

December 2002

Optimal Liquidation of Venture Capital Stakes

Robert Dubil
San Jose State University

Follow this and additional works at: <https://digitalcommons.pepperdine.edu/jef>

Recommended Citation

Dubil, Robert (2002) "Optimal Liquidation of Venture Capital Stakes," *Journal of Entrepreneurial Finance and Business Ventures*: Vol. 7: Iss. 2, pp. 65-82.

DOI: <https://doi.org/10.57229/2373-1761.1091>

Available at: <https://digitalcommons.pepperdine.edu/jef/vol7/iss2/6>

This Article is brought to you for free and open access by the Graziadio School of Business and Management at Pepperdine Digital Commons. It has been accepted for inclusion in The Journal of Entrepreneurial Finance by an authorized editor of Pepperdine Digital Commons. For more information, please contact bailey.berry@pepperdine.edu.

Optimal Liquidation of Venture Capital Stakes

Robert Dubil⁺⁺
San Jose State University^{*}

We model the optimal liquidation behavior of a venture capital or non-diversified asset management firm faced with a sale of concentrated security holdings. As the firm's stake is large, its sales can lead to permanent and temporary price depressions. At the optimum, the institution chooses the liquidation interval to balance the exposure to the market return variance against the impact of its own sales on the realized return. We obtain closed-form solutions for power impact functions uncorrelated with returns. We also consider market impact correlated with the return process, i.e. a case where liquidity evaporates during severe price dislocations.

By the very nature of the business, venture capitalists find themselves holding significant stakes in the companies they own. Cashing out of their positions carries significant liquidity costs and risks. Their sales are closely watched and are perceived to have informational content. Large liquidations are subject to price pressures and time delays. This paper offers a theoretical model of the optimal disposition strategy for a venture capital (VC) firm.

Market microstructure research studies the role of trading mechanisms on the price-setting process. Most papers, the best known being Kyle (1985) and Glosten and Harris (1988), investigate the size of the bid-ask spread in the listed stock markets. Their goal is largely descriptive. We build on the behaviorally prescriptive work of Bertsimas and Lo (1998), and Almgren and Chriss (2000). They consider an optimization problem of an agent who liquidates a large block of securities. As the agent trades, he affects the market price. His impact on the market price can be temporary (disappears when the market absorbs the quantity he supplies) and permanent (persists due to its informational content). The models solve for the optimal sales trajectory (a sequence of sales quantities) over the total liquidation time. They assume that the final liquidation horizon is exogenously given, and so their results are particularly useful for proprietary traders where the holding time is pre-determined by a policy limit or some relative-value strategy

⁺⁺ Robert Dubil is an Associate Professor of Finance at San Jose State University. He holds a Ph.D. and M.B.A. from U. of Connecticut, and MA (ABD) from Wharton. He is a former Director of Risk Analytics in the Corporate Risk Management group at Merrill Lynch (1999-2001), Head of Exotic Fixed Income Derivatives Trading at UBS (1996-99) and Chase Manhattan (1994-95), an equity and debt derivatives trader at Merrill Lynch (1992-94), and a quantitative researcher at Nomura (1990-92) and JP Morgan (1989-90). His recent research interests cover liquidity risks and banking regulation. His most recent articles on those subjects can be found in the Financial Services Review and the Journal of Investing.

^{*} Special thanks for helpful comments to Shantaram Hegde, my colleagues in the Corporate Risk Management group at Merrill Lynch, the participants at the 2001 FMA Doctoral Colloquium in Toronto and the 2002 EFA meetings in Baltimore.

horizon. In this paper, we take the perspective of an agent that routinely, but infrequently, trades illiquid securities. This could include a venture capitalist, a private equity firm, or a high-yield fund. For such an agent, the best total time in which to liquidate any given security position, is not exogenously given, but is itself a solution to an optimization problem. The agent solves for the optimal length of time over which he engages in a sequence or a continuum of equal-size trades. He trades off between the impact of his trades on the realized sale value and the risk of holding the assets for an extended period of time. The focus is on the concept of endogenous liquidity, as defined by Bangia et al. (1999). An agent who decides to sell a large position in the market, where ‘large’ is defined as exceeding the standard quote depth, will adversely affect the price at which he transacts if he sells too quickly. This sale discount will be the price he will pay for avoiding the market risk of the position. At the optimal time to liquidate, the change in his marginal utility will be zero.

The only simplifying assumption we make is that of constant sales speed which may lead in non-trivial cases to sub-optimality in the stochastic control sense. Almgren and Chriss (2000) show that the deviation is relatively small for the type of “utility” function we choose. The restriction is motivated by practical considerations. First, it introduces a simple one-dimensional liquidity metric with which to assess the combined market and liquidity risk of different equity stakes a venture capital firm may hold. This avoids having to account for several risk factors like the stock return variances, the sizes of the positions relative to the quote depths, the impact parameters, etc. Second, for concentrated stakes, it is hard to segregate permanent and temporary impacts. A VC firm with positions in several assets is likely to ignore the extra risk of a linear strategy relative to its optimal non-linear alternative for a given final time, and instead will focus on the total time.

