Stochastic Control Theory and High Frequency Trading

March 1, 2010

Doug Borden
Electronic Trading Group
Knight Equity Markets
dborden@knight.com

Contents

Introduction

Knight

Who am I?

I work in the Electronic Trading Group (ETG) at Knight Capital. In my career I previously worked at Credit Suisse and at Goldman Sachs. I started my career in derivatives, and later switched to quantitative cash trading. At Knight I'm responsible for execution across all of our electronic strategies.

Who is Knight?

- Knight is a Broker-Dealer originally established to cater to the retail community
- Knight is the leading source of off-exchange liquidity in U.S. equities across all market segments
- Knight provides market making and agency-based trading in U.S., European and Asian equities, ADRs, ETFs, futures and options
- Knight is the largest U.S. market-maker, trading in more than 19,000 U.S. Equities
- Knight provides connectivity to more than 100 external market centers worldwide, including exchanges, ECNs, ATSs, dark pools, ATFs, MTFs and broker-dealers
- In 2009, Knight traded approximately 2.5 trillion shares and executed more than 980 million trades, an average of 600,000 trades per hour
- Knight is #1 in shares traded of Listed securities with 17.3% market share*
- Knight is #1 in shares traded on NASDAQ Capital Market, Global Market, and Global Select Market securities with a combined market share of 23.9%*
- Knight is #1 in shares traded of Bulletin Board and Pink Sheet securities with 86.5% market share*
- Knight is #1 in shares traded of S&P 500 securities with a 15.1% market share*
- Knight also trades in Fixed Income, Currencies & Commodities, and offers Capital Markets and Investment Banking services as well.

*Autex, 2009

High Frequency Trading Decisions

	Long-dated Trading	High Frequency Trading
Holding Period	Weeks to Years	Seconds to Minutes
Expected Returns	1% -- 10% per position >> typical bid-offer spread	.01% -- .10% per position ~ typical bid-offer spread
Trading Style	Liquidity Taking	Liquidity Providing

I'm going to be discussing high frequency trading in the context of providing liquidity to the markets – i.e. voluntarily placing bids and offers on a large number of stocks into public and semi-public trading venues (exchanges, ECN's, dark pools, etc.).

When engaging in such activity, the high frequency trader hopes to get paid the bid-ask spread in return for suffering negative selection (i.e. he or she is more likely to be a buyer in a falling market and a seller in a rising market).

Generally the difference between profit and loss is very thin, and making real-time decisions on size and price is a difficult and delicate affair.

Parametric Approach

Imagine that for every stock in my universe, at all times, I have a short-dated return forecast for the mid-market price. The return forecast may be zero, positive or negative, and will typically be smaller than or similar in size to the stock's half-spread.

Here is a simple rule for a high frequency trading strategy:

1. Size: 100 shares to buy; 100 shares to sell
2. Price: BuyPrice = Mid-Market + A x Forecast – B x MCR – Current half-spread

   SellPrice = Mid-Market + A x Forecast – B x MCR + Current half-spread

where MCR is the stock's marginal contribution to my portfolio risk.

I then have two parameters, A and B, which I determine by backtesting my strategy.

Generally this is a sound approach, albeit in practice the models will have many, *many* more parameters than the two I've specified here.

This approach is sometimes referred to as Global Parametric Optimization.

Parametric Approach (cont.)

What are the strengths and weaknesses of the parametric approach?

	Strength	Weakness
Rules are intuitive	✓	
Optimization tends to be stable	✓	
Edge cases handled poorly		✓
Rule-sets can grow to be very complex		✓
Not adaptive to changes in market structure		✓

Could Stochastic Control Theory offer a better approach?

Not a new idea:

“Applications of Mathematical Control Theory to Finance: Modeling Simple Dynamic Cash Balance Problems”, S. Sethi and G. L. Thompson (1970), *Journal of Financial and Quantitative Analysis*, **5**, pp. 381-394.

