Value Under Liquidation

Robert Almgren and Neil Chriss

October 5, 1999

1 Introduction

Liquidity is the readiness by which an asset can be converted into cash. Readiness describes the degree of transaction costs, time and uncertainty that must be borne in order to effect a transaction. Consider a portfolio held by a financial institution. The value of this portfolio to the institution will depend in certain situations on its liquidity. This may be the case, for example, if the firm needs to liquidate the portfolio to fund short term obligations coming due. In such a case, the true value of the portfolio will depend on the expectation of turning the portfolio into cash over a pre-specified period of time. This paper examines this situation under the guise of what we call *value under liquidation* and lays down a framework for determining the true value of a portfolio when it has to be liquidated over a fixed period of time.

Our main conclusion is that value under liquidation is best viewed as an expected value with an associated risk. The finite liquidity of markets makes it impossible to instantly liquidate a portfolio and therefore the value of any portfolio under liquidation is random with a given mean (expected value) and standard deviation. That is, since the value of a portfolio under liquidation will be realized over a finite trading horizon, the total value captured is necessarily uncertain. As we cannot know for sure ahead of time what value we will actually receive; the best we can do is determine the statistical average or expected value that liquidation will yield. This value, naturally, is uncertain, and this uncertainty represents a risk to any organization that intends to make use of funds derived from liquidation. What we will show is that the level of uncertainty, called the *revenue risk*, may be controlled by modulating the way the portfolio is traded. To decrease risk one simply has to trade more rapidly. This however has a downside: while increasing the speed of transacting decreases risk, it also increases transaction costs. Thus, there is an intrinsic tradeoff between revenue risk and the cost of trading. Decrease one and you naturally increase the other.

This paper is about the tradeoff between trading cost and trading risk and how it applies to the value of a portfolio under liquidation. In determining a portfolio's value under liquidation one must start by determining his or her degree of risk aversion. This risk aversion is simply an expression of the level of tolerance for discrepancies between the realized trading revenue and the actual trading revenue. This tolerance, then, should be an expression of a firm's policy for valuing portfolios. The message is clear: one can raise the value of a portfolio under liquidation simply by being willing to take on more risk. Once the risk tolerance is determined, our methods show there is a unqiue *optimal trading trajectory* associated to it. This trajectory has the property that it is *efficient,* namely that trading along it has the minimal possible trading cost for its level of market exposure. Given all of this, the value of a portfolio under liquidation is nothing more than the expected cost of trading along an optimal trajectory for a given level of risk aversion.

This paper describes the elements of optimal execution theory necessary to compute optimal trading trajectories and portfolio values under liquidation. In the next section we describe this theory in a non-technical manner. In the following section we describe some of the mathematics, and end with an example using several US equities. For a more detailed mathematical treatment we refer the reader to *Optimal Execution of Portfolio Transactions,* Almgren and Chriss (1999).

2 Non-Technical Introduction to Value Under Liquidation

In this section we give a non-technical overview of optimal liquidation, omitting mathematical details and focussing on the qualitative aspects of the theory. The aim of the theory of optimal execution is to compute the value of a portfolio under liquidation, supposing that trading is conducted from the present time to some fixed future time. To get at this, we need two pieces of information about the portfolio:

1. The variance/covariance structure of the portfolio
2. The transaction cost functions of the assets in the portfolio

Transaction costs functions express the relationship between trading rate and the premium or discount one must pay or concede to effect a trade. Put simply, the more rapidly one trades, the more demand this places on the market and these demands will only be met at a premium (for a purchase) or a concession (for a sale) to the current market price. The specific premium/concession depends on the extent of the demand and particular features of the security, these together are captured in the security's transaction cost function. Calculating transaction cost functions is a part of *market microstructure theory,* and has been studied extensively. It turns out to be very difficult to predict in detail the cost of a single trade. However, over a large number of trades, the prediction errors become small relative to costs so that well-specified trading models are useful tools for predicting aggregate trading costs.

