being executed by large institutional investors. The trader begins with a belief about the direction of this imbalance, and expresses a level of confidence in this belief that may range anywhere from perfect knowledge to no knowledge. This belief is then updated using observations of the price process during trading. Under the assumptions of the model, our solutions deliver better performance than non-adaptive strategies.

It is natural to ask whether this model can be justified by empirical data, but we would like to highlight some of the difficulties of doing such a study. In our model, the random daily drift is superimposed on the price volatility caused by small random traders. In theory, these two sources of randomness can be disentangled by measuring volatility on an intraday time scale and comparing it to daily volatility. If daily volatility is higher than intraday, then the difference can be attributed to serial correlation of the type considered here.

In practice, because real price processes are far from Gaussian, it is difficult to do this comparison with any degree of reliability, even if one restricts attention to days when there is large institutional flow.

We therefore justify our model not by empirical study, but by the practical observation that some fraction of traders do express interest in using strategies similar to the ones described here. Our model provides a conceptual framework for designing optimal strategies that capture this preference; without any such framework it is impossible to design algorithms except by completely ad hoc methods.

A Stochastic Optimal Control

Now we support our claim that the trade velocity (10) of the Bayesian adaptive strategy is in fact the optimal strategy for (6). For that, we will formulate the problem in a full dynamic programming framework.

The control, the state variables, and the stochastic differential equations of problem (6) are given by

$$
\begin{aligned}
v &= \text{rate of buying} & & \\
x &= \text{shares remaining to buy} & dx &= -v\,dt \\
y &= \text{dollars spent so far} & dy &= (s + \eta v)\,v\,dt \\
s &= \text{stock price} & ds &= \alpha\,dt + \sigma\,dB
\end{aligned}
$$

where $\alpha \sim N(\bar{\alpha}, \nu^2)$, chosen randomly at $t = 0$. We begin at $t = 0$ with shares $x(0) = X$, cash $y(0) = 0$, and initial stock price $s(0) = S$. The strategy $v(t)$ must be adapted to the filtration of B, and must satisfiy $x(T) = 0$. We focus on the unconstrained case and don't require $0 \le x(t) \le X$.

We want to find a control function $v(t)$ to minimize the final amount of dollars spent,

$$\min_{v(\tau)\ \text{s.t.}\ x(T)=0} \mathbb{E}[y(T)].$$

This is a common problem in stochastic optimal control, and solved by dynamic programming. Standard techniques lead to the Hamilton-Jacobi-Bellman (HJB) partial differential equation

$$0 = u_t + \frac{1}{2}\sigma^2 u_{ss} + \alpha_* u_s + \min_v \Big((s\, u_y - u_x)\, v + \eta u_y v^2 \Big)$$

for the value function

$$u(t,x,y,s) = \min_{v(\tau),\ t\le\tau\le T,\ \text{s.t.}\ x(T)=0} \mathbb{E}[y(T)].$$

$\alpha_* = \alpha_*(t,s)$ denotes the estimate of α at time t as computed in (2). The optimal trade velocity is found as

$$v_*(t,x,y,s) = \frac{u_x - s\,u_y}{2\eta\, u_y} \tag{12}$$

and we have the final HJB partial differential equation

$$0 = u_t + \frac{1}{2}\sigma^2 u_{ss} + \alpha_* u_s - \frac{(s\,u_y - u_x)^2}{4\eta\, u_y} \tag{13}$$

for $u(t,x,y,s)$ together with the boundary condition

$$u(T,0,y,s) = y \qquad \text{for all } y, s. \tag{14}$$

It is straightforward to check that

$$\begin{aligned} u(t,x,y,s) = {} & y + xs + \frac{x^2\eta}{T-t} + \frac{x\,\alpha_*(t,s)\,(T-t)}{2} \\ & - \frac{(T-t)^3\alpha_*(t,s)^2}{48\eta} - \int_t^T \frac{\sigma^2\nu^4(T-\tau)^3}{48\eta(\tau\nu^2+\sigma^2)^2}\,d\tau \end{aligned} \tag{15}$$

satisfies the PDE (13) and the boundary conditions (14). Moreover, the corresponding optimal trade velocity (12) reads

$$v_*(t,x,y,s) = \frac{x(t,x,y,s)}{T-t} + \frac{\alpha_*(t,s)\cdot(T-t)}{4\eta}$$

which is exactly the trade velocity (10). That is, the Bayesian adaptive strategy is in fact the optimal strategy for the optimization problem (6).

For the constrained case, the optimal velocity (12) becomes

$$v_*(t,x,y,s) = \max\left\{ \frac{u_x - s\,u_y}{2\eta\,u_y},\ 0 \right\}.$$

This makes the corresponding PDE even more highly nonlinear, and we do not know how to derive explicit solutions.

References

Almgren, R. and N. Chriss (2000). Optimal execution of portfolio transactions. *J. Risk 3*(2), 5–39. 6, 14

Almgren, R. and J. Lorenz (2006). Adaptive arrival price. Preprint. 2, 6

Almgren, R., C. Thum, E. Hauptmann, and H. Li (2005). Equity market impact. *Risk 18*(7, July), 57–62. 6

Bertsimas, D. and A. W. Lo (1998). Optimal control of execution costs. *J. Financial Markets 1*, 1–50. 3, 9, 14

Brunnermeier, M. K. and L. H. Pedersen (2005). Predatory trading. *J. Finance 60*(4), 1825–1863. 2

Carlin, B. I., M. S. Lobo, and S. Viswanathan (2005). Episodic liquidity crises: Cooperative and predatory trading. Preprint. 2

Kissell, R. and R. Malamut (2006). Algorithmic decision-making framework. *J. Trading 1*(1), 12–21. 2