“Optimal Investment Models and Risk-Sensitive Stochastic Control”, W. H. Fleming (1995), *IMA Volume of Mathematical Finance*, **65**, pp. 35-45.

“An Application of Stochastic Control Theory to Financial Economics”, W. H. Fleming and T. Pang (2003), *SIAM Journal on Control and Optimization*, **43**, pp. 502-531.

… and many more.

Remember back to Freshman Calculus, and the problem of the guy and the rowboat:

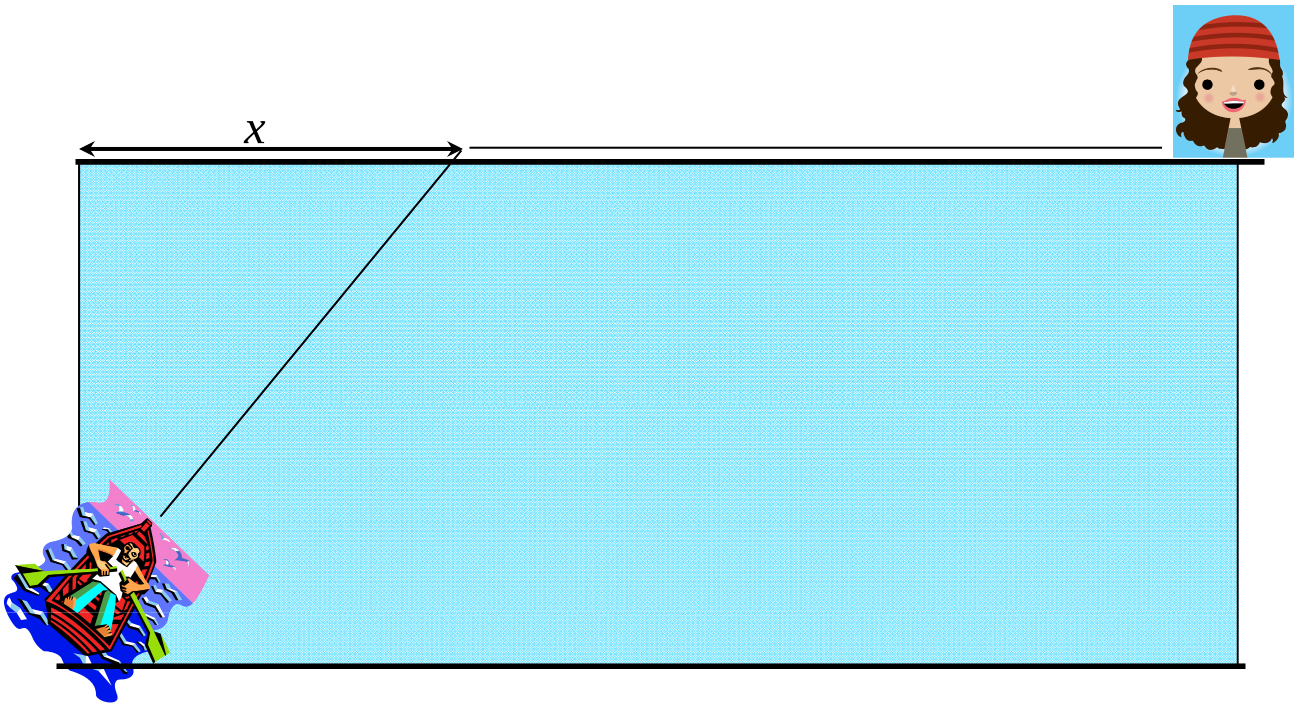

Jack wants to get to Jill as fast as he can, who is 2 miles down-river from him, but he needs to cross the ½-mile wide river to get to her. He can row at 3mph, but once across he can run at 9mph along the bank. Toward which point along the bank should Jack row so as to reach Jill in the smallest amount of time *“T”*?

$$T = \sqrt{\tfrac{1}{4} + x^2} \,/\, 3 + (2 - x) \,/\, 9 \qquad dT/dx|_{x=x_o} = 0 \Rightarrow x_o = 0.1767 \text{ miles}$$

We say that the optimal *point “x_o”* is a solution to the algebraic equation *dTime/dx = 0*.