In liquidating a basket of securities over a fixed period of time two things occur. First, for a particular choice of trading path there will be market impact determined by the rate of trading at each point along the path. The total cost of trading is the sum of the market impact at each point in time. Second, at each point along the trading path the unexecuted portion of the portfolio is exposed to subsequent market moves. Thus, the form of the trading path determines the total uncertainty in trading revenue. For example, trading rapidly in the first instant of trading decreases risk rapidly, because the reduction in market exposure is felt for the entire life of the trade. Naturally, of course, the decrease in risk comes at the cost of higher market impact. This said, the problem at the heart of optimal execution is simple: for a given level of trading risk, find the trading trajectory that minimizes cost. This sounds simple, and in fact it turns out that fairly straightforward calculations yield explicit formulas for optimal trading paths if two conditions are met. One, the stocks must follow a simple random walk, and two, transaction costs must be quadratic in total trading cost. In a moment we will discuss the ramifications of these assumptions with regard to the theory, but first we finish our exposition of value under liquidation.

With the above background out of the way, it is straightforward to compute value under liquidation. We proceed in the following steps:

- Determine the most tolerable level of uncertainty in trading revenue for the liquidation at hand,
- Compute the optimal trading trajectory for this level of uncerrtainty
- Compute the expected value of trading revenue liquidating over this trajectory

Value under liquidation is the expected value of trading along the optimal trajectory. Note that value is not actually a solid, single value: it is the statistical expected value of the liquidation and is subject to uncertainty. The uncertainty is unavoidable as it is due to the market movements that will be experienced while trading. The only thing one can control is the level of uncertainty, and this is determined by risk management policy ahead of trading, and thus is a manageable uncertainty.

3 A Technical Overview of Optimal Execution

We now proceed with a more technical treatment of the above exposition, focussing on the mathematical assumptions, the formulas for optimal trajectories and finally the ramifications and realism of the assumptions. The treatment here is brief. For a more detailed discussion, see Almgren and Chriss (1999).

First, assume we have a basket of securities $S_1, S_{2,} ..., S_N$, each following an arithmetic Brownian motion

$$dS_i = \mu_i dt + \sigma_i dz_i$$

where the coefficient μ_i is the drift and the coefficent σ_i is the volatility and dz_i

is a standard Brownian motion. Moreover, suppose there is a known covariance matrix Σ relating the dz_i's. In what follows, we set $\mu_i = 0$, reflecting the idea that at the time of liquidation we will have no particular view of what direction the stock will move. The formulas and results that follow are predicated on the fact that $\mu_i = 0$. To see what happens in the case where $\mu_i \neq 0$ see Almgren and Chriss (1999). Assume we want to liquidate from now (time 0) to later (time T) and that we will trade over equally spaced periods of time $t_1, t_2, ..., t_M$, where $\tau = t_i - t_{i-1}$ is constant. We describe the transaction cost functions in terms of the number of shares traded over the period of time τ.

Microstructure theory identifies two components to the cost of transacting: permanent and temporary. Permanent costs refer to market impact that persists for the life of the liquidiation. That is, if we execute a trade and in so doing the market moves its concensus view of the price, then this impact will be felt in subsequent transaction levels. On the other hand, temporary costs refer to purely liquidity based costs that reflect the market's short horizon premium for providing liquidity. These costs are not reflected in subsequent transaction prices and are thus referred to as temporary. For the temporary impact function of security i we want to know the "price" of trading n shares (where $n > 0$ means buy and $n < 0$ means sell) in the period of time τ. We express this as a function of the rate of trading per unit time τ

$$\text{temporary impact} = \epsilon\,\text{sgn}(n) + \ \eta_i \frac{n}{\tau}$$

where ϵ is a fixed cost, and *sgn* is the sign function. This says simply that the cost per share of trading n shares is a fixed cost plus a cost proportional to the size of the trade. Note that since the cost per share is linear in the trade size, the cost is quadratic in total quantity traded. Similarly, we set the permanet impact function (a function of the rate of trading per unit time τ) to be linear in the shares traded:

$$\text{permanent impact} = \gamma_i \frac{n}{\tau}$$

Now, we describe the trading trajectory for asset i as a list of numbers n_{ji}, $j = 1, 2, ..., M$ representing the number of shares to trade between time t_{j-1} and t_j. We write $N_i = \sum_j n_{ji}$ for the total number of units of asset i to liquidate. Given this we have the expected cost of trading asset i along the trajectory given by $\mathcal{N}_i = (n_{1i}, n_{2i}, ..., n_{T-1i})$ is written $E[\text{cost}|\mathcal{N}_i]$ and is given by the formula:

$$
\begin{aligned}
E[\text{cost}|\mathcal{N}_i] &= \frac{1}{2}\gamma N_i^2 + \epsilon \sum |n_{ji}| + \frac{\widetilde{\eta}}{\tau} \sum n_{ji}^2 \qquad (1) \\
\widetilde{\eta} &= \eta - \frac{1}{2}\gamma\tau
\end{aligned}
$$

Next write x_{ji} for the number of units of asset i held in the portfolio at time j. Then the variance of cost trading along the trajecory $\mathcal{N}_i$ written $Var[\text{cost}|\mathcal{N}_i]$ is given by the formula:

$$
Var[\text{cost}|\mathcal{N}_i] = \sigma^2 \sum \tau x_{ji}^2
$$

Note that the value of security i under liquidation is simply the initial value less the expected cost of trading, that is, if S_{i0} is the value of security i at time t_0 then the value under liquidation is given by

$$
N_i S_{i0} - E[\text{cost}|\mathcal{N}_i]
$$

and the value of the portfolio under liquidation is simply given by summing over the different trajectories. Now write $\mathcal{N}$ for the trajectories of all the assets in the portfolio and write $E[\text{cost}|\mathcal{N}]$ for $\sum E[\text{cost}|\mathcal{N}_i]$ where the sum is across all securities, and write $V[\text{cost}|\mathcal{N}] = \sum V[\text{cost}|\mathcal{N}_i]$. Now for a fixed constant $\lambda > 0$ define the *cost* of $\mathcal{N}$ as

$$
\mathcal{C}[\text{cost}|\mathcal{N}] = E[\text{cost}|\mathcal{N}] + \lambda V[\text{cost}|\mathcal{N}]
$$

The parameter λ is called the trading risk aversion; one can see that for a fixed level of expected cost and variance of cost we see increasing risk aversion increases cost. Now, the idea of optimal liquidation is to find the trajectory that minimizes cost. That is, to solve the problem

$$
\min_{\mathcal{N}} E[\mathcal{N}] + \lambda V[\mathcal{N}]
$$

where the minimization takes place across all trading trajectories. The problem may be solved explicitly in the case where securities follow an arithmetic random walk and impact functions are quadratic in total transaction size. For a portfolio consisting of a single security, the optimal trajectory is given by the formula:

$$
\begin{aligned}
n_j &= \frac{\sinh(\kappa(T - t_j))}{\sinh(\kappa T)} N \qquad (2) \\
\kappa &\backsim \sqrt{\frac{\lambda\sigma^2}{\eta}} + O(\tau)
\end{aligned}
$$

where n_j is the number of shares to be traded at time t_j and N is the total number of shares of the security in the portfolio. The constant κ,which for reasonably small τ is well approximated by the above formula, determines the curvature of the optimal trajectory and depends only on the level of risk aversion, the temporary impact parameter and the variance of the security. Note that qualitatively the path becomes more curved as risk aversion and variance increase and becomes less curved as trading costs grow. The more curved the trading trajectory, the more rapidly one trades, and thus the more the risk is reduced. Optimal trajectories for portfolios have more complicated formulas which depend on the variance/covariance matrix Σ. We refer the reader to Almgren and Chriss 1999 for more details.