Following the 1996 mandate from the Federal Reserve, money center banks with trading portfolios started disclosing market risk statistics in their annual statements and setting aside capital against those risks based on internal Value-at-Risk (VaR) models. For definitions, see Bank of International Settlements (1996). Increasingly, private equity and asset management firms have started to adopt similar models to manage the non-diversifiable (market price and liquidity) risks of their portfolios. Jarrow and Subramanian (1997), and Bangia et al. (1999) extend the market VaR model to exogenous liquidity risk factors, i.e. those beyond the firm’s control (e.g. quote size, bid-ask spread). Dubil (2001) extends it further to endogenous liquidity and provides a method of aggregating VaR across the firm’s sub-portfolios, each subject to a different liquidation horizon. This paper contributes indirectly to this line of research in that it offers a method of parameterizing the necessary inputs for the enhanced VaR model. At the same time, we make a practical use of the VaR concept: as a simple “utility” function which relates the risk (profit variance) to the return (proceeds) of the liquidation program. The agent reveals his risk preferences by choosing a confidence interval parameter, α , which determines the worst-loss tail probability he is willing to accept. The choice of a “utility” function affects the behavior of the optimizing agent, but our results do not depend on that particular choice. We could easily obtain closed-form solutions for alternative formulations involving a general HARA class as defined by Huang and Litzenberger (1988).

The paper is organized as follows. First, we review the research examining the optimal liquidation behavior of an agent selling a large block to the market. The control parameter is always the sales trajectory over a given a fixed sales horizon, but the studies differ in the assumptions they make about the underlying price processes and objective functions. Then, we develop, from first principles, a model of the optimal liquidation time under the assumption of a constant speed trajectory. The model is recast from units and prices into the terms of dollar exposures and returns. The latter offers the preferred log-normality of prices and is applicable to the general case of liquidating a large block of shares or a private equity stake prior to an IPO. We consider two formulations of the impact functions: general power uncorrelated with the price and linear correlated with the price process. We obtain closed-form solutions for the first group and numerical ones for the second which is in a sense more general, as it allows for the market impacts to change (i.e. increase) during large market dislocations. We provide intuitions for the example solutions and include the discussion of the alternatives for the optimization functions.

I. Optimal Liquidation Trajectories in Finance Literature

As Chan and Lakonishok (1995) showed, a typical large investor's trade in the stock market is broken down into smaller packages and executed over a period of four days or more. Presumably, such protracted liquidation is designed to minimize the adverse impact on the overall transaction price, but exposes the trader to market risk. Two studies – Bertsimas and Lo (1998), and Almgren and Chriss (2000) – examine the best execution strategies for a stock trader under the constraint that he decides in advance on the final close-out date. The trader's objective is to acquire (liquidate) a fixed number of shares, X , in a fixed time period, T . As market conditions change he speeds up or slows down his purchases (sales), but his holdings are always equal to the targeted amount (zero) on the final date.

Bertsimas and Lo (1998) consider an expected trading cost minimization for a program designed to acquire a fixed number of shares, X , by the final time, T . While the program is in effect, new information arrives in the market in the form of random shocks to the trading price. The price is also affected by each trade executed in the program. The authors write the basic Bellman equation for this dynamic programming optimization and employ the optimal control machinery to solve recursively for the best trading trajectory. A trajectory is defined as a sequence of the amounts purchased in each of the N equally spaced time intervals. They show that the best strategies are often linear combinations of a 'naïve' strategy, of breaking the total size X into T identical packages of size X/T , and a correction portion reflecting the new information. In the absence of private stock-specific information, although the naïve strategy is not optimal in general, it is the best under the assumption that the price follows an arithmetic random walk and the impact of the liquidation strategy is linear. The authors consider alternative formulations of the price process and solve for optimal execution strategies for portfolios of correlated assets. They also illustrate the difficulties of imposing constraints on the optimization parameters.

Holthausen et al. (1987 and 1990) first introduced the intra-period distinction between temporary and permanent impacts. They estimated the impact of large block trades on NYSE on the stock prices. They broke each large block transaction, whether

seller- or buyer-initiated, into two estimable variables. The temporary price effect was defined as the continuously compounded return earned on the difference of the block transaction price and the equilibrium price prior to the block transaction. The permanent price effect was defined as the return on the difference between the equilibrium prices after and before the block transaction. For a seller-initiated block, these are depicted in Figure 1 (see Holthausen et al., 1987, Fig. 1).

In Almgren and Chriss (2000), who follow the same set-up, the stock price is subject to both effects in each of the N equally spaced intervals (t_{k-1}, t_k) , $k = 1, \dots, N$, of length $\tau = t_k - t_{k-1} = T/N$, $\forall k$. A trader faces selling X shares over the fixed total time T through a sequence of sales in each of the N intervals. His holdings at the end points of the intervals are $x_0 = X$, x_1, x_2, \dots, x_{N-1} , $x_N = 0$, his sales during the intervals are $x_k - x_{k-1}$, $k = 1, \dots, N$, and the speed with which he sells in each interval is denoted by $v_k = (x_k - x_{k-1})/\tau$. The equilibrium price follows an arithmetic random walk with no drift, but is subject to a permanent impact effect resulting from the trader's action. The equation of motion for the equilibrium price is:

$$S_k = S_{k-1} + \sigma\sqrt{\tau} \Delta \tilde{z}_k - \tau g(v_k) \quad (1)$$

where $g(v_k)$ represents the permanent impact function, σ is the annualized normal volatility of the stock price, and $\Delta \tilde{z}_k$ is a standard normal deviate, with $E[\Delta \tilde{z}_k \Delta \tilde{z}_{k'}] = 0$, $k \neq k'$. Note that $g(v_k)$ is pre-multiplied by τ to emphasize that the total permanent impact effect depends more on the total number of shares, τv_k , sold in each interval, than on the pure intensity parameter v_k . The trading price, \tilde{S}_k , the trader faces in each interval is subject to a temporary impact function $h(v_k)$:

$$\tilde{S}_k = S_{k-1} - h(v_k) \quad (2)$$

Instead of simply minimizing the expected cost, the trader cares about the risk of the strategy (variance of the liquidation cost). He minimizes the negative utility associated with the cost of the protracted liquidation over the time interval $(0, \dots, T)$. That cost, C , is defined as the difference between an instantaneous sale of all of his X units at the time-zero price S_0 and the sum of the proceeds from a sequence of sales of $(x_k - x_{k-1})$, each at the trading price \tilde{S}_k , over the entire time interval $(0, T)$, i.e. $C = XS_0 - \sum_{k=1}^N (x_k - x_{k-1}) \tilde{S}_k$. The cost is subject to the sequence of the random shocks $\sigma\sqrt{\tau} \Delta \tilde{z}_k$. Since the shocks are independent, the mean and the variance of the cost function can be derived as functions of the strategy sequence. The impact functions are

assumed to be linear in the sales intensity parameter v_k , and parameterized as $g(v_k) = \gamma_k$ and $h(v_k) = \varepsilon \text{sgn}(n_k) + \eta v_k$. Hence, $h(v_k)$ is denominated in \$/share, ε can be thought of as half the bid-ask spread denominated in \$/share, and η has a peculiar denomination of [\$/share]/[share/time]. Almgren and Chriss (2000) define a mean-variance efficient frontier for a trader who minimizes the expected cost given a level of variance. This is equivalent to an unconstrained minimization of

$$\min_{\{x_k, k=1, \dots, N\}} \{E[C] + \alpha V[C]\} \quad (3)$$

where the Lagrange multiplier α can be interpreted as the relative risk aversion coefficient. When $\alpha > 0$, a unique solution $\{x_k\}^*$ is guaranteed by the strict concavity of the minimand. For a risk-neutral trader ($\alpha = 0$), the optimal trajectory is a straight line of declining holdings over time defined by the decrement $n_k = X / N$. Risk averse traders follow convex lines below that line; risk loving traders follow concave lines above the straight line. Risk averse traders sell relatively more up front, and less later. They incur higher impact costs in order to avoid the exposure to the random shocks.

The two papers differ in their choices of the optimized functions. In Bertsimas and Lo (1998), the agent is an expected cost minimizer. The expected cost depends on the variance of the price process, but the agent does not care about the variance of the cost itself. The optimal control methodology allows for the agent to change his strategy during the execution of the program. Almgren and Chriss (2000) explicitly make their agent worry about risk by choosing only those solutions which lie on the mean-variance efficient frontier. Equivalently, they adopt a constant relative risk aversion utility function. If the latter is of a particular form (e.g. log-utility), the static up-front optimization yields the same solution as the less restrictive dynamic program. The authors discuss the use of VaR utility which does not guarantee that. It does avoid, however, their somewhat counterintuitive result whereby large and small baskets are liquidated identically by a trader with a given risk aversion coefficient. It also produces solutions closer to the straight line. In our optimization, we adopt VaR as the objective function¹ even though the analysis does not depend on this assumption. This allows us to relate risk to return in a straightforward, intuitively parameterized, formulation. Dubil (2002) argues that this is the best choice for concentrated wealth liquidations in a non-general-equilibrium setting. He shows that it is equivalent to a mean-standard-deviation ‘efficient frontier’ selection.

In order for a liquidity model to be of practical use, the market impact parameters have to be computed. The liquidity effects of this paper can be readily estimated by adapting the methodology of Glosten and Harris (1988) who consider two components of the bid-ask spread of NYSE stocks. The transitory component allows market-makers to cover inventory costs, clearing fees and monopoly profits and explains the negative serial correlation of closing prices. The persistent component due to adverse selection costs allows market-makers to recover from liquidity traders losses on trades with informed

¹ Basak and Shapiro (2001) examine some unappealing outcomes when the VaR-derived utility function is used in a continuous-time partial equilibrium.

traders. This information asymmetry component does not contribute to the serial correlation and is thus separable from the first. The paper relates the adverse selection cost to order size and applies a likelihood method to estimate and decompose the spread from time-series prices into the inventory and asymmetric information parts. Data requirements are quite modest and the specification is extendible to other markets of interest (VC stakes, private placements, high yield, etc.). We refer the reader to O'Hara (1997) for a more extensive list of references.

II. The Optimal Liquidation Horizon Model

In the reviewed studies of Section I, the agents face an externally imposed final liquidation time and solve for the optimal sales trajectory between time zero and that final time. The problems are therefore posed from a perspective of a proprietary trader who must liquidate by the end of a certain time interval (say, by the end of the day). In our model, the venture capital or high-yield investor does not face an exogenous final sale time. As our agent's problem is more general and a lot more difficult – to choose the best time and the best strategy – we choose to make a more-fitting simplifying assumption. He solves for the optimal final horizon by adopting a simple strategy to liquidate at a constant rate per unit of time. His strategy trajectory (a plot of his remaining position against time) is a straight line. We solve for the final time point where the line crosses the horizontal axis. This is equivalent to solving for the one constant rate of his sales expressed in units sold per time. His behavior is motivated by the fact that he has no special information about the timing or the 'lumpiness' of other trades coming to the market beyond knowing the market impact of his own actions. He chooses the simplest strategy to sell the same amount in each time interval, but is concerned about the total time of liquidation.

We also do not follow the set-up of an agent optimally liquidating X number of shares under the assumption that the equilibrium stock price is locally an arithmetic Brownian motion with no drift². Instead, we consider an agent optimally reducing, instead of the number of shares, his dollar exposure to the return on a stock under the more palatable assumption that the return, not the price, follows a locally arithmetic Brownian motion. The mathematics are essentially the same, but the parameter interpretations are slightly different. Note that the main difference is the stochastic process assumption (log-normal price) and the form of the impact functions (cost defined as a known function of returns instead of the price).

Within the model, we consider two different forms of the market impact formulations: a general power function, which is uncorrelated with the price process, with some of its special cases (e.g. linear, square root), and a stochastic linear impact function correlated with the price process. The latter case allows for a feedback loop, whereby a significant drop in returns (prices) can cause a deterioration of liquidity. The agent's objective is to choose optimally the final time by which his holdings will be reduced to zero under the assumption that he sells at a constant speed over that horizon. His sales cause a permanent market return change at the "end" of each interval $(t - dt, t)$. They also cause a temporary shock to the realized return which deviates from the general market return, during the interval $(t - dt, t)$, but dissipates completely by the "end" of

² The equilibrium price has no drift except that due the permanent impact of the agent's sales.

that interval. We work in continuous time, but it is easy to show that the results are a natural limit of the discrete version in which the number of trading intervals within the fixed final horizon increases to infinity³.

Let W_t be the dollar amount of investment in the underlying asset (VC stake, listed stock or high-yield bond) at time t . The initial exposure to the asset return is denoted by W_0 , and the final exposure is zero, i.e. $W_T = 0$. Let us define the agent's holdings at time t , as the ratio of his remaining dollar exposure to the original investment, i.e. let $x_t = \frac{W_t}{W_0}$. The agent starts with holdings $x_0 = X = 1$ at time $t = 0$ and liquidates all of them by the final time T , i.e. $x_T = 0$. In each interval $(t - dt, t)$, the agent sells $-dx_t = \lim_{dt \rightarrow 0} (x_{t-dt} - x_t)$ of his holdings. The speed of trading in each interval is defined as $v_t = -\frac{dx_t}{dt}$. The equilibrium return on the underlying asset R_t , representing all public information in the market, follows an arithmetic Brownian motion process⁴. The only drift is due to the accumulated permanent impact on the price from the sales executed by our agent. The equilibrium return, R_t , at time $t \in [0, T]$ can thus be written:

$$R_t = \sigma z_t - \int_0^t g(v_s) ds \quad (4)$$

where $g(v_t)$ represents the permanent impact of the agent's sales on the equilibrium return, and z_t is the standard Wiener process. The cumulative trading return, \tilde{R}_t , realized by the agent through a trade within each interval $(t - dt, t)$ ⁵, is also subject to the temporary impact $h(v_t)$ as a function of the speed of sales within the interval:

$$\tilde{R}_t = R_t - h(v_t) \quad (5)$$

Let us further assume that the agent sells at a constant speed $v_t \equiv v$, so that

$$-dx_t = v \, dt \quad (6)$$

³ In continuous time, it is less intuitive to imagine the temporary impact dissipating by the end of an infinitesimal interval while the permanent impact persists.

⁴ The cumulative return, R_t , is computed over the entire interval $(0, t)$. In discrete time, the cumulative return through period k , would be equivalently defined by $1 + R_k = \prod_{i=1}^k (1 + r_i)$, where r_i is the return on the asset in period i .

⁵ Note that both R_t and \tilde{R}_t are random. The tilde symbol merely distinguishes the 'intra-period' return realized by our agent from the equilibrium return in the market.

The excess profit (most likely negative, i.e. a cost) due to the non-instantaneous liquidation, as a fraction of his original wealth, is defined as:

$$\Pi = \int_0^T \tilde{R}_t(-dx_t) = \int_0^T \tilde{R}_t v dt \quad (7)$$

In order to introduce risk into the agent's decision process, we adopt a concave utility function. We choose the mean-standard deviation VaR framework. The VaR of a set of positions is defined as a profit value, Π^* , for which the probability of the profit, Π , falling below that critical value, Π^* , is equal to a given quantity. The latter is equal to one minus the chosen confidence interval. For a normally distributed profit function, the VaR, Π^* , is defined explicitly as:

$$\Pi^* = E[\Pi] - \alpha \sqrt{V[\Pi]} \quad (8)$$

where α is a known number from the standard normal table corresponding to the chosen confidence level, e.g. 1.645 for a 5% left-tail probability. $E[\Pi]$ and $V[\Pi]$ denote the expected value and the variance of the profit. The confidence level-related parameter α embodies an implicit set of risk preferences of the agent. His objective is to choose an optimal final liquidation time to maximize the VaR of his profit:

$$\max_T \Pi^* = \max_T \{E[\Pi] - \alpha \sqrt{V[\Pi]}\} \quad (9)$$

Eq. (9) is analogous to Eq. (3) as $\Pi = -C$, with a change to the square root term in the variance, and is equivalent to the minimization of the VaR of the agent's excess cost due to the non-instantaneous liquidation. Dubil (2002) argues the appropriateness of using VaR in liquidity-constrained wealth liquidation cases, e.g. for venture capitalists and executives with vested stock holdings. The interpretation of the agent's 'utility' function is straightforward and appealing. He maximizes the expected profit of the liquidation (sale revenue net of the market impact), but assigns a penalty function to the risk defined as the standard deviation of the profit. The penalty parameter, α , depends directly on his 'worry' level. An institutional argument in favor of choosing Eq. (9) as the specification is that the optimand in Eq. (8) is often viewed as a cost determinant rather than a risk measure. Bank and asset management firms multiply the dollar VaR by their borrowing rate as the amount of capital they "carry" to cover potential cost of liquidation. This is the prescribed Bank for International Settlements definition of the market risk capital.