What if there is a variable current with a velocity *v(y)* [and the shore is too rocky to run along]?

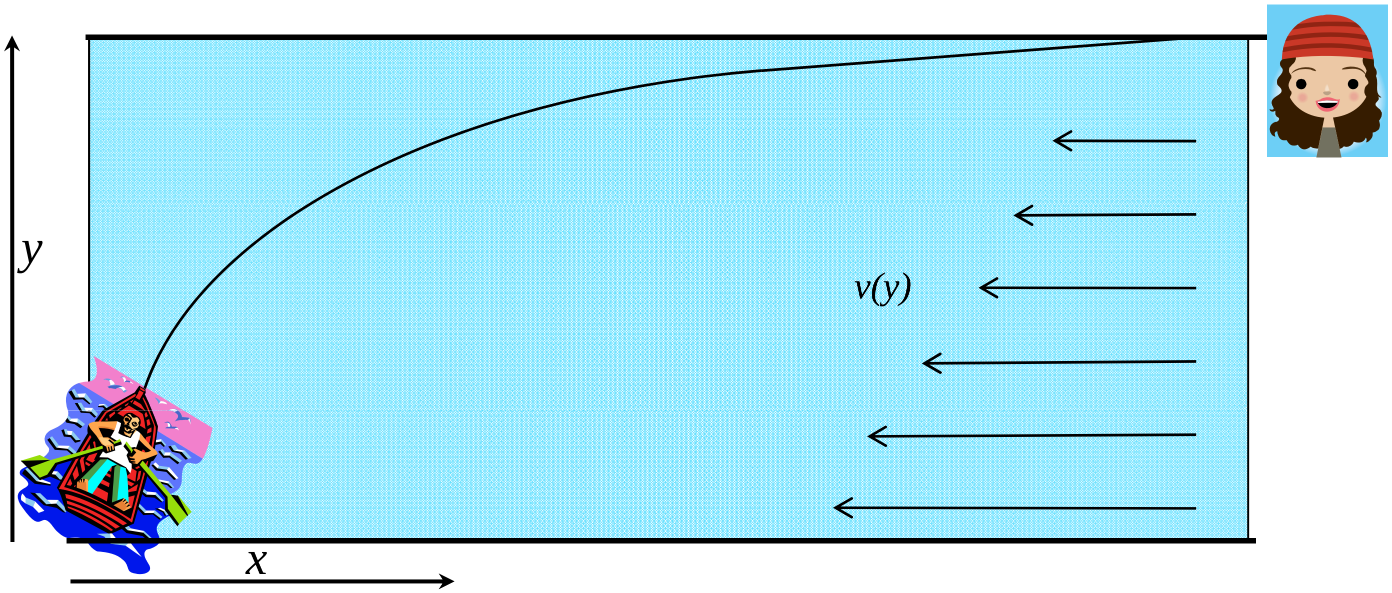

Now we need to find the optimal ***path*** $y_o(x)$ for Jack to row. For any path $y(x)$, the time it takes Jack to reach Jill is given by:

$$T = \int_0^2 \sqrt{\frac{(1+y'^2)^{3/2}}{\left(9+v(y)^2\right)\sqrt{1+y'^2}-6v(y)}}dx \equiv \int_0^2 L(y,y')dx$$

and the *optimal path* $y_o(x)$ must satisfy the Euler-Lagrange Equation:

$$\left.\frac{d}{dx}\frac{\partial L}{\partial y'}-\frac{\partial L}{\partial y}\right|_{y(x)=y_o(x)}=0$$

In this case we say the optimal path $y_o(x)$ is a solution to the differential equation $\frac{d}{dx}\frac{\partial L}{\partial y'}-\frac{\partial L}{\partial y}=0$

Now let's add turbulence to the river, so that as he rows across, Jack's rowboat is tossed about by the rapids.