Now with the formulas for the n_j in hand, it is straightforward to compute the expected cost of liquidation and therefore the value under liquidation. To review, all one needs to know are four parameters:

1. The total time for liquidation $T - t_0$,
2. The risk aversion coefficient λ
3. The temporary impact cost parameter η
4. The permanet impact cost parameter γ
5. The securities volatility σ,and
6. The total number units to liquidate N

Now, using equations (1) and (2) we can compute the value under liquidation.

Before turning to the examples, we note that the parameter κ has an interesting interpretation. Write

$$\theta = 1/\kappa$$

and call this the *half-life* of the trade. It is a measure of the liquidity of an asset under liquidation and intuitively represents the time-scale for liquidating a security in the case where there are no time constraints placed on trading. Technically it represents the amount of time one would take to liquidate a factor of e of the security with no time constraints. In the next section we display a graph of half-lives of securities for a cross-section of US stocks.

4 Examples

In this section we provide some simple pictures and examples illustrating the the main points of optimal execution and value under liquidation. To start, we note that for each level of risk aversion $\lambda > 0$ there is a unique optimal trading trajectory $\mathcal{N}_\lambda$ that is optimal, that is, that has minimal variance (that is, uncertainty) for its level of cost. For a trajectory $\mathcal{N}$ we can compute its expected value $E[\text{cost}|\mathcal{N}]$ and its variance $V[\text{cost}|\mathcal{N}]$. If we plot these points in two dimensions for every possible trajectory, we will end up with a curved region whose boundary is the set of optimal trajectories. We call this boundary the *efficient frontier of optimal execution.*

We now illustrate the entire theory for the stock IBM based on data for April 30, 1997 (data supplied by ITG Inc.). In the following pages there are two figures, depicting respectively the efficient frontier of optimal strategies for trading 1,000,000 shares of IBM, three trading trajectories with varying degrees of risk aversion (and one *risk loving trajectory*), and a display of the half-lives of a cross-section of US stocks. What follows is a more detailed description of each figure:

- **Figure 1:** This figure displays the efficient frontier for trading 1,000,000 shares of IBM over a 5 day time horizon. There are three points on the frontier labeled A, B and C. The trajectory corresponding to point A represents a very risk averse trader, point B represents a moderately risk averse trader and point C represents risk neutral trader.

- **Figure 2:** This figure displays the trajectories represented in the efficient frontier in figure 1. Trading along the trajectories in this figure represents three different trading strategies. Notice that strategy A, the very risk averse strategy, demands very rapid trading, while strategy C, the risk-neutral strategy, trades as slowly as possible while still completely liquidating within five days.

- **Figure 3:** This figure shows the characteristic times for a stratified sample of US stocks. Each point in the graph represents a different US stock. Its position is determined by its temporary impact parameter (vertical axis) and its volatility in dollars per day (horizontal access). The dashed lines are lines of constant characteristic time. Two stocks lying on the same line have the same characteristic time and therefore would have the same value under liquidation for a given level of risk aversion and a given number of shares. Each adjacent pair of lines differs by a factor of $\sqrt{10}$.

- **Table 1:** This table displays the values under liquidation for 1,000,000 shares of IBM where liquidation will take place over a 5 day horizon. The values are derived from the trajectories in figure 2. Note that depending on the level of risk aversion trading costs differ by almost a factor of over 3. However, the uncertainties, as measured in standard deviation of dollars, differ by a factor of over five. Note also that the risk neutral trader has an expected cost of liquidation of \$0.93mm with a standard deviation of \$1.977mm. The moderately risk averse trader (trader B) has an expected cost of \$1.345mm with \$1.039mm standard deviation, significantly less risk. This is a result worth noting: because the efficient frontier (figure 1) is smooth at its minimum point, one can obtain a first order reduction in risk while only paying a second order increase in cost.