III. A General Power Function Market Impact

Let us assume that the market impact functions are of the following general form:

$$\begin{aligned} g(v_t) &= \mathcal{W}_t^G \\ h(v_t) &= \varepsilon \operatorname{sgn}(n_t) + \eta v_t^H \end{aligned} \quad (10)$$

where G and H are given constants⁶. In a pure sale strategy, $\text{sgn}(n_t) = 1$ for all t . The trading return in the interval $(t - dt, t)$ is equal to:

$$\tilde{R}_t = \sigma z_t - \varepsilon - \eta v_t^H - \gamma \int_0^t v_s^G ds \quad (11)$$

The last three terms in Eq. (11) represent a total liquidity discount. Using the Def. (7), we can write the profit function for the agent liquidating at a constant speed, v , as⁷:

$$\Pi = v\sigma \int_0^T z_t dt - \varepsilon v T - \eta v^{H+1} T - \frac{1}{2} \gamma v^{G+1} T^2 \quad (12)$$

The expected value of the profit is equal to⁸:

$$E[\Pi] = -\varepsilon v T - \eta v^{H+1} T - \frac{1}{2} \gamma v^{G+1} T^2 \quad (13)$$

and the variance of the profit is equal to⁹:

$$V[\Pi] = \frac{1}{3} T^3 \sigma^2 v^2 \quad (14)$$

Since $X = vT$, the mean and the variance can be written as:

$$\begin{aligned} E[\Pi] &= -\varepsilon X - \eta X^{H+1} T^{-H} - \frac{1}{2} \gamma X^{G+1} T^{-G+1} \\ V[\Pi] &= \frac{1}{3} T X^2 \sigma^2 \end{aligned} \quad (15)$$

Now we can set up the agent's optimization problem as the trade-off between the total liquidity discount, affecting the mean of the profit, and the 'market risk' of returns, represented in the total variance of the profit, multiplied by the penalty parameter α , representing the agent's risk tolerance:

$$\max_T \left\{ -\varepsilon X - \eta X^{H+1} T^{-H} - \frac{1}{2} \gamma X^{G+1} T^{-G+1} - \alpha \sqrt{\frac{1}{3} T X^2 \sigma^2} \right\} \quad (16)$$

⁶ Almgren (2001) extends original linear temporary impact to power specifications.

⁷ Write $\Pi = v\sigma \int_0^T z_t dt - \varepsilon v \int_0^T dt - \eta v^{H+1} \int_0^T dt - \gamma v^{G+1} \int_0^T \left(\int_0^t ds \right) dt$ and reduce to Eq. (12).

⁸ Use integration by parts to show that $E[\int_0^T z_t dt] = E[tz_t |_0^T] - E[\int_0^T t dz_t] = 0$.

⁹ Show that $\int_0^T z_t dt = tz_t |_0^T - \int_0^T t dz_t = Tz_T - \int_0^T t dz_t = T \int_0^T dz_t - \int_0^T t dz_t = \int_0^T (T-t) dz_t$. Then show that

$V[\int_0^T z_t dt] = E[\int_0^T (T-t) dz_t]^2 = E[\int_0^T (T-t)^2 dz_t^2] = E[\int_0^T (T-t)^2 dt] = \int_0^T (T-t)^2 dt = \frac{1}{3} T^3$.

The first-order condition for this optimization is the following equation in T :

$$H\eta X^{H+1}T^{-H-1} + \frac{1}{2}(G-1)\gamma X^{G+1}T^{-G} - \frac{1}{2\sqrt{3}}\alpha\sigma XT^{-\frac{1}{2}} = 0 \quad (17)$$

Note that $X = 1$. It is carried here to emphasize that all the parameters in (10) are relative and can be defined for the holdings of the firm or a market. In general, Equation (17) has to be solved by numerical methods. However, we can obtain closed-form solutions for some special cases.

A. Linear Market Impact Functions

First let $G = 1$. The permanent impact function is linear in the speed of sales, i.e. $g(v_t) = \gamma v_t$ and $h(v_t) = \varepsilon + \eta v_t^H$. Eq. (17) can be solved explicitly for T :

$$T = \left(\frac{2\sqrt{3}H\eta X^H}{\alpha\sigma} \right)^{\frac{1}{H+\frac{1}{2}}} \quad (18)$$

If we further assume that $G = H = 1$, so that $g(v_t) = \gamma v_t$ and $h(v_t) = \varepsilon + \eta v_t$, then the explicit solution for the optimal liquidation time becomes:

$$T = \left(\frac{2\sqrt{3}\eta X}{\alpha\sigma} \right)^{\frac{2}{3}} \quad (19)$$

Let us consider a few numerical examples for the linear market impact $G = H = 1$. Let $\alpha = 1.645, \sigma = 0.15, \eta = 0.001714$. This corresponds to a 95% confidence interval VaR utility, an annualized return volatility of 15%, and a temporary impact of 17.14 bp when selling at a pace of full holdings over a year. The optimal liquidation time is equal to $T = 0.0833$ or 1 month. If the annualized volatility is changed to 10% ($\alpha = 1.645, \sigma = 0.10, \eta = 0.001714$), the agent optimally sells over 40 days as the risk of staying in the market is reduced. Instead, if he wants the same profit threshold at a 99% confidence interval, i.e. he becomes more “risk-averse” ($\alpha = 2.33, \sigma = 0.15, \eta = 0.001714$) he will reduce his liquidation time to about 24 days.

The solutions (18)-

(19) do not depend on the permanent impact coefficient, γ , as it affects the trading return only through the cumulative sales total and not the sale amount in each interval. They are also independent of the mean temporary impact, ε , which is a constant subtraction from the return no matter what the speed of sales is.

B. Other Special Cases

Let us consider the square root case for the permanent impact with a general form of the temporary impact function. We let $G = \frac{1}{2}$ and H be general. The impact functions are $g(v_t) = \gamma\sqrt{v_t}$ and $h(v_t) = \varepsilon + \eta v_t^H$, and the closed-form solution is

$$T = \left(\frac{2\sqrt{3}H\eta X^H}{\alpha\sigma + \frac{\sqrt{3}}{2}\gamma\sqrt{X}} \right)^{\frac{1}{H+\frac{1}{2}}} \quad (20)$$

If the temporary impact is also of a square root form, $H = \frac{1}{2}$, so that the impact functions are $g(v_t) = \gamma\sqrt{v_t}$ and $h(v_t) = \varepsilon + \eta\sqrt{v_t}$, then the closed-form solution is

$$T = \frac{\sqrt{3}\eta}{\alpha\sigma + \frac{\sqrt{3}}{2}\gamma} \quad (21)$$

The choice of the exponent in the power function is subject to empirical research. The exponent of less than one, the square root in particular, is quite intuitive, as it exhibits diminishing returns. As one increases the speed of liquidation, initially the market impact grows rapidly, but at high levels of sales, the growth in the market impact decreases. In the square root case the solution is linear in the temporary impact parameter η and is independent of the initial size of the holdings to be liquidated, X .

What stands out in (20) and (21) is that the different components of the liquidity discount do not have the same effect on the solution. The higher the temporary impact parameter η , the longer the optimal final time. At a given speed, an increase in the temporary impact parameter reduces the agent's realized return for each sale. To compensate for that, the agent slows down his sales and extends his selling horizon. This can be seen in Eq. (17) with $G=H=1/2$. The part of the agent's maximand related to the temporary impact is an increasing function of T . In contrast, the objective function is a decreasing function of the final time T . The permanent impact parameter γ also reduces the realized return in each interval, but only via an expression related to the cumulative amount transacted. It thus acts like a deterministic time-proportional 'drag' on the expected returns and profit, not unlike the market risk penalty function on the agent's total utility. Were this 'drag' per unit of time to increase, the agent would speed up his sales to minimize its negative effect on total profit. In the linear case, this 'drag' was simply proportional to the total amount for sale and independent of time. With $G=1$, it dropped out of the Eq. (17). As such, it operated similarly to a constant bid-ask spread and did not affect the solution at all.

Let us now set $\gamma = 0$. As it is often difficult to distinguish between the information-driven equilibrium price movement and one's own permanent impact, it is convenient to assume no permanent impact to ease the estimation of the temporary effects due to one's inventory problems. We allow a general form of the temporary impact, i.e. $g(v_t) = 0$ and $h(v_t) = \varepsilon + \eta v_t^H$, and still get a closed-form solution.

$$T = \left(\frac{2\sqrt{3}H\eta X^H}{\alpha\sigma} \right)^{\frac{1}{H+\frac{1}{2}}} \quad (22)$$

The solution (22) is identical to that of the linear case. There the permanent impact affected the equilibrium return only through the sales total, and not in each interval. Here that reduction is zero.

IV. Stochastic Market Impact Correlated With Returns

The specification in this section captures a systemic feedback loop where, as prices and realized returns in the market drop, liquidity deteriorates, leading to further temporary depression of the returns. An obvious example here is the bursting of the tech bubble in 2000. This can be accomplished by assuming

$$\begin{aligned} g(v_t) &= \gamma v_t \\ h(v_t) &= \varepsilon + (\eta + \theta \tilde{z}_t) v_t \end{aligned} \quad (23)$$

where \tilde{z}_t is another Wiener process correlated with that driving the returns in (4), i.e.

$$\begin{aligned} E[dz_s d\tilde{z}_t] &= \rho dt, \quad \text{if } s = t \\ &= 0, \quad s \neq t \end{aligned} \quad (24)$$

and η and θ are two constants. The trading return is a random variable dependent on the realizations of two Brownian motions:

$$\tilde{R}_t = \sigma z_t - \varepsilon - (\eta + \theta \tilde{z}_t) v_t - \gamma \int_0^t v_s ds \quad (25)$$

The profit function can be written as¹⁰:

$$\Pi = v\sigma \int_0^T z_t dt - \varepsilon vT - \eta v^2 T - \theta v^2 \int_0^T \tilde{z}_t dt - \frac{1}{2} \gamma v^2 T^2 \quad (26)$$

To set up the optimization, we derive, using $X = vT$, the expected profit as

$$E[\Pi] = -\varepsilon X - \frac{\eta X^2}{T} - \frac{1}{2} \gamma X^2 \quad (27)$$

¹⁰ Write $\Pi = v\sigma \int_0^T z_t dt - \varepsilon v \int_0^T dt - \eta v^2 \int_0^T dt - \theta v^2 \int_0^T \tilde{z}_t dt - \gamma v^2 \int_0^T (\int_0^t ds) dt$ and reduce.

and the variance of the profit¹¹ as

$$V[\Pi] = \frac{1}{3} X^2 \left(\sigma^2 T + \frac{\theta^2 X^2}{T} - 2\rho\theta\sigma X \right) \quad (28)$$

The agent maximizes the VaR of his profit:

$$\max_T \left\{ -\varepsilon X - \frac{\eta X^2}{T} - \frac{1}{2} \gamma X^2 - \alpha \sqrt{\frac{1}{3} X^2 \left(\sigma^2 T + \frac{\theta^2 X^2}{T} - 2\rho\theta\sigma X \right)} \right\} \quad (29)$$

He faces the following first-order condition:

$$\frac{\eta X^2}{T^2} - \frac{1}{2} \alpha \left(\frac{1}{3} X^2 \right) \left(\sigma^2 - \frac{\theta^2 X^2}{T^2} \right) \left[\frac{1}{3} X^2 \left(\sigma^2 T + \frac{\theta^2 X^2}{T} - 2\rho\theta\sigma X \right) \right]^{-\frac{1}{2}} = 0 \quad (30)$$

Eq. (30), equivalent to a 5-th degree polynomial in T , can be solved numerically.

Table I shows the results using the numerical values of the previous section, $\alpha = 1.645$, $\sigma = 0.15$, $\eta = 0.001714$. Recall that the optimal liquidation time without the correlated term was $T = 0.0833$ which was equivalent to 1 month. Here $\theta = 0$ is equivalent to the non-stochastic linear case. One common feature of all the results is that, for all correlations levels, the optimal liquidation time attains a minimum at some level of θ ¹² and increases rapidly as we move away from that level in either direction. That is, the larger the shocks in absolute value, the more the agent compensates for them by staying longer in the market. In the combinations $\theta > 0$, $\rho < 0$ (and $\theta < 0$, $\rho > 0$), the liquidity deteriorates when the returns in the markets drop. The agent's profit has high variability as the combined shocks to his trading returns are strengthened. It is optimal for him to extend the liquidation period in order to reduce the temporary impact on the realized return. When the 'reinforcement' coefficient is low, $\theta = 0.001$, and the correlation of the market impact with the return process is negative, $\rho = -0.50$, the agent liquidates quickly within a little over a month, $T^* = 0.08596$. When the feedback loop is amplified,

¹¹ To derive the variance we need to evaluate the following expression:

$$I = E\left[\int_0^T z_t dt \cdot \int_0^T \tilde{z}_t dt\right] = E\left[\int_0^T \int_0^T z_u \tilde{z}_w dudw\right] = \int_0^T \int_0^T E[z_u \tilde{z}_w] dudw.$$

We can split the inside integral w.r.t. du in two regions $u \in (0, w)$ and $u \in (w, T)$. Noting that

$E[z_u \tilde{z}_w] = \rho \cdot \min[u, w]$, the expression I evaluates to:

$$I = \int_0^T \int_0^w \rho u dudw + \int_0^T \int_w^T \rho w dudw = \frac{1}{3} \rho T^3.$$

¹² That level depends on all the other parameters through the first-order condition (17). In the example, it is slightly negative as can be seen from Table I.

as θ increases to 0.050, the agent's profits suffer greater variability, $Std(\Pi) = 0.0867 \gg 0.0264$. He compensates, by extending his sales over time to $T^* = 0.36772$, in order to minimize the utility penalty for that. The profit VaR would have declined to $VaR(\Pi) = -0.2063$ ¹³, had he stayed at the no-longer-optimal $T = 0.08596$. Instead it drops only to $VaR(\Pi) = -0.1474$ at the optimal $T^* = 0.36772$.

V. Conclusion

This paper develops a model of the liquidation behavior of a venture capitalist or hedge fund manager holding a concentrated stake in an asset. At the optimum, the liquidator balances the desire to limit the risk of random market factors against the negative impact his sales have on the realized return. The position is large enough so that it is exposed not only to market, but also liquidity risks. The realizable value of the asset may be significantly smaller than that based on marked-to-market prices. We consider two impact function characterizations: power, uncorrelated with returns, and linear, correlated with returns. We obtain closed-form solutions for the first. We examine numerical solution to the second case which provides for a feedback loop between price dislocations and liquidity. The solutions illustrate the effect of the interaction of liquidity and market conditions on the liquidator's optimal strategy.

¹³ Not shown in Table I. Plug in the original $T=0.08596$ to the VaR Equation (16).

REFERENCES

- Almgren, Robert and Neil Chriss, 2000, Optimal Execution of Portfolio Transactions, Journal of Risk 3 (Winter 2000/2001), 5-39.
- Almgren, Robert, 2001, Optimal Execution with Nonlinear Impact Functions and Trading-Enhanced Risk, University of Toronto Working Paper
- Bangia Anil, Francis X. Diebold, Til Schuermann, and John D. Stroughair, 1999, Liquidity on the outside, Risk (June).
- Bangia Anil, Francis X. Diebold, Til Schuermann, and John D. Stroughair, 1999, Modeling Liquidity Risk, With Implications for Traditional Market Risk Measurement and Management, Working Paper 99-06, Financial Institutions Center, The Wharton School.
- Bank of International Settlements, 1996, Amendment to the Capital Accord to Incorporate Market Risks, January 1996, updated April 1998
- Basak Suleyman and Alex Shapiro, 2001, Value-at-Risk-Based Risk Management: Optimal Policies and Asset Management, Review of Financial Studies 14(2), 371-406.
- Bertsimas, Dimitris and Andrew W. Lo, 1998, Optimal Control of Execution Costs, Journal of Financial Markets 1, 1-50.
- Chan, Louis K.C. and Josef Lakonishok, 1995, The behavior of stock prices around institutional trades, Journal of Finance 50, 1147-74.
- Dubil, Robert, 2001, A liquidity-adjusted Value-at-Risk model based on a market-specific transaction clock, Working Paper, U. of Connecticut.
- Dubil, Robert, 2002, The risk-return tradeoff when converting wealth into cash, Working Paper, U. of Connecticut.
- Glosten, Lawrence R. and Lawrence E. Harris, 1988, Estimating the Components of the Bid-Ask Spread, Journal of Financial Economics 21, 123-142.
- Holthausen, Robert W., Richard W. Leftwich, and David Mayers, 1987, The effect of large block transactions on security prices, Journal of Financial Economics 19, 237-267.
- Holthausen, Robert W., Richard W. Leftwich, and David Mayers, 1990, Large-block transactions, the speed of response, and temporary and permanent stock-price effects, Journal of Financial Economics 26, 71-95.
- Huang, Chi-fu and Robert Litzenberger, 1988, Foundations for Financial Economics (North Holland).
- Jarrow, Robert and Ajay Subramanian, 1997, Mopping up Liquidity, Risk (December).
- Kyle, Albert S., 1985, Continuous Auctions and Insider Trading, Econometrica 53(6), 1315-1335.
- O'Hara, Maureen, 1997, Market Microstructure Theory (Blackwell Publishers, Inc.).

Table I.

Table I. Optimal Liquidation Times, the Expected Value, Variance and VaR of the Profit as Functions of Θ and ρ

Θ	Optimal T					E[Π]			Std[Π]			VaR[Π]		
	$\rho = -0.99$	$\rho = -0.5$	$\rho = 0$	$\rho = 0.5$	$\rho = 0.99$	$\rho = -0.99$	$\rho = -0.5$	$\rho = 0$	$\rho = -0.99$	$\rho = -0.5$	$\rho = 0$	$\rho = -0.99$	$\rho = -0.5$	$\rho = 0$
-0.050	0.33627	0.35359	0.36165	0.36772	0.37266	-0.0051	-0.0048	-0.0047	0.0071	0.0501	0.0708	-0.0168	-0.0872	-0.1213
-0.030	0.20380	0.22547	0.23520	0.24239	0.24818	-0.0084	-0.0076	-0.0073	0.0055	0.0390	0.0551	-0.0175	-0.0718	-0.0980
-0.010	0.07462	0.10636	0.11805	0.12624	0.13260	-0.0230	-0.0161	-0.0145	0.0040	0.0247	0.0342	-0.0296	-0.0568	-0.0707
-0.003	0.07093	0.08122	0.08783	0.09296	0.09714	-0.0242	-0.0211	-0.0195	0.0167	0.0223	0.0263	-0.0516	-0.0577	-0.0628
-0.002	0.07486	0.08080	0.08541	0.08920	0.09246	-0.0229	-0.0212	-0.0201	0.0195	0.0229	0.0256	-0.0550	-0.0588	-0.0622
-0.001	0.07902	0.08164	0.08390	0.08596	0.08782	-0.0217	-0.0210	-0.0204	0.0223	0.0238	0.0252	-0.0584	-0.0601	-0.0618
0.000	0.08338	0.08338	0.08338	0.08338	0.08338	-0.0206	-0.0206	-0.0206	0.0250	0.0250	0.0250	-0.0617	-0.0617	-0.0617
0.001	0.08782	0.08596	0.08390	0.08164	0.07902	-0.0195	-0.0199	-0.0204	0.0276	0.0264	0.0252	-0.0649	-0.0634	-0.0618
0.002	0.09246	0.08920	0.08541	0.08080	0.07486	-0.0185	-0.0192	-0.0201	0.0301	0.0280	0.0256	-0.0680	-0.0653	-0.0622
0.003	0.09714	0.09296	0.08783	0.08122	0.07093	-0.0176	-0.0184	-0.0195	0.0325	0.0297	0.0263	-0.0711	-0.0672	-0.0628
0.010	0.13260	0.12624	0.11805	0.10636	0.07462	-0.0129	-0.0136	-0.0145	0.0473	0.0414	0.0342	-0.0907	-0.0816	-0.0707
0.030	0.24818	0.24239	0.23520	0.22547	0.20380	-0.0069	-0.0071	-0.0073	0.0777	0.0675	0.0551	-0.1348	-0.1181	-0.0980
0.050	0.37266	0.36772	0.36165	0.35359	0.33627	-0.0046	-0.0047	-0.0047	0.0999	0.0867	0.0708	-0.1689	-0.1474	-0.1213

Figure 1. Temporary and permanent price effects of a seller-initiated block.