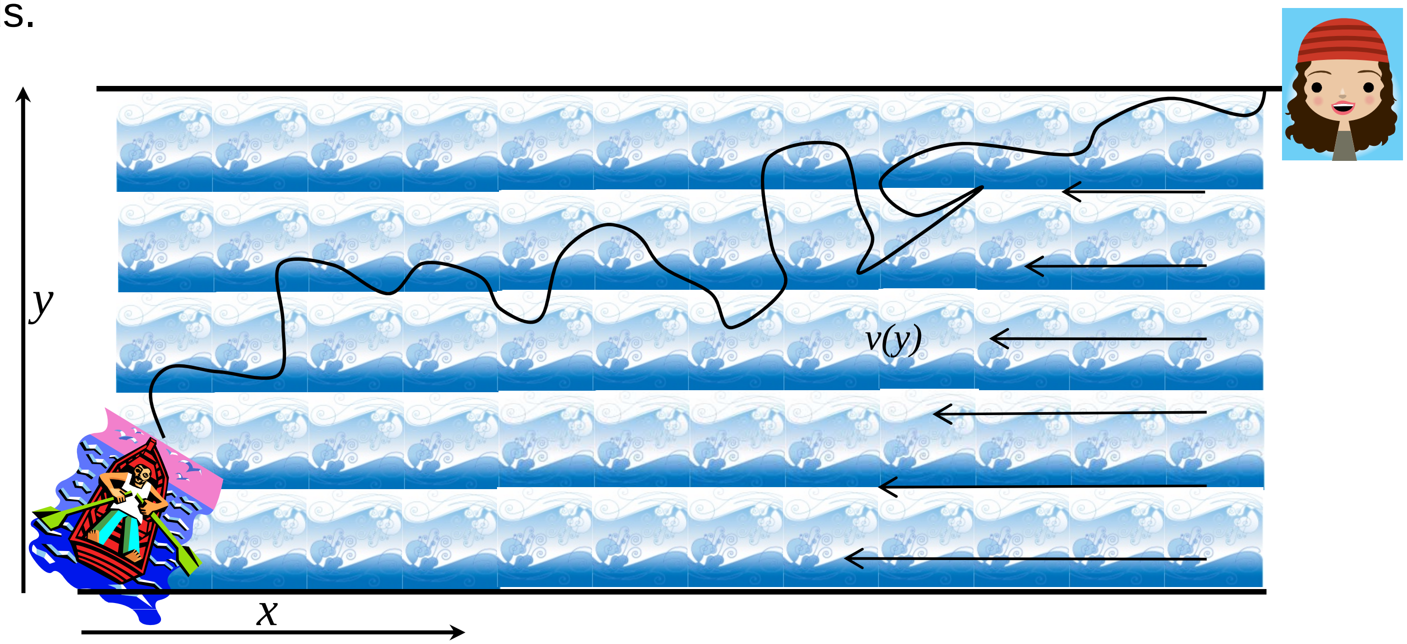

How do we help Jack get to Jill in this situation? Jack's path to Jill is no longer deterministic – though his boat's motion is partially determined by the direction in which Jack rows, there is also a random component to the motion due to the turbulence. In a case such as this we need to provide Jack with an ***optimal policy*** which he can follow – instructions for Jack which will optimize his path to Jill no matter where in the river he finds himself. This policy will take as inputs Jack's current position and velocity, and will output the optimal direction that Jack should row. Finding optimal *policies* is the job of Stochastic Control Theory.

Review of Variational Calculus (cont.)

In general, finding an optimal policy requires three specifications:

1. System dynamics for the state variables
2. A quantity to minimize or maximize
3. Relationships between “controls” and system state variables; these often appear as constraints.