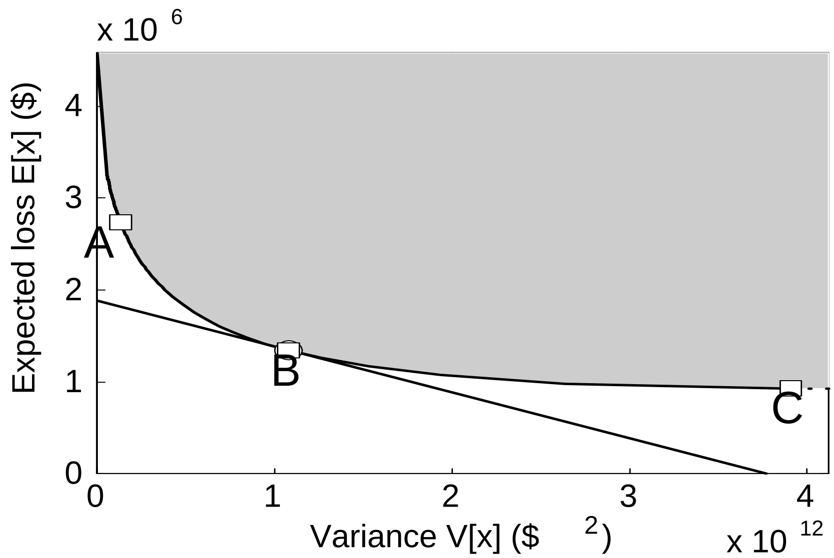

Figure 1:

Table1

Risk Aversion	Value	Cost	Std Dev
Very Risk Averse	$155.6	$2.736	$0.361
Moderately Risk Averse	$156.8	$1.345	$1.039
Risk Neutral	$157.07	$0.93	$1.977

5 Background and Extensions

This paper is based on some results in a larger work entitled *Optimal Execution of Portfolio Transaction* by Robert Almgren and Neil Chriss. This paper was originally submitted to the FEN web site November 24, 1997 under the name *Optimal Liquidation.* In the full work we discuss at length some of the ramifications of the assumptions we make herein concerning asset price movements. The work herein was presented in the section "Liquidity Risk Issues" by Neil Chriss at the Wharton Financial Engineering Round Table, April 1999. We thank ITG Inc. for supplying the transaction cost data used in the examples. Also, ITG Inc. recently introduced a product called ACE (Agency Cost Estimator) that is similar to the idea of value under liquidation (but is used for the purpose of computing execution costs), and is based on the ideas in Almgren and Chriss

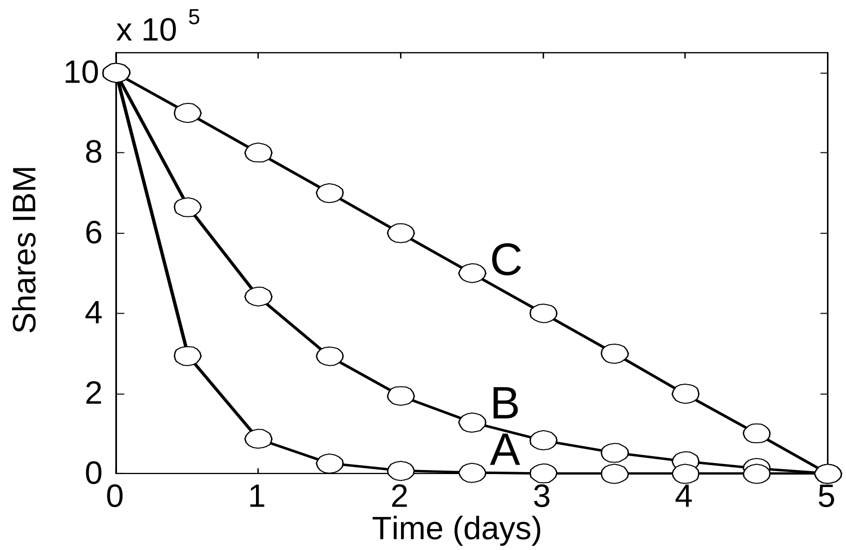

Figure 2:

(1999) with several extensions. Essentially, ACE computes the expected cost of executing a basket of securities along the multivariate optimal liquidation path described by the multi-asset version of optimal execution. ACE uses a proprietary parametrization of the temporary and permanent impact functions to compute transaction costs, and assumes that the price impact functions and volatilities vary within each trading day according to the familiar u-shape pattern concentrating greater liquidity and volatility just after the open and just before the close.

