

12-2003

A Simple Utility Approach to Private Equity Sales

Robert Dubil
San Jose State University

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

Recommended Citation

Dubil, Robert (2003) "A Simple Utility Approach to Private Equity Sales," *Journal of Entrepreneurial Finance and Business Ventures*: Vol. 8: Iss. 1, pp. 103-110.
Available at: <http://digitalcommons.pepperdine.edu/jef/vol8/iss1/7>

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 administrator of Pepperdine Digital Commons. For more information, please contact paul.stenis@pepperdine.edu.

A Simple Utility Approach to Private Equity Sales

Robert Dubil⁺
San Jose State University

The paper examines the liquidity risk of a private equity firm that decides to dispose of a large holding in its portfolio. As the sale takes time, it requires a careful balancing act of the exposure to the fluctuations in the market value of the investment against the large sale-induced price depression. A mean-standard deviation utility framework is an appealing decision tool for optimizing protracted asset dispositions. The firm maximizes the expected profit from the sale strategy net of the price concession minus a penalty function for exposure to the price risk, with the penalty weight related to a loss confidence interval.

Introduction

Economists and finance researchers do not distinguish between asset wealth and cash wealth. The two are fungible: assets can be sold for cash at a prevailing market price. The only problem left to solve then is the maximization of wealth through by choosing from a set of “efficient” asset allocations within a portfolio. A private equity firm maximizes the total value of its holdings by selecting assets with growth potential. But when the growth is realized, it faces another problem of converting the paper wealth into disposable cash. The conversion is costly and risky. It takes time during which the value of the paper wealth fluctuates. “Optimal” liquidation requires one to solve a dynamic utility maximization similar to the efficient frontier setup used in asset allocation decisions.

This article defines “efficiency” in the private equity liquidation context and provides an intuitive, theoretically sound, tool for making “efficient” paper-to-cash conversion decisions. We adapt the Value-at-Risk (VaR) methodology used by banks to risk-manage their trading portfolios to the specific problem under consideration. We ask the private equity holder to reveal his risk tolerance by choosing the probability of loss he is willing to accept. That allows us to prescribe the “best” liquidation path which maximizes the expected sales proceeds adjusted for

⁺ **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 an 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), a quantitative researcher at Nomura (1990-92) and JP Morgan (1989-90). His recent articles covering liquidity risks and banking regulation can be found in the Journal of Applied Finance, Financial Services Review and the Journal of Investing.

the risk of that loss. The approach leads to answers to questions like: how long should the sales program take, or what amount should be sold in each transaction.

Some theoretical solutions are provided in Dubil (2002). Here we merely guide the reader through the process of building a decision-making tool based on the microeconomic theory of risk preferences. In a related paper, Dubil (2003) considers the personal finance problem of an executive or a trustee with accumulated vested stock. We follow a similar line of thought, but our case is restricted to that of a venture capital or a private equity firm: buys and sells of illiquid assets are not one-off transactions at retirement, but are the bread-and-butter of the firm's business. They are closely monitored by other investors and often have significant adverse impact on the realized value of the holdings. They are perceived as trades by insiders. The use of VaR as the risk-return prism is more easily justified here than in the personal finance context. Tail risk as defined in VaR may be seen as the cost of running the business. On an ongoing basis, banks set capital against that risk, and so should private equity firms. Standard investment performance measures assume that a seller's actions do not affect his realized return, i.e. he is merely a "price-taker". Our private equity holder is not. However, he is "rational" in the standard microeconomic sense. He chooses only from efficient investments, i.e. those with lower risks for a given expected return the way a risk-averse agent selects an efficient frontier portfolio in the CAPM. Unlike in the perfect capital markets setup, the private equity firm holds a significant portion of an asset and its own actions affect the market price formation process. Although it chooses rationally, its return-to-risk reward does not simply depend on market-driven returns, but also on its own actions over time which affect the privately realized returns. The subject of this paper is to model the risk preferences for this situation in order to be able formulate an efficient liquidation strategy.

I. Liquidation "efficiency"

Liquidation can be broadly defined as a trading strategy over time. As such, from a utility theory perspective, it is no different from any dynamic problem of portfolio choice. In order to use the efficient frontier machinery, let us describe in more detail the financial behavior of the liquidating private equity firm. The firm's total liquidated value is going to be a function of two random variables: the fluctuating market price and the sale concession. While the sale program is being executed the market price of the asset can move adversely independent of the firm's actions. If the firm sells too quickly, it may cause additional deterioration of the realized return.

The firm can sell all at once or more likely choose a multiple-sale strategy. The asset price fluctuates randomly from sale to sale. The firm's sales have random impacts on the realized returns, which depend on the sizes of the transactions. To relate risk to expected return, we can define our firm's or agent's risk preferences in a form of a utility function. Huang and Litzenberger (1988) show that, if we adopt a HARA utility class coupled with some mild distributional assumptions, we can turn directly to Markowitz's (1959) notion of the mean-variance efficient frontier. Within that formulation, a strategy is considered efficient, if it is a solution to the quadratic programming maximization of expected terminal wealth (profit or return) for a given level of risk (defined as the variance of the wealth).

For our private equity agent, a strategy s may be simply defined as a sequence of stock sales quantities, and the sales profit $\Pi(s)$ as the sum of the revenues in dollars. Alternatively, if we denote as s_t the proportion of the original amount still invested in the stock, where t runs

from zero, the start, to T , the end of the liquidation program, then, by definition, $s_0 = 1$ and $s_T = 0$, and $-ds_t$ is the quantity sold in the interval dt . We can denote the cumulative return, relative to the start of the liquidation program, on each sale as r_t . Then the profit can be defined as the return realized on all the sales in strategy $s = \{s_t\}$ as

$$\Pi(s) = \int_0^T r_t(-ds_t) \quad (1)$$

In discrete time, the integral sign is replaced by summation. Ex-post, $\Pi(s)$ is a known total value realized from the sales. Ex-ante, it is a random variable, as we do not know at what prices (or returns) the agent will sell. He maximizes the expected profit minus a penalty function for risk, i.e. he seeks to find a solution to the following optimization:

$$\max_s E[\Pi(s)] - \lambda V[\Pi(s)] \quad (2)$$

where λ is a ‘shadow price’ of risk, $E[\square]$ denotes the expectations and $V[\square]$ the variance operator. If we use (2) in a quadratic programming setup, where the expected profit is maximized for each level of the variance, then the solution set forms a familiar frontier in the mean-variance space.

Banks use Value-at-Risk as a single summary statistic of how much money they can lose in an unlikely left-tail event defined by a confidence interval percentage. Let us use the same concept, but in a reverse way. The private equity firm specifies what probability of “loss” it is willing to accept in the liquidation and we will solve for the strategy that minimizes that loss. Formally, we define the VaR of a strategy as a critical (most likely negative) profit value, Π^* , for which the probability of the profit, Π , falling below that critical value, Π^* , is equal to a given quantity $\Phi(\alpha)$. The latter is equal to one minus the chosen confidence interval. For example, if the confidence level is 95%, then $\Phi(\alpha) = 0.05$ and the critical loss level Π^* is exceeded only 5% of the time. The VaR concept is an extremely simple predictive measure. If the firm holds \$10 million worth of a stock with the expected return of 20% and the standard deviation of 30% over a year, then, using a standard normal $\alpha = 1.645$ corresponding to the 95% one-sided confidence interval, the VaR of the position is $20 - 1.645 \cdot 30 = -29.35\%$, or $\Pi^* = \$10mm \cdot (1 - 0.2935) = \$7,065,000$ in dollars. The expected value of the position is \$12 mm. However, one can only be sure with 95% probability that the value of the position after one year will exceed \$7,065,000. This is a static view used by most banks. The portfolio is held constant and is simply subject to random market fluctuations. Let us take a dynamic view. While market fluctuations take place, the firm can alter the portfolio, perhaps by selling a portion of it. It is not able to change the distribution of the market returns, but it is able to shift and contort the resultant profit distribution in such a way that at the same 95% confidence interval the value of the portfolio will exceed \$9,000,000. See Fig. 1 for illustration.

Once the agent defines his profit as a function of the strategy, he is able to make statistical predictions that his strategy, with 95% probability, will yield at least Π^* dollars. If the returns on the asset are normally distributed, then profit Π will also be normally distributed. A

table of the cumulative standard normal $\Phi(\cdot)$ can be used to find $\alpha = 1.645$ corresponding to a 5% tail. The VaR, Π^* , can be defined *explicitly* as:

$$\Pi^* = E[\Pi(s)] - \alpha \sqrt{V[\Pi(s)]} \quad (3)$$

We further assume that, at the chosen confidence parameter, the agent seeks a strategy that maximizes that critical profit level, Π^* , i.e. he solves:

$$\max_s E[\Pi(s)] - \alpha \sqrt{V[\Pi(s)]} \quad (4)$$

The difference between (4) and (2) is the substitution of the standard deviation for the variance and the fact that α is not an implicit shadow price (related to the concavity of the utility function), but it is determined from the firm's chosen confidence level. It reflects an implicit set of risk preferences linear in standard deviation instead of in variance. The set of solutions to (4) forms an efficient frontier in a mean-standard deviation space.

II. Optimal liquidation under VaR

A seller of a concentrated asset, whether using (2) or (4) as his optimization specification, would start by defining the liquidation profit as a function of his strategy. Suppose he owns $W_0 = \$10$ million of the stock and decides to sell over a 12-month period. His strategy is defined as a vector sequence of holding proportions still outstanding $s = \langle s_0 = 1; s_1, s_2, \dots, s_{11}, s_{12} = 0 \rangle$, $-\sum_{i=1}^{12} \Delta s_i = 1$. For example s might be $\langle 100\%, 95\%, 90\%, 80\%, 70\%, \dots, 20\%, 15\%, 10\%, 0\% \rangle$. The proportion decrements add up to one since by the end he sells his entire position. Let us denote the cumulative returns over the 1-, 2-, ..., 12-month period as $\tilde{R} = \langle \tilde{R}_1, \tilde{R}_2, \dots, \tilde{R}_{12} \rangle$. Ex-post, these returns and the return on the entire strategy will be known. Ex-ante, the seller does not know what the returns will be, but he chooses his strategy s to get the 'best' expected return-risk combination. If his sales did not affect the realized return, then his trading profit might be defined as:

$$\Pi(s) = W_0 \sum_{i=1}^{12} (-\Delta s_i) \tilde{R}_i \quad (5)$$

As his equity stake is concentrated, his sales affect adversely the realized returns by amounts related to the size of the sales and the market returns, $\theta(s) = \langle \theta(s, R_1), \dots, \theta(s, R_{12}) \rangle$. In general, the impact function $\theta(\cdot)$ can be quite complicated as the impact of subsequent sales may depend on prior sales. It can also be quite simple, e.g. $\theta(\cdot)$ could be linear in the sale amounts $-\Delta s_i$ to reflect the fact that the more the agent sells each time the more he depresses the price. If the returns and the sales discounts are normally distributed, then his trading profit defined as

$$\Pi(s) = W_0 \sum_{i=1}^{12} (-\Delta s_i) \left[\tilde{R}_i - \theta(s, \tilde{R}_i) \right] \quad (6)$$

will also be normally distributed. All one has to do now to complete the trading model is to parameterize more explicitly the discount function $\theta(s)$ in order to compute the mean and the standard deviation of the profit. The simplest way would be to define each discount as a proportion of the realized sale value itself, i.e. $\tilde{\theta}_i = \delta(-\Delta s_i) \tilde{R}_i$. In that case, $\delta = 0.1$ would mean that you give up 10% of each realized return in order to liquidate a large quantity. For this and other more complicated examples, see Bertsimas and Lo (1998), Almgren and Chriss (2000), and Dubil (2002). The model is not yet complete. We still need to choose a utility function, i.e. we must specify the tradeoff of future return to risk. In the previous section, we offered two alternatives (2) and (4). Eq. (2) implicitly surrenders the choice of weight of the risk penalty λ to a particular class of utility functions, e.g. HARA with constant relative risk aversion. Eq. (4)'s maximization specifically adopts VaR as the risk-return tradeoff. The agent simply sets the confidence interval level, α , to a numerical value which reveals his worst-loss tail probability, $\Phi(\alpha)$, he is willing to accept. Then if the chosen impact function is simple, he solves (4) explicitly. If not he does it numerically, perhaps by creating a histogram of potential profit outcomes given different strategy combinations and reading off the 5% tail outcome, similarly to Figure 1.

III. The practical appeal of VaR for private equity firms

Most risk-return models in finance rely on the assumption of perfect capital markets in partial or general equilibrium framework. Agents maximize their own utilities while taking prices as given. Agent's optimal strategies affect prices only collectively through the market clearing condition. This is not the case in market microstructure studies where what is examined is the price formation process itself. Price behavior is explained by the interaction of disparate objectives of different groups of agent, perhaps with asymmetric access to information. The case of a private equity firm is akin to that last case, though even more specialized. The firm's objective itself must be specified. That objective needs to be consistent with the predictions of the general theory and the observed price behavior in the market place. However, the theory does not offer any suggestions as to the choice of the preference specifications correct for this situation. At the same time, a concentrated asset holder may need a quantitative tool to assess the efficiency of available strategies and to select an optimal one. The VaR utility function is a theoretically sound and yet a very practical choice for examining the risk-return tradeoff in a large asset liquidation program. It is expressed in dollars of potential loss as opposed to a score based on statistical terms of variances, betas, etc. It takes into account the fact that, with undiversifiable risks, the holder of the asset has already determined the fundamental riskiness of his investment when he decided to take a position in the asset. During the final sale, however, he does not view upside volatility the same way as downside volatility. Once he has decided to sell his concentrated holdings, he worries a lot more about losing his accumulated investment return than about potential additional market gains. The VaR requires only one parameter. The seller can easily express his risk preferences by choosing the confidence level for a loss he is willing to

tolerate (choosing the constant relative risk aversion function might be more difficult). All he needs to do is to state the amount of loss he can accept and the probability of that loss. The latter defines the confidence parameter used in VaR and the optimization process solves for the strategy that minimizes the potential loss. Very risk-averse sellers, concerned with very low probability loss events, may choose small tail regions (i.e. high confidence levels). Regulated banks promulgate 95% and 99% VaR numbers to constrain their traders' behavior, but they supplement them with other position limits. A private equity firm with concentrated management and only a few large, and closely monitored, positions is typically free from agency problems of a bank whose traders may have their own objectives. However, the exit strategy is a lot more important for the private equity firm as it may affect its return on assets to a greater degree. The success or failure of the exit strategies for different asset stakes needs to be quantitatively monitored. For that purpose, the firm may choose more unorthodox confidence percentages, like 50%.

IV. Conclusion

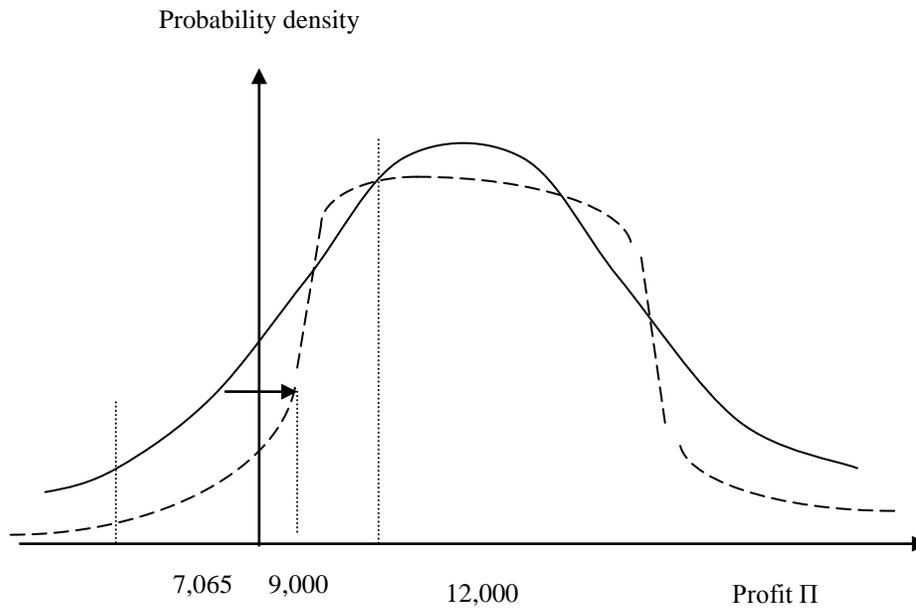
Exit strategies for private equity firms are normally examined from the perspective of optimal timing of sales after the holding period restrictions for insiders have expired. In this paper, we proposed a risk-return framework for analyzing the liquidity aspect of the sale of a concentrated asset stake. Typically, even past the holding period, the market for the asset is thin relative to the size of the holding of the private equity firm. Liquidation, even partial, may significantly affect the realized return on the investment. The VaR utility properly used is shown to be an excellent and simple metric with which to gauge the efficiency and optimality of the exit strategy.

REFERENCES

- Almgren, Robert and Neil Chriss, 2000, Optimal Execution of Portfolio Transactions, Journal of Risk 3 (Winter 2000/2001), 5-39.
- Bertsimas, Dimitris and Andrew W. Lo, 1998, Optimal Control of Execution Costs, Journal of Financial Markets 1, 1-50.
- Dubil, Robert, 2002, How quickly should you sell your vested stock?, Financial Services Review 11(1), 65-84.
- Dubil, Robert, 2003, Wealth to cash conversion in personal finance, Journal of Investing, forthcoming in Summer 2003.
- Huang, Chi-fu and Robert Litzenberger, 1988, Foundations for Financial Economics (North Holland).
- Markowitz, Harry M., 1959, Portfolio Selection: Efficient Diversification of Investments, New York: John Wiley & Sons.

Figure 1

VaR of a \$10 million position with mean 20% and st.dev. 30%. Amt. in thousands.



Solid and dashed lines represent two different strategies. 5% probability mass under the solid line to the left of 7,065. 5% probability mass under the dashed line to the left of 9,000.