“Controls”, as their name implies, are variables under the control of the system operator (e.g. Jack) while the state variables attain their values as a consequence of the controls and the system dynamics. In helping Jack get to Jill, we might have the following:

1. $dx = (v_o \cos\theta - v(y))\, dt + \sigma_x\, dz_x$

   $dy = v_o \sin\theta\, dt + \sigma_y\, dz_y$

2. $T = \int dt$ *we want to minimize T*

3. $2 = \int dx$ *we need to travel two miles upriver*

   $\frac{1}{2} = \int dy$ *we need to cross the ½-mile wide river*

where θ is the angle with respect to the shoreline indicating the direction Jack is rowing.

More generally, with $\boldsymbol{x}$ a vector of state variables, $\boldsymbol{u}$ a vector of controls, and R the quantity to minimize or maximize:

$$d\boldsymbol{x} = \boldsymbol{F}[\boldsymbol{x}(t), \boldsymbol{u}(\boldsymbol{x}, t), t]dt + \boldsymbol{G}[\boldsymbol{x}(t), \boldsymbol{u}(\boldsymbol{x}, t), t]dz$$

$$R[\boldsymbol{u}] = \int_0^T L[\boldsymbol{x}(t), \boldsymbol{u}(\boldsymbol{x}, t), t]dt + B[\boldsymbol{x}(T), T]$$

We have the following system of equations for $\boldsymbol{u}^*$, the optimal controls:

$$\nabla_{\boldsymbol{u}}\left[\nabla_{\boldsymbol{x}} H \cdot \boldsymbol{F}(\boldsymbol{x}, \boldsymbol{u}, t) + \tfrac{1}{2}\boldsymbol{G}^T(\boldsymbol{x}, \boldsymbol{u}, t) \cdot \nabla_{\boldsymbol{x}}^2 H \cdot \boldsymbol{G}(\boldsymbol{x}, \boldsymbol{u}, t) + L(\boldsymbol{x}, \boldsymbol{u}, t)\right]_{\boldsymbol{u}=\boldsymbol{u}^*} = 0$$

Where *"H"* is called the Value Function and satisfies the differential equation:

$$\frac{\partial H}{\partial t} + \nabla_{\boldsymbol{x}} H \cdot \boldsymbol{F}(\boldsymbol{x}, \boldsymbol{u}^*, t) + \tfrac{1}{2}\boldsymbol{G}^T(\boldsymbol{x}, \boldsymbol{u}^*, t) \cdot \nabla_{\boldsymbol{x}}^2 H \cdot \boldsymbol{G}(\boldsymbol{x}, \boldsymbol{u}^*, t) + L(\boldsymbol{x}, \boldsymbol{u}^*, t) = 0$$

with boundary condition: $H(\boldsymbol{x}, T) = B(\boldsymbol{x}, T)$

The above equations are known as the **Hamilton-Jacobi-Bellman Equations**.

Let's consider the following simplified model for Stock ABC:

- The mid-market price of ABC follows algebraic Brownian motion ($\mu(t)$ is presumed *known*).

$$dS = \mu(t)dt + \sigma dz$$

- The spread of ABC is constant with a width 2δ.
- Market participants can place passive **limit** orders to buy at the bid or to sell at the ask, or they may place active **market** orders to buy at the ask or sell on the bid. Note that limit orders "earn" spread but market orders "pay" spread.
- Limit orders to buy and sell are filled at rates given by:

$$\frac{dv_s}{v_s} = \beta(v_o - v_s)dt + \varsigma dz$$

$$\frac{dv_b}{v_b} = \beta(v_o - v_b)dt - \varsigma dz$$

- Market orders are filled at a constant rate $f_o >> v_o$.
- We assume no market impact.

As market makers, what do we get to control? We can control four variables:

1. Whether we place a limit order to buy → Let's define this as $\lambda_b(t)$ which takes values of either 0 or 1
2. Whether we place a limit order to sell → Let's define this as $\lambda_s(t)$ which takes values of either 0 or 1
3. Whether we place a market order to buy → Let's define this as $m_b(t)$ which takes values of either 0 or 1
4. Whether we place a market order to sell → Let's define this as $m_s(t)$ which takes values of either 0 or 1

Let's denote $N(t)$ the number of shares we hold in inventory at time t. Then N(t) is given by:

$$dN = \left[-\lambda_s(t)v_s(t) + \lambda_b(t)v_b(t) - m_s(t)\,f_o + m_b(t)\,f_o\right]dt$$

What quantity do we want to extremize? We want to maximize our risk-adjusted PnL:

$$U = \int_0^T \left[\left(S(t)+\delta\right)\lambda_s(t)v_s(t) - \left(S(t)-\delta\right)\lambda_b(t)v_b(t) + \left(S(t)-\delta\right)m_s(t)\,f_o - \left(S(t)+\delta\right)m_b(t)\,f_o - \eta\sigma^2 S^2(t)N^2(t)\right]dt$$

with η our risk aversion coefficient.

So we have 4 state variables (S, v_s, v_b, N), each with an evolution equation, and 4 controls (λ_b, λ_s, m_b, m_s). Plugging into the HJB equations …

… gives a somewhat intimidating 2nd order partial differential equation for $H=H(S,v_s,v_b,N,t)$:

$$\frac{\partial H}{\partial t}+\mu(t)\frac{\partial H}{\partial S}+\beta(v_o-v_s)v_s\frac{\partial H}{\partial v_s}+\beta(v_o-v_b)v_b\frac{\partial H}{\partial v_b}-\left(\lambda_s v_s-\lambda_b v_b+m_s f_o-m_b f_o\right)$$
$$+\tfrac{1}{2}\sigma^2\frac{\partial^2 H}{\partial S^2}+\tfrac{1}{2}\zeta^2 v_s^2\frac{\partial^2 H}{\partial v_s^2}+\tfrac{1}{2}\zeta^2 v_b^2\frac{\partial^2 H}{\partial v_b^2}+\sigma\zeta v_s\frac{\partial^2 H}{\partial S\partial v_s}+\sigma\zeta v_b\frac{\partial^2 H}{\partial S\partial v_b}+\zeta^2 v_s v_b\frac{\partial^2 H}{\partial v_b\partial v_s}$$
$$+\left(S+\delta\right)\lambda_s v_s-\left(S-\delta\right)\lambda_b v_b+\left(S-\delta\right)m_s f_o-\left(S+\delta\right)m_b f_o-\eta\sigma^2 S^2 N^2=0$$

This PDE, though a bit scary to look at, is solvable using standard numerical techniques from derivative pricing, operations research, physics, etc.

So … let's assume we've solved our HJB equations using one of the excellent open-source PDE solvers available for download, and have used the optimality conditions shown earlier to get our optimal controls.

- *Is the calculation fast enough?* Probably not.
- *Can we speed it up?* Yes.

We've now found the optimal policies (our λ's and m's), and also the Value Function, H. We know what to do with our optimal policies – but **what good is the Value Function**?

Imagine you are happily using your optimal λ's and m's to make markets, when a salesperson calls asking you bid on a block trade. It is a large number of shares to buy (N_0), but you can execute it outside the NBBO at a price P. Should you take the trade, or pass?

<u>The Value Function serves as a decision engine for questions like this one.</u>

You take your current estimate of $\mu(t)$ and the current values of S, v_s, v_b, & N. Now evaluate the Value Function for two different starting conditions: N shares and zero dollars; or $N + N_0$ shares and $-N_0P$ dollars. Whichever is higher tells you whether to do the trade or not! Very valuable information.

Summary

- Stochastic Control Theory provides a rigorous framework for making decisions under conditions of uncertainty.
- High Frequency Trading decisions lend themselves to be cast into such a framework.
- Even the simplest of market models leads to very complicated differential equations
- BUT: Physics, Engineering, Operations Research, and two decades of derivatives pricing all provide a wealth of tools for solving the resulting HJB equations.
- Very interesting approach which is only now being explored.

If you are interested in exploring employment opportunities with ETG, please send your CV to Jessica Wang at jewang@knight.com 201-386-2897