6 Bibliogaphy

Almgren and Chriss (1999): *Optimal Execution of Portfolio Transactions*, working paper.

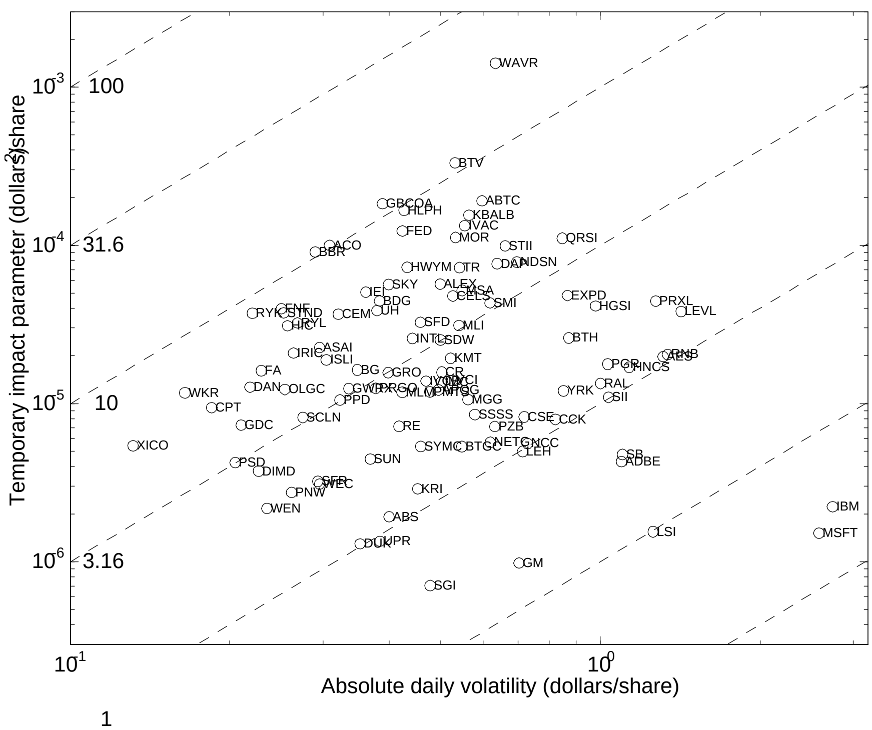
Temporary impact parameter (dollars²)/share
Absolute daily volatility (dollars/share)
10³
10⁴
10⁵
10⁶
10¹
10⁰
100
31.6
10
3.16
1
WAVR
BTV
ABTC
GBCOA
HLPH
KBALB
IVAC
FED
MOR
QRSI
STII
ACO
BBR
HWYM
TR
ADSN
SKY
ALEX
MSA
CELS
SMI
EXPD
HGSI
PRXL
LEVL
BDG
CEM
UH
SFD
MLI
INTC
LSDW
BTH
ASAI
ISLI
KMT
FA
BG
GRO
PGR
HNCS
WKR
DAN
OLGC
PPD
MGG
YRK
RAL
SII
CPT
GDC
SCLN
RE
SSSS
PZB
CSE
CCK
XICO
SYMC
BTGC
NETG
NCC
LEH
SB
ADBE
PSD
DIMD
SUN
PNW
WEC
KRI
WEN
ABS
DUK
UPR
GM
SGI
LSI
IBM
MSFT

Figure 3: