Mean-Variance Trade-offs in Supply Contracts[1]

Victor Martínez-de-Albéniz and David Simchi-Levi

Operations Research Center, MIT, U.S.A.

Abstract

We study the trade-offs faced by a manufacturer signing a portfolio of option contracts with its suppliers and having access to a spot market. The manufacturer incurs inventory risk when purchasing too many options and spot price risk when buying too few. We quantify these risks for a single selling period by studying mean and variance of the profit for a given contract. Similarly to Markowitz's Capital Asset Pricing Model, we characterize the efficient portfolios of options that the manufacturer must hold in order to obtain dominating mean-variance pairs. Among these, we emphasize the maximum expectation portfolio, obtained by solving the classical newsvendor problem, and the corresponding minimum variance portfolio. We provide bounds on the efficient frontier. Finally, we characterize the upper-level sets of a mean-variance utility function and prove that they are connected. Hence, a *greedy method* will find the portfolios on the efficient frontier.

1 Introduction

For most manufacturers, effective supply chain strategies require careful consideration of procurement decisions. These decisions need to take into account not only the many aspects of the manufacturing process, but also the balance between overstocking and shortages. Traditionally, the academic literature has modeled this trade-off by introducing, in particular, the newsvendor model: in this framework, it is assumed that the consequence of a shortage is lost revenues while the impact of overstocking is the production of an item that must be disposed or sold at a loss. This tool has allowed manufacturers to quantify stocking and purchasing decisions.

One of the drawbacks of this approach is that it is based on the assumption that decision makers, the manufacturers, are risk neutral and hence optimize their expected profit. However, recent experiences, such as Cisco's $2.5 billion inventory write-off in April 2001, suggest that risk matters, as illustrated in [2]. Thus, the challenge for the operations management community is to develop risk management models to complement current purchasing tools. This is especially important when, together with *inventory and shortage risk*, manufacturers face *spot price risk*. That is, by committing in advance to a given contract, manufacturers take the risk of not being able to take advantage of a low spot market price; similarly, when they do not secure enough supply in advance, they take the risk of paying a high spot market price.

This is especially obvious when one looks at the way the financial industry addresses similar issues. Since the emergence of the Capital Asset Pricing Model (CAPM), financial theory has

[1]Research supported in part by the Center of eBusiness at MIT, ONR Contracts N00014-95-1-0232 and N00014-01-1-0146, and by NSF Contracts DMI-9732795, DMI-0085683 and DMI-0245352.

developed models that quantify the trade-offs between risky returns with higher average returns and risk-free returns. Thus, the market is risk-averse in the sense that it requests larger average profits when these profits are uncertain. The first financial models in the 1960s were quickly transferred into industry, triggering the development of mutual funds. A company's returns risk could be evaluated through its "beta", i.e. its correlation with the market index variations. With these methods, investors could manage market risk while seeking high returns. This is in contrast with the way managers make purchasing decisions.

The objective of this paper is to present a model where risk, measured by the variance of profit, is considered together with expected profit. We assume that a variety of supply options, including firm-commitment contracts, are available to the manufacturer, with different reservation and execution unit prices, and these can be purchased in advance of knowing demand and spot market price. We define, in a single-period setting, the set of efficient mean-variance profit pairs, similarly to the CAPM, as a function of the amounts purchased for every option.

We emphasize the challenges of a newsvendor model compared to traditional financial model. Indeed, in finance, portfolio decisions can be reversed in the sense that investors can sell whatever asset they hold at any time; in manufacturing, however, once inventories are ordered, they are written off at high cost, i.e., sold at a low salvage value, since the components are somehow engineered for a particular manufacturer, see the Cisco case alluded to before.

Our analysis of the trade-off between profit mean and variance highlight two special portfolios: a maximum expected profit portfolio, obtained by the classical risk-neutral newsvendor model, and a minimum profit variance portfolio. The characterization of this second portfolio is new, and is similar to the minimum variance portfolio in the CAPM without risk-free assets. Between these two "extreme" portfolios, we find a set of mean-variance efficient portfolios, for which we give bounds. The main difference between our model and the financial CAPM is that in the later model the efficient frontier is typically unbounded, while in our newsvendor model we show that the efficient frontier is bounded. In addition, we characterize the upper-level sets of mean-variance utility functions and prove that they are connected. Hence, a *greedy method* will find the portfolios on the efficient frontier.

This research is a natural extension of recent development in the analysis of portfolio contracts for supply chains. Martínez-de-Albéniz and Simchi-Levi [9] first presented a general multi-period framework where portfolios could be optimized in terms of expected profit. These techniques have been applied in a number of high-tech firms, and will no doubt be increasingly popular as a way to better react to uncertain demands and spot market prices.

The literature on risk in supply contracts is quite limited. An exception is the work of Lau [7], who proposes alternative optimization objectives for the newsboy problem, and in particular the objective of maximizing the probability of achieving a given level of profit. For this purpose, Lau presents formulas for the moments of the profit for a general demand distribution.

The most common approach to deal with risk in an industrial setting has come from the financial world. The objective in this stream of literature is to analyze a stochastic inventory model together with trading in the market, and to hedge the inventory project with a financial operation. The single period problem has been studied by Anvari [1], followed by Chung [4]. They apply the financial CAPM theory to provide a market valuation of an inventory project, and derive an optimal ordering quantity provided that the market index and the demand follow a bivariate normal distribution. In a more general setting, a lot of research has been done on the so-called mean-variance hedging problem. Duffie and Richardson [5], Schweizer [13], Gouriéroux, Laurent and Pham [6], for instance, have analyzed the problem of dynamically hedging some asset with the available assets in the market. The objective is to minimize the final deviation between the asset and the hedge with respect to some stochastic metric. Caldentey and Haugh [3] have applied this approach to the newsvendor problem, thus generating a dynamic trading strategy in the financial markets together with a single-period inventory decision. We should point out that this research does not consider the intrinsic risk of the inventory project, and uses the correlation between customer demand and the market returns to reduce risk. In contrast, the purpose of our paper is to describe the intrinsic risk of an inventory decision, as a function of inventory and component price risk, without the risk associated with financial markets.

We start by presenting in Section 2 the traditional financial formulation, i.e., following the CAPM, of the profit as a function of the amounts of options purchased by the manufacturer. In this part, overstocked inventory can be returned at the spot market price. We remove this assumption in Section 3, where we present the newsvendor model for a portfolio of options. We analyze the model in Section 4 and define the maximum expectation and minimum variance portfolios, both for the single and multiple suppliers cases, together with bounds on the efficient portfolios. We then characterize the level sets of mean-variance objectives and show that a greedy algorithm always finds the optimal portfolio.

Finally, we discuss in Section 5 the challenges of a multi-period extension and provide concluding remarks in Section 6.

2 Hedging Earnings with Financial Options

Consider a firm that sells a product at a predetermined price p. To manufacture the product the firm uses a component that can be found in a spot market at a spot price S, and production only takes place when $p > S$. The firm's profit is thus

$$\Pi^0 = (p - S)^+ Q,$$

and it depends on a stochastic demand, Q, and a stochastic spot price, S.

Assume now that this firm is able to sign contracts in advance with suppliers. This enables the firm to become less sensitive to spot price fluctuations. For example, buying supply through an forward contract can reduce the exposure of the firm to the spot price.

We assume that n suppliers are available, each offering a supply option contract for a reservation fee of v_i per unit and an execution fee of w_i per unit, $i = 1, \ldots, n$. Traditionally, these options have been used as a side-bet with no real consequences: they are mechanisms that arrange side-payments as a function of the spot price. That is, in this setting the firm is using the contracts not only to provide raw-material for production, but also to sell back excess capacity to the spot market.

Under this framework, the profit of the buyer, denoted by the superscript F (for financial), can be written as

$$\Pi^F = (p - S)^+ Q + \sum_{i=1}^{n} \Big\{ (S - w_i)^+ - v_i \Big\} x_i, \tag{1}$$

where $x_i \geq 0$ is the amount of options purchased from supplier i, $i = 1, \ldots, n$. These options can therefore be treated like assets with returns $(S - w_i)^+ - v_i$ but no up-front cost. This setting is equivalent to the Capital Asset Pricing Model (CAPM), where x_i is the amount of stock i purchased and $(S - w_i)^+ - v_i$ is the return of the stock.

In Markowitz's CAPM, a mean-variance optimal trade-off is found, the so-called *efficient frontier*. One can similarly determine an optimal mean-variance frontier in our case. Specifically,

$$\mathbb{E}\Pi^F = \mathbb{E}\Big[(p - S)^+ Q\Big] + \sum_{i=1}^{n} \mathbb{E}\Big[(S - w_i)^+ - v_i\Big] x_i, \tag{2}$$

and

$$Var\Pi^F = \begin{cases} Var\Big[(p - S)^+ Q\Big] \\ + \sum_{i=1}^{n} Var\Big[(S - w_i)^+ - v_i\Big] x_i^2 \\ +2 \sum_{i=1}^{n} Covar\Big[(S - w_i)^+ - v_i, (p - S)^+ Q\Big] x_i \\ +2 \sum_{i<j} Covar\Big[(S - w_i)^+ - v_i, (S - w_j)^+ - v_j\Big] x_i x_j. \end{cases} \tag{3}$$

The mean-variance optimal trade-off curve is defined by the points $(Var\Pi, \mathbb{E}\Pi)$ found in the optimization programs

$$\min_{x \geq 0} \; Var\Pi \text{ subject to } \mathbb{E}\Pi \geq \mu, \tag{4}$$

for $\mu \in \mathbb{R}$, or equivalently

$$\max_{x \geq 0} \; \mathbb{E}\Pi \text{ subject to } Var\Pi \leq \sigma^2, \tag{5}$$

for $\sigma^2 \in \mathbb{R}_+$, or also,

$$\max_{x \geq 0} \; \mathbb{E}\Pi - \lambda Var\Pi, \tag{6}$$

for $\lambda \in \mathbb{R}_+$.

Figure 1 shows the trade-off curve (efficient frontier) between expected profit and variance of profit for the model discussed in this section. In this setting, we do not have a risk-free asset, which is often used in the financial literature. The "mutual fund" theorem, described in Sharpe [14], Mossin [12] or Merton [10], does not hold, since there is no riskless bond; there is no market price for systematic risk. Instead, we can work with the efficient frontier, and have the manufacturer (and every other buyer in this market) choose its own efficient portfolio.

3 A Newsvendor Model

Of course, the financial model described above does not capture some of the constraints associated with real world purchasing practices. These constraints are either contractual, i.e., the supplier does not allow the buyer to resell the component, or operational, i.e., design constraints that are typically associated with product specification requiring specializing the component to the buyer's needs. This type of constraint implies that it is difficult or costly for the buyer to resell the components back to the market.

To capture these issues, we propose the following model which is a generalization of the newsvendor problem to portfolio contracts. Details can be found in Martínez-de-Albéniz and Simchi-Levi [9]. The manufacturer can not sell back excess supply to the spot market, and thus we we restrict the use of capacity to serving the demand faced by the manufacturer. Hence, unused capacity is lost, in the sense that the manufacturer is unable to exercise the options and sell the corresponding supply units to the spot market, at the spot price. We thus rule out the possibility of the manufacturer becoming a trader and gaining financial advantage from speculation.

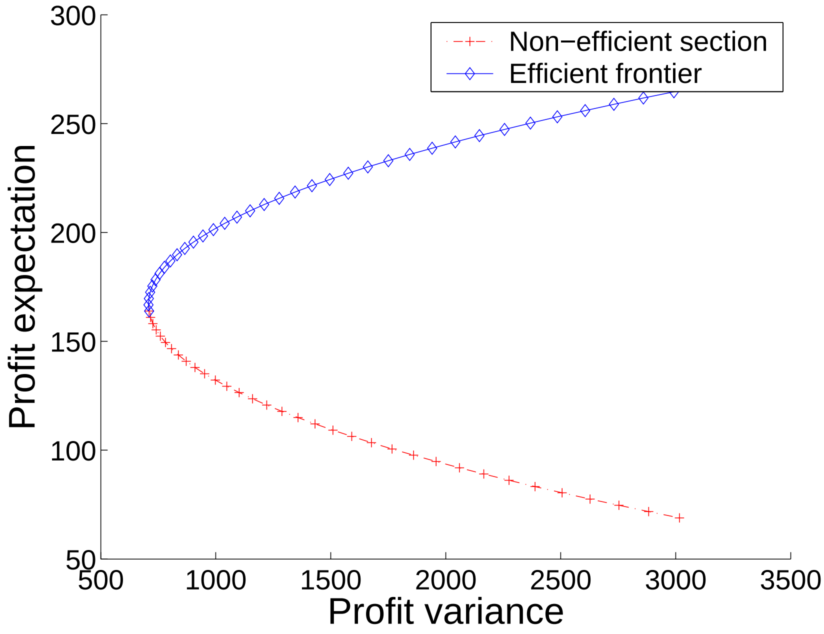

Figure 1: Simulation of mean-variance curve for a single supplier as a function of the amount bought from this supplier

Assuming without loss of generality that $w_1 \leq \ldots \leq w_n$, and letting $v_{n+1} = 0$, $w_{n+1} = p$, we define for $i = 1, \ldots, n$, $y_i = \sum_{j=1}^{i} x_j$, and

$$Z_i = \left\{ \min(S, w_{i+1}) - \min(S, w_i) \right\}, \quad Z_0 = (p - S)^+ = p - \min(S, p).$$

The manufacturer's profit, denoted now with the superscript N (for newsvendor), can be written as

$$\Pi^N = p \min(Q, y_n) + (p - S)^+ (Q - y_n)^+ - \sum_{i=1}^{n} v_i x_i - \sum_{i=1}^{n} \min(S, w_i) \min \left\{ x_i, (Q - y_{i-1})^+ \right\}$$

Using the relationship $(x - y)^+ = x - \min(x, y)$ for any x, y, we can reformulate

$$\begin{aligned} p \min(Q, y_n) + (p - S)^+ (Q - y_n)^+ &= p \min(Q, y_n) + \{p - \min(S, p)\}\{Q - \min(Q, y_n)\} \\ &= Z_0 Q + \min(S, p) \min(Q, y_n), \end{aligned}$$

and

$$\begin{aligned} \min \left\{ x_i, (Q - y_{i-1})^+ \right\} &= \min \left\{ y_i - y_{i-1}, (Q - y_{i-1})^+ \right\} \\ &= \min(Q, y_i) - \min(Q, y_{i-1}). \end{aligned}$$

We can plug these expressions into Π^N and observe that $\min(Q, y_n) = \sum_{i=1}^{n} \{ \min(Q, y_i) - \min(Q, y_{i-1})\}$. We thus obtain

$$
\begin{aligned}
\Pi^N = &\quad Z_0 Q + \min(S, p) \min(Q, y_n) - \sum_{i=1}^{n} v_i x_i - \sum_{i=1}^{n} \min(S, w_i)\{ \min(Q, y_i) - \min(Q, y_{i-1})\} \\
= &\quad Z_0 Q - \sum_{i=1}^{n} v_i(y_i - y_{i-1}) + \sum_{i=1}^{n} \{ \min(S, p) - \min(S, w_i)\}\{ \min(Q, y_i) - \min(Q, y_{i-1})\} \\
= &\quad Z_0 Q + \sum_{i=1}^{n} (v_{i+1} - v_i) y_i + \sum_{i=1}^{n} Z_i \min(Q, y_i).
\end{aligned}
\tag{7}
$$

Of course, a similar transformation shows that

$$\Pi^F = Z_0 Q + (S - p)^+ y_n + \sum_{i=1}^{n} (v_{i+1} - v_i) y_i + \sum_{i=1}^{n} Z_i y_i.$$

Evidently, the difference between the two profit functions stems from the assumption that in the newsvendor model the buyer is not able to sell excess supply to the spot market.

4 Properties of Supply Option Portfolios

To proceed with the mean-variance analysis of a portfolio, we make the following assumption.

Assumption 1 *The customer demand Q is independent of the spot market price S.*

The assumption thus implies that demand faced by the buyer does not drive spot market prices. This is typically the case when the buyer's product does not capture a large portion of the component's demand.

Assumption 2 *Q and S follow respectively distributions with p.d.f. $f_Q > 0$ (on $\mathbb{R}$) and f_S that are continuous and c.d.f. F_Q and F_S.*

In addition, we define $\overline{F}_Q = 1 - F_Q$ and $\overline{F}_S = 1 - F_S$.

The analysis of mean and variance of profit is similar to Equations (2) and (3):

$$\mathbb{E}\Pi^N = \mathbb{E}[Z_0 Q] + \sum_{i=1}^{n} \Big\{ \mathbb{E}[Z_i] \int_0^{y_i} \overline{F}_Q(q) dq - (v_{i+1} - v_i) y_i \Big\}, \tag{8}$$

and

$$
Var\Pi^N = \begin{cases} Var\Big[Z_0 Q\Big] + \sum_{i=1}^{n} Var\Big[Z_i \min(Q, y_i)\Big] \\ +2\sum_{i=1}^{n} Covar\Big[Z_i \min(Q, y_i), Z_0 Q\Big] \\ +2\sum_{i<j} Covar\Big[Z_i \min(Q, y_i), Z_j \min(Q, y_j)\Big]. \end{cases} \tag{9}
$$

As derived in Martínez-de-Albéniz and Simchi-Levi [9], the expected profit of the buyer is a concave function of the variables $(x_1, \ldots, x_n)$, or equivalently of a linear transformation of them, $(y_1, \ldots, y_n)$. In particular, using the independence of Q and S,

$$
\frac{d\mathbb{E}\Pi^N}{dy_i} = v_{i+1} - v_i + \mathbb{E}[Z_i]\overline{F}_Q(y_i). \tag{10}
$$

To understand the behavior of variance, define

$$
\begin{aligned}
A_i &= \frac{1}{2} Var\Big[Z_i \min(Q, y_i)\Big], \\
B_i &= Covar\Big[Z_i \min(Q, y_i), Z_0 Q\Big], \\
C_{ij} &= Covar\Big[Z_i \min(Q, y_i), Z_j \min(Q, y_j)\Big].
\end{aligned}
$$

Thus, the variance can thus be expressed as

$$
Var\Big[Z_0 Q\Big] + 2\sum_{i=1}^{n}(A_i + B_i) + 2\sum_{i<j} C_{ij},
$$

and hence,

$$
\frac{1}{2}\frac{dVar\Pi^N}{dy_i} = \frac{dA_i}{dy_i} + \frac{dB_i}{dy_i} + \sum_{j<i}\frac{dC_{ji}}{dy_i} + \sum_{i<j}\frac{dC_{ij}}{dy_i}.
$$

The formulas for such expressions are presented below, and are detailed in the appendix.

$$
\begin{aligned}
\frac{dA_i}{dy_i} &= \overline{F}_Q(y_i)\Big(\mathbb{E}[Z_i^2]y_i - \mathbb{E}[Z_i]^2 \int_0^{y_i} \overline{F}_Q(u)du\Big) \\
\frac{dB_i}{dy_i} &= \overline{F}_Q(y_i)\Big(\mathbb{E}[Z_0 Z_i](y_i + \frac{\int_{y_i}^{\infty} \overline{F}_Q(u)du}{\overline{F}_Q(y_i)}) - \mathbb{E}[Z_0]\mathbb{E}[Z_i]\mathbb{E}[Q]\Big) \\
\frac{dC_{ij}}{dy_i} &= \overline{F}_Q(y_i)\Big(\mathbb{E}[Z_i Z_j](y_i + \frac{\int_{y_i}^{y_j} \overline{F}_Q(u)du}{\overline{F}_Q(y_i)}) - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\int_0^{y_j} \overline{F}_Q(u)du\Big) \\
\frac{dC_{ij}}{dy_j} &= \overline{F}_Q(y_j)\Big(\mathbb{E}[Z_i Z_j]y_i - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\int_0^{y_i} \overline{F}_Q(u)du\Big)
\end{aligned} \tag{11}
$$

4.1 Single supplier case

In the case of a single supplier, we define the quantity y^E as the portfolio maximizing expected profit, obtained by solving the equation

$$\overline{F}_Q(y^E) = \frac{v}{\mathbb{E}[(S-w)^+]}. \tag{12}$$

Observe that y^E is well defined since $\overline{F}$ is decreasing. Of course, when $v = 0$, we select $y^E = \infty$, and when $v \geq \mathbb{E}[(S-w)^+]$, we choose $y^E = 0$.

Similarly, y^V is defined as the portfolio minimizing profit variance. Using Equation (11), we observe that the sign of the variance is the same as the right-hand side of the following equation. Thus, y^V is determined by solving

$$\begin{aligned} 0 = \;& \mathbb{E}[(S-w)^{+2}]y^V - \mathbb{E}[(S-w)^+]^2 \int_0^{y^V} \overline{F}_Q(u)du \\ &+\mathbb{E}[(p-S)^+(S-w)^+(y^V + \frac{\int_{y^V}^{\infty} \overline{F}_Q(u)du}{\overline{F}_Q(y^V)}) \\ &-\mathbb{E}[(p-S)^+]\mathbb{E}[(S-w)^+]\mathbb{E}[Q] \end{aligned} \tag{13}$$

The right-hand side of Equation (13) is increasing: its derivative is

$$\begin{aligned} &\mathbb{E}[(S-w)^{+2}]F_Q(y^V) + Var[(S-w)^+]\overline{F}_Q(y^V) \\ &+\mathbb{E}[(p-S)^+(S-w)^+]f_Q(y^V)\frac{\int_{y^V}^{\infty} \overline{F}_Q(u)du}{[\overline{F}_Q(y^V)]^2} \\ &> 0. \end{aligned}$$

In addition, we notice that the right-hand side of Equation (13) tends to infinity when $y \to \infty$ and is equal to $\mathbb{E}[(p-S)^+(S-w)^+]\mathbb{E}[Q] - \mathbb{E}[(p-S)^+]\mathbb{E}[(S-w)^+]\mathbb{E}[Q] \leq 0$ when $y = 0$, since $Covar[(p-S)^+, (S-w)^+] \leq 0$. This implies that y^V is well-defined. Moreover, this shows that the variance is quasi-convex as a function of y.

y	$< y^E$	$> y^E$	$< y^V$	$> y^V$
Expected Profit	$\uparrow$	$\downarrow$		
Variance of Profit			$\downarrow$	$\uparrow$

This implies the following proposition.

Proposition 1 *When the buyer maximizes its quadratic utility function, defined in Equation (6), it always purchases an amount in the interval $[y^E, y^V]$ (if $y^E \leq y^V$) or $[y^V, y^E]$ (if $y^E \geq y^V$).*

This proposition is illustrated by the crossed curve in Figure 2. We observe that in this case $y^V < y^E$. As y increases, starting from 0, variance decreases and expectation increases at first; then, when $y > y^V$, variance starts to increase while expectation keeps growing; finally, when $y > y^E$, variance still increases (but reaches a limit though) as expectation starts to decrease.

When compared to the benchmark, i.e. the financial model following the CAPM, we observe that the expected profit is always smaller for the same level of risk. Also, as the amount y bought from the supplier increases, variance is bounded from above in the newsvendor model but grows unbounded for CAPM. In other words, using the financial model in a situation where the newsvendor model is more appropriate may lead to a major underestimation of profit risk, for a given level of expected profit.

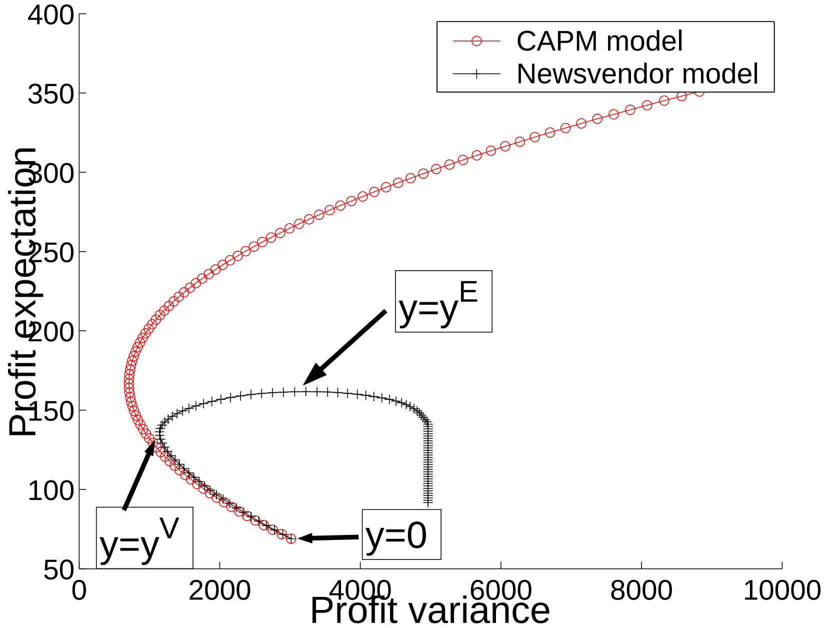

Figure 2: Comparison of the mean-variance curves for the different models: financial (circled line) and newsvendor (crossed line)

4.2 Multiple suppliers case

When multiple suppliers are available, the expected profit behaves in a similar way to the single supplier case. Indeed, in Martínez-de-Albéniz and Simchi-Levi [9] we show the following.

Proposition 2 *The expected profit is a strictly concave function of the amounts purchased from each contract. It attains a unique maximum for a portfolio* $\mathbf{y^E}$.

Unfortunately, the variance of profit behaves in a more complicated way. The following proposition provides some information about the first and second moments of the variance.

Proposition 3 *For each i, $i = 1, \ldots, n$, we have that*

$$\frac{dVar\Pi^N}{dy_i} = \overline{F}_Q(y_i)\Phi_i(\mathbf{y})$$

for some functions Φ_i. Moreover, we have that

$$H = \Big(\frac{d^2Var\Pi^N}{dy_i dy_j}\Big)_{(i,j)} \succ \begin{pmatrix} \ddots & & 0 \\ & -f_Q(y_i)\Phi_i(\mathbf{y}) & \\ 0 & & \ddots \end{pmatrix},$$

where $A \succ B$ means that the matrix $A - B$ is definite positive.

All the proofs are presented in the appendix.

We use this result to characterize the behavior of the profit variance. Let

$$F = \Big\{\mathbf{y}\Big| 0 \leq y_1 \leq \ldots \leq y_n\Big\}. \tag{14}$$

Define also the following notation.

Definition 1 *Consider a twice-differentiable function $f : F \to \mathbb{R}$, where F is defined in Equation (14). For any $I \subset \{1, \ldots, n\}$, let*

$$A_I = F \cap \Big\{\mathbf{y}\Big| y_i = y_{i-1} \text{ for } i \notin I\Big\}. \tag{15}$$

We can write $I = \{i_1, \ldots, i_m\}$, and thus define, for $0 \leq z_1 \leq \ldots \leq z_m$,

$$y_i(\mathbf{z}) = z_j \text{ for every } j \text{ such that } i_j \leq i < i_{j+1}.$$

Let

$$g(z_1, \ldots, z_m) = f\Big(\mathbf{y}(\mathbf{z})\Big).$$

*Let $\mathbf{y} \in A_I$. $\mathbf{y}$ is a I-*unconstrained critical point *of f if and only if for $j = 1, \ldots, m$,*

$$\frac{dg}{dz_j}\Big(\mathbf{y}(\mathbf{z})\Big) = 0.$$

Intuitively, an I-unconstrained critical point for a function f is a portfolio $\mathbf{y}$ such that the function f, restricted to A_I, has a critical point at $\mathbf{y}$.

Definition 2 *Given a function $f : F \to \mathbb{R}$, the* lower-level set *at c is*

$$\Big\{\mathbf{y} \in F | f(\mathbf{y}) \leq c\Big\},$$

and the upper-level set *at c is*

$$\Big\{\mathbf{y} \in F | f(\mathbf{y}) \geq c\Big\}.$$

Proposition 4 *For some I, let $\mathbf{y} \in A_I$ and assume that $\mathbf{y}$ is an I-unconstrained critical point of the variance. Then, portfolio $\mathbf{y}$ is a strict local minimizer of the variance in the set A_I.*

The previous result implies the following important qualitative result.

Proposition 5 *The lower-level sets of the variance of a portfolio are connected. That is, for any c, for any $\mathbf{y^0}, \mathbf{y^1}$, if $Var\Pi^N(\mathbf{y^0}) \leq c$ and $Var\Pi^N(\mathbf{y^1}) \leq c$, then there is a continuous path $\mathbf{y}(\cdot)$ such that $\mathbf{y}(0) = \mathbf{y^0}$, $\mathbf{y}(1) = \mathbf{y^1}$ and for all $t \in [0,1]$, $Var\Pi^N(\mathbf{y}(t)) \leq c$.*

Corollary 1 *Any portfolio that minimizes the variance locally is a global minimizer. Moreover, such a portfolio, $\mathbf{y^V}$, is unique. Thus, a greedy search method will lead to the global minimum, $\mathbf{y^V}$.*

An interesting by-product of the proof of Proposition 5 is the existence of $M < \infty$ such that the global minimum belongs to a bounded "box" B,

$$B = \Big\{\mathbf{y}\Big| 0 \leq y_1 \leq \ldots \leq y_n \leq M\Big\}.$$

We construct this bound M as follows. For every $k = 1, \ldots, n$, the expression

$$\mathbb{E}[(Z_k + \ldots + Z_n)^2]y_{(k)} - \mathbb{E}[Z_k + \ldots + Z_n]^2 \int_0^{y_{(k)}} \overline{F}_Q(u)du$$
$$+\mathbb{E}[Z_0(Z_k + \ldots + Z_n)](y_{(k)} + \frac{\int_{y_{(k)}}^{\infty} \overline{F}_Q(u)du}{\overline{F}_Q(y_{(k)})}) - \mathbb{E}[Z_0]\mathbb{E}[Z_k + \ldots + Z_n]\mathbb{E}[Q]$$

is strictly increasing and tends to $+\infty$ when $y_{(k)} \to +\infty$. Thus, there is $M_k > 0$ such that the expression is non-negative for $y_{(k)} > M_k$. We then define $M = \max\{M_1, \ldots, M_n\}$.

In general, the lower-level sets of variance are not convex. Figure 3 shows this observation for the following data. We consider two different options, with reservation and execution fees equal to $v_1 = 5$, $w_1 = 0$ and $v_2 = 1$, $w_2 = 6$ respectively. Customer price is $p = 10$ and spot market price follows a truncated normal distribution with mean 8 and standard deviation 1. Finally demand follows a truncated normal distribution with mean 60 and standard deviation 20.

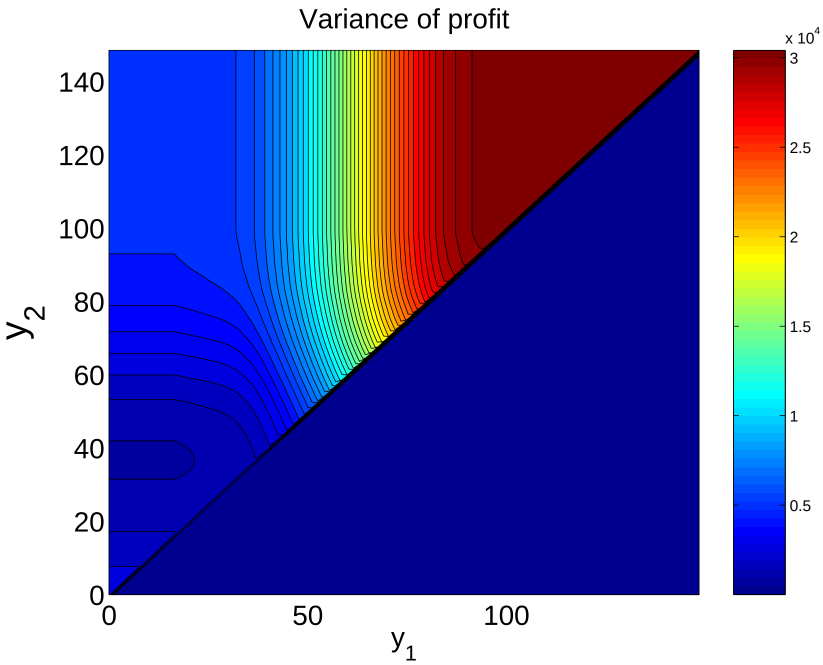

Figure 3: The lower-level sets of variance are not always convex

The non-convexity of the lower-level sets raises the following challenge. How does the variance behave when the feasible portfolio set is not F, but a smaller set? In particular, what happens when the feasible portfolio $\mathbf{y}$ is constrained within a line, i.e. there are $\mathbf{y^0}$ and $\mathbf{\Delta y}$ such that all feasible portfolios can be expressed as $\mathbf{y} = \mathbf{y^0} + t\mathbf{\Delta y}$ for some t?

This situation may arise when the manufacturer owns a portfolio $(y_1, \ldots, y_{i-1}, y_{i-1}, y_{i+1}, \ldots, y_n)$ already and is approached by a new supplier, which offers a new contract i, with parameters (w_i, v_i). Notice that at this point $x_i = 0$. Increasing x_i to $x_i + t$ implies changing $(y_i, \ldots, y_n)$ to $(y_i + t, \ldots, y_n + t)$. Of course, in this instance, we have $\Delta y_j = 1$ for $j \geq i$ and 0 otherwise.

Using Proposition 3, we have that

$$\frac{dVar\Pi^N}{dt} = \sum_{i=1}^{n} \Delta y_i \overline{F}_Q(y_i)\Phi_i(\mathbf{y})$$

and

$$\frac{d^2Var\Pi^N}{dt^2} \geq -\sum_{i=1}^{n} \Delta y_i^2 f_Q(y_i)\Phi_i(\mathbf{y}).$$

This inequality does not, in general, allow us to characterize the structure of $Var\Pi^N(t)$. However, in specific cases, we can show that this function is quasi-convex. This happens when the following two conditions are satisfied.

(1) The demand is exponentially distributed, i.e. there is $\mu > 0$ such that $\overline{F}_Q(u) = e^{-\mu u}$, and hence $f_Q(y) = \mu \overline{F}_Q(u)$.

(2) $\Delta y_i = 0, 1$ for all i, $i = 1, \dots, n$.

Indeed, observe that $\sum_{i=1}^{n} \Delta y_i \overline{F}(y_i)\Phi_i(\mathbf{y}) = 0$ implies that

$$\frac{d^2 Var\Pi^N}{dt^2} \geq -\sum_{i=1}^{n} \Delta y_i^2 f(y_i)\Phi_i(\mathbf{y}) = -\mu \sum_{i=1}^{n} \Delta y_i \overline{F}_Q(y_i)\Phi_i(\mathbf{y}) = 0.$$

Thus, the variance is a quasi-convex function of t in this case.

The situation depicted above can be easily treated using the quasi-convexity of the variance. In other words, the problem of adding a contract i to an existing portfolio is as easy as the single supplier case of Section 4.1. This is true since here $\Delta y_j = 1$ for $j \geq i$ and 0 otherwise.

4.3 The efficient frontier

We can now turn to solving the problem posed by Equation (6), finding the mean-variance trade-offs where more expectation of profit is preferred to less, and less variance is preferred to more. The efficient frontier, in terms of average profit and profit variance, will clearly be defined between the minimum variance portfolio $\mathbf{y}^{\mathbf{V}}$, defined in Corollary 1, and the maximum expectation portfolio $\mathbf{y}^{\mathbf{E}}$, defined in Proposition 2. Since expectation and variance of profit are continuous, we can find a continuum of mean-variance pairs in the efficient frontier, between these two portfolios, i.e. between $\mathbf{y}^{\mathbf{E}}$ and $\mathbf{y}^{\mathbf{V}}$. These belong in the bounded set

$$E = \Big\{\mathbf{y}\Big| 0 \leq y_1 \leq \ldots \leq y_n \leq \bar{y}\Big\},$$

where $\bar{y} = \max(M, y_n^E)$.

The next proposition summarizes the results regarding this problem.

Proposition 6 *For $\lambda \geq 0$, let the manufacturer's utility function be*

$$U(\mathbf{y}) = \mathbb{E}\Pi^N - \lambda Var\Pi^N.$$

Then, for each i, $i = 1, \dots, n$, we have that

$$\frac{dU}{dy_i} = -a_i + \overline{F}_Q(y_i)\Psi_i(\mathbf{y})$$

for some functions Ψ_i and the scalars $a_i = v_i - v_{i+1}$. Moreover, we have that

$$H = \Big(\frac{d^2 U}{dy_i dy_j}\Big)_{(i,j)} \prec \begin{pmatrix} \ddots & & 0 \\ & -f_Q(y_i)\Psi_i(\mathbf{y}) & \\ 0 & & \ddots \end{pmatrix}.$$

When $v_1 \geq \ldots \geq v_n \geq 0$, the utility upper-level sets are connected and the utility has a unique local maximum $\mathbf{y}^*$. A greedy algorithm will lead to the global maximizer of the utility function, $\mathbf{y}^*$.

This first part of the proposition is similar to Proposition 3, and the final part is derived using the same ideas of Proposition 5 and Corollary 1.

Observe that, in the proposition, we assumed that $v_1 \geq \ldots \geq v_n \geq 0$. This is not a restrictive assumption, given that $0 \leq w_1 \leq \ldots \leq w_n$. Indeed, an option that has a lower execution price should have a higher reservation price. Otherwise, we would be able to identify some option that is dominated by some other, i.e. find i, j such that $v_i \leq v_j$ and $w_i \leq w_j$. We could then eliminate option j from the pool of "acceptable" contracts, because it is too expensive.

We must point out that even though the mean-variance frontier is continuous, the corresponding efficient portfolios might not change continuously. Since the variance level sets are connected but not convex, as observed in Figure 3, we might find discontinuous jumps in the efficient portfolios.

5 Discussion

5.1 Scope of application

So far we have discussed a single period model where a manufacturer chooses a portfolio of options based on a profit mean-variance trade-off. As seen in Martínez-de-Albéniz and Simchi-Levi [9], when the manufacturer maximizes its expected profit, it is possible to extend the analysis to a multi-period case. However, as we will discuss in the following section, we have encountered important difficulties in extending the analysis for mean-variance objectives.

The single period model is nevertheless applicable in itself, in a multi-period environment. This can typically be done when the component is perishable or non-storable. In this case, since inventory is not transferred from one period to the next, one can apply this model for every period independently. For example, the model is applicable for fashion items or other products with short life-cycle, e.g., laser printers.

5.2 Multi-period models

In the financial literature, the CAPM has been extended to a multi-period setting by Merton [11] and others.

However, extending mean-variance objectives to a multi-period environment has encountered significant problems. For instance, recently, Li and Ng [8] have formulated, in a dynamic programming framework, the problem of maximizing expected return of investment under variance constraints. They provide through a clever relaxation of the problem a way to keep track of the initial variance constraint as we go by in the dynamic program. Unfortunately, the model formulation is not completely intuitive. Specifically, it assumes that once the problem is posed at the beginning of the horizon, the investor, as it moves forward in a scenario tree, keeps record of the variance constraint posed at the root of this tree. This implies that when it moves to a branch with more wealth than average, it will consciously choose less return in order to reduce the variance of returns across all branches of the tree. This point makes such formulation inappropriate for real situations.

A multi-period extension of the present newsvendor model has some additional difficulties. Indeed, when dealing with financial models, the portfolio holder can liquidate its portfolio at the beginning of every time period and reinvest the corresponding cash into a brand new portfolio. However, this is not possible in the model analyzed in this paper since we focus on components that cannot be sold back to the market since they are tailored for the buyer or due to contractual constraints.

We now present an example that illustrates the challenges with a multi-period model. Consider a two-period model with a single supplier providing a fixed commitment contract, i.e. an option with zero execution price, to be executed in the first period. The sequence of events is the following:

(i) Before period 1, the manufacturer can reserve x_1 units of capacity (to be used in period 1) at price v_1 per unit.

(ii) At the beginning of period 1, it observes the realization of demand Q_1 and spot price S_1. The manufacturer then uses the available capacity x_1 (at zero execution price) and the spot market to purchase supply, serve demand and stock I units of inventory.

(iii) At the end of period 1, I units of inventory are left.

(iv) At the beginning of period 2, Q_2 and S_2 are observed. Demand is satisfied using the on-hand inventory I and any units purchased in the spot market.

The stocking decision I is similar to the single-period mean-variance trade-off discussed in this paper. Since it is likely that $I \geq x_1 - Q_1$, because the manufacturer can raise the inventory up to the level $x_1 - Q_1$ for free, I can be dependent of Q_1, x_1 and S_1.

Thus,

$$\Pi \quad = p_1 Q_1 - v_1 x_1 - S_1 (Q_1 + I - x_1)^+ + p_2 Q_2 - S_2 (Q_2 - I)^+.$$

Once Q_1 and S_1 become known, at the beginning of the first period, the decision on I involves a trade-off between the second period's expected profit and variance. Observe that when $I < x_1 - Q_1$ an additional unit of inventory costs 0, whereas, when $I \geq x_1 - Q_1$, it costs S_1. It is clear that S_1, together with $x_1 - Q_1$, is important in determining the level I^E that maximizes expected profit, as described in Proposition 1. We also notice that the variance of the second period's profit is independent of S_1: this implies that I^V is independent of Q_1, x_1 and S_1. In general, the manufacturer's decision on I will be within $[I^E, I^V]$ (or $[I^V, I^E]$), and thus depend on the values of Q_1, x_1 and S_1.

It is now clear that when we analyze the variance of profit at the beginning of period 1, as a function of x_1, most of the complications come from the fact that the decision $I(Q_1, x_1, S_1)$ is random. This implies that the function I depends on the level of risk that the manufacturer takes as a function of Q_1, S_1. For instance, if the manufacturer is risk-averse (i.e., selects $I = I^V$) when S_1 is high and risk-neutral (i.e., selects $I = I^E(x_1 - Q_1, S_1)$) when S_1 is low, then $I(Q_1, x_1, S_1)$ will have a high variance. As a result, this variance will influence the total variance of Π.

Formally, by using the conditional variance formula, and assuming that (Q_2, S_2) are independent of (Q_1, S_1), we have

$$\begin{aligned}
Var\Pi = \quad & \mathbb{E}\left[\, Var\left(\begin{array}{c} p_1Q_1 - v_1x_1 - S_1(Q_1 + I(Q_1,x_1,S_1) - x_1)^+ \\ +p_2Q_2 - S_2(Q_2 - I(Q_1,x_1,S_1))^+ \end{array} \middle| \, S_1, Q_1 \right) \right] \\
+ \quad & Var\left[\, \mathbb{E}\left(\begin{array}{c} p_1Q_1 - v_1x_1 - S_1(Q_1 + I(Q_1,x_1,S_1) - x_1)^+ \\ +p_2Q_2 - S_2(Q_2 - I(Q_1,x_1,S_1))^+ \end{array} \middle| \, S_1, Q_1 \right) \right] \\
= \quad & \mathbb{E}\Big[\, Var\Big(\; p_2Q_2 - S_2(Q_2 - I(Q_1,x_1,S_1))^+ \; \Big| \, S_1, Q_1 \Big) \Big] \\
+ \quad & Var\left[\begin{array}{c} p_1Q_1 - v_1x_1 - S_1(Q_1 + I(Q_1,x_1,S_1) - x_1)^+ \\ -\mathbb{E}\Big(S_2(Q_2 - I(Q_1,x_1,S_1))^+ | S_1, Q_1 \Big) \end{array} \right]
\end{aligned}$$

We see that now the choice of x_1 will have an influence on the random inventory decision $I(Q_1, x_1, S_1)$ and this implies that controling the variance is difficult. In the case where $I = I^V$, and thus independent of x_1, the problem of minimizing variance in terms of $x_1 \geq 0$ is equivalent to minimizing the function

$$Var\Big[p_1Q_1 - v_1x_1 - S_1(Q_1 + I^V - x_1)^+\Big]. \tag{16}$$

In this particular case, the analysis is similar to the single period case, except that we must work with a modified demand, equal to $Q_1 + I^V$. Similarly, we can extend this example to a multiple period case, provided that at every time period we minimize the "variance-to-go", disregarding the expected profit. The manufacturer in this case finds the minimum variance portfolio for a demand $Q_t + I^V_{t+1}$.

To summarize, the insight provided here is the following. A single period model allows a manufacturer to consider mean-variance trade-offs, and the problem can be described by some interesting properties, involving a maximum-profit-expectation portfolio $\mathbf{y^E}$, and a minimum-profit-variance portfolio $\mathbf{y^V}$. Between these two portfolios, a continuum of efficient mean-variance pairs exist. However, in a multi-period setting, we face many complications. First, similarly to financial theory, it is difficult to pose an optimization problem since intermediate decisions depend on intermediate mean-variance trade-offs, thus requiring a complete specification of the preferences of the manufacturer in all the states of the world. Second, by using an inventory model that does not allow to sell back inventory to the spot market, past inventory decisions create constraints on present inventory decisions, and this in turn modifies the present mean-variance trade-off.

Hence, we see that multi-period models are significantly more difficult than single-period models. There are two important exceptions to this conclusion: when the manufacturer only cares about expected profit or when it only cares about future variance. The first case is studied in Martínez-de-Albéniz and Simchi-Levi [9]. A solution to the second case is suggested by Equation (16).

6 Conclusion

Most inventory decisions imply tremendous risks for a buyer, especially when this stock is limited to in-house production and there is no way to get rid of it after purchasing, e.g., the recent Cisco case. Thus, this type of decisions should take into account not only expected profit but also the associated risk. For this purpose, we propose, as has been done in the financial literature, to apply a mean-variance analysis to procurement contracts.

Our focus in this paper is on a single-period inventory setting where purchasing decisions create both inventory risk, i.e. created by demand uncertainty, and price risk, i.e. created by alternative spot sourcing uncertainty. The contracts used in our model are portfolios of simple option contracts, which can replicate fixed commitment contracts, quantity flexibility contracts or buy-back contracts, as shown in Martínez-de-Albéniz and Simchi-Levi [9].

We show that there is an efficient frontier bounded by the maximum expectation portfolio and the minimum variance portfolio, and provide bounds for this frontier. These two portfolios would be selected by a risk-neutral buyer and a risk-obsessed, i.e. with infinite risk aversion, buyer, respectively. We investigate structural properties of mean-variance utility objectives, which, even though they are not concave in general, can be shown to have connected upper-level sets. Such a result provides a theoretical foundation to the use of greedy algorithms to solve these trade-offs to optimality.

Finally, we need to point out an important extension of our model to situations when the buyer can sell back to the spot market any remaining inventory at the end of the season. If the market is perfect, i.e., the buying and selling spot prices are equal, then the financial model analyzed in Section 2 applies. If, on the other hand, there is a bid-ask spread in the spot market the situation is significantly more complex.

References

[1] Anvari M. 1987. "Optimality Criteria and Risk in Inventory Models: The Case of the Newsboy Problem." *J. Opl. Res. Soc.*, **38**(7), pp. 625-632.

[2] Barrett L. 2001. "Cisco's $2.5 billion mea culpa." *News.com*, May 9, 2001.

[3] Caldentey R. and M. Haugh 2003. "Optimal Control and Hedging of Operations in the Presence of Financial Markets." Working Paper.

[4] Chung K.H. 1990. "Risk in Inventory Models: The Case of the Newsboy Problem - Optimality Conditions." *J. Opl. Res. Soc.*, **41**(2), pp. 173-176.

[5] Duffie D. and H.R. Richardson 1991. "Mean-Variance Hedging in Continuous Time." *The Annals of Applied Probability*, **1**(1), pp. 1-15.

[6] Gouriéroux C., J.P. Laurent and H. Pham 1998. "Mean-Variance Hedging and Numéraire." *Mathematical Finance*, **8**(3), pp. 179-200.

[7] Lau H.S. 1980. "The Newsboy Problem under Alternative Optimization Objectives." *J. Opl. Res. Soc.*, **31**(1), pp. 525-535.

[8] Li D. and D.L. Ng 2000. "Optimal Dynamic Portfolio Selection: Multiperiod Mean-Variance Formulation." *Mathematical Finance*, **10**(3), pp. 387-406.

[9] Martínez-de-Albéniz V. and D. Simchi-Levi 2002. "A Portfolio Approach to Procurement Contracts." Submitted.

[10] Merton R.C. 1972. "An Analytical Derivation of the Efficient Portfolio Frontier." *The Journal of Financial and Quantitative Analysis*, **7**(4), pp. 1851-1872.

[11] Merton R.C. 1973. "An Intertemporal Capital Asset Pricing Model." *Econometrica*, **41**(5), pp. 867-887.

[12] Mossin J. 1966. "Equilibrium in a Capital Asset Market." *Econometrica*, **34**(4), pp. 768-783.

[13] Schweizer M. 1992. "Mean-Variance Hedging for General Claims." *The Annals of Applied Probability*, **2**(1), pp. 171-179.

[14] Sharpe W. 1964. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *The Journal of Finance*, **19**(3), pp. 425-442.

A Variance analysis

The following equations describe the gradient of the variance as a function of $y_1, \ldots, y_n$.

$$\begin{aligned}
A_i(y) &= \frac{1}{2}\mathbb{E}[Z_i^2]\mathbb{E}[\min(Q, y_i)^2] - \frac{1}{2}\mathbb{E}[Z_i]^2\mathbb{E}[\min(Q, y_i)]^2 \\
\frac{dA_i}{dy_i} &= \mathbb{E}[Z_i^2]y_i\overline{F}_Q(y_i) - \mathbb{E}[Z_i]^2\mathbb{E}[\min(Q, y_i)]\overline{F}_Q(y_i) \\
&= \overline{F}_Q(y_i)\Big(\mathbb{E}[Z_i^2]y_i - \mathbb{E}[Z_i]^2 \textstyle\int_0^{y_i} \overline{F}_Q(u)du\Big) \\
\frac{dA_i}{dy_j} &= 0 \text{ for } j \neq i
\end{aligned}$$

$$\begin{aligned}
B_i(y) &= \mathbb{E}[Z_0 Z_i]\mathbb{E}[Q \min(Q, y_i)] - \mathbb{E}[Z_0]\mathbb{E}[Z_i]\mathbb{E}[Q]\mathbb{E}[\min(Q, y_i)] \\
\frac{dB_i}{dy_i} &= \mathbb{E}[Z_0 Z_i](y_i\overline{F}_Q(y_i) + \int_{y_i}^{\infty} \overline{F}_Q(u)du) - \mathbb{E}[Z_0]\mathbb{E}[Z_i]\mathbb{E}[Q]\overline{F}_Q(y_i) \\
&= \overline{F}_Q(y_i)\Big(\mathbb{E}[Z_0 Z_i](y_i + \frac{\int_{y_i}^{\infty} \overline{F}_Q(u)du}{\overline{F}_Q(y_i)}) - \mathbb{E}[Z_0]\mathbb{E}[Z_i]\mathbb{E}[Q]\Big) \\
\frac{dB_i}{dy_j} &= 0 \text{ for } j \neq i
\end{aligned}$$

$$\begin{aligned}
C_{ij}(y) &= \mathbb{E}[Z_i Z_j]\mathbb{E}[\min(Q, y_i)\min(Q, y_j)] - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\mathbb{E}[\min(Q, y_i)]\mathbb{E}[\min(Q, y_j)] \\
\frac{dC_{ij}}{dy_i} &= \mathbb{E}[Z_i Z_j](y_i\overline{F}_Q(y_i) + \int_{y_i}^{y_j} \overline{F}_Q(u)du) - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\overline{F}_Q(y_i)\int_0^{y_j} \overline{F}_Q(u)du \\
&= \overline{F}_Q(y_i)\Big(\mathbb{E}[Z_i Z_j](y_i + \frac{\int_{y_i}^{y_j} \overline{F}_Q(u)du}{\overline{F}_Q(y_i)}) - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\int_0^{y_j} \overline{F}_Q(u)du\Big) \\
\frac{dC_{ij}}{dy_j} &= \mathbb{E}[Z_i Z_j]y_i\overline{F}_Q(y_j) - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\overline{F}_Q(y_j)\int_0^{y_i} \overline{F}_Q(u)du \\
&= \overline{F}_Q(y_j)\Big(\mathbb{E}[Z_i Z_j]y_i - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\int_0^{y_i} \overline{F}_Q(u)du\Big)
\end{aligned}$$

B Proofs

B.1 Proposition 3

Proof. We have that for all $i = 1, \ldots, n$,

$$\frac{1}{2}\frac{dVar\Pi^N}{dy_i} = \overline{F}_Q(y_i)\begin{pmatrix} \mathbb{E}[Z_i^2]y_i - \mathbb{E}[Z_i]^2\int_0^{y_i}\overline{F}_Q(u)du \\ \\ +\mathbb{E}[Z_0Z_i](y_i + \dfrac{\displaystyle\int_{y_i}^{\infty}\overline{F}_Q(u)du}{\overline{F}_Q(y_i)}) - \mathbb{E}[Z_0]\mathbb{E}[Z_i]\mathbb{E}[Q] \\ \\ +\displaystyle\sum_{j>i}\mathbb{E}[Z_iZ_j](y_i + \dfrac{\displaystyle\int_{y_i}^{y_j}\overline{F}_Q(u)du}{\overline{F}_Q(y_i)}) - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\int_0^{y_j}\overline{F}_Q(u)du \\ \\ +\displaystyle\sum_{j<i}\mathbb{E}[Z_iZ_j]y_j - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\int_0^{y_j}\overline{F}_Q(u)du \end{pmatrix}$$

Thus, we can define Φ_i where

$$\begin{aligned} 2\Phi_i(\mathbf{y}) = \quad & \mathbb{E}[Z_i^2]y_i - \mathbb{E}[Z_i]^2\int_0^{y_i}\overline{F}_Q(u)du \\ & +\mathbb{E}[Z_0Z_i](y_i + \frac{\int_{y_i}^{\infty}\overline{F}_Q(u)du}{\overline{F}_Q(y_i)}) - \mathbb{E}[Z_0]\mathbb{E}[Z_i]\mathbb{E}[Q] \\ & +\sum_{j>i}\mathbb{E}[Z_iZ_j](y_i + \frac{\int_{y_i}^{y_j}\overline{F}_Q(u)du}{\overline{F}_Q(y_i)}) - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\int_0^{y_j}\overline{F}_Q(u)du \\ & +\sum_{j<i}\mathbb{E}[Z_iZ_j]y_j - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\int_0^{y_j}\overline{F}_Q(u)du \end{aligned}$$

such that $\dfrac{dVar\Pi^N}{dy_i} = \overline{F}_Q(y_i)\Phi_i(\mathbf{y})$.

We can know compute the Hessian of the variance. For $i, j = 1, \dots, n$, we have

$$\frac{1}{2}\frac{d^2Var\Pi^N}{dy_i^2} = -\frac{1}{2}f_Q(y_i)\Phi_i(\mathbf{y}) + \overline{F}_Q(y_i)\begin{pmatrix} \mathbb{E}[Z_i^2] - \mathbb{E}[Z_i]^2\overline{F}_Q(y_i) \\ \\ +\mathbb{E}[Z_0Z_i]\dfrac{f_Q(y_i)\displaystyle\int_{y_i}^{\infty}\overline{F}_Q(u)du}{\overline{F}_Q(y_i)^2} \\ \\ +\displaystyle\sum_{j>i}\mathbb{E}[Z_iZ_j]\dfrac{f_Q(y_i)\displaystyle\int_{y_i}^{y_j}\overline{F}_Q(u)du}{\overline{F}_Q(y_i)^2} \end{pmatrix}; \qquad (17)$$

for $i < j$,

$$\begin{aligned}\frac{1}{2}\frac{d^2 Var\Pi^N}{dy_i dy_j} &= \overline{F}_Q(y_i)\left(\ \mathbb{E}[Z_iZ_j]\frac{\overline{F}_Q(y_j)}{\overline{F}_Q(y_i)} - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\overline{F}_Q(y_j)\ \right)\\ &= \overline{F}_Q(y_j)\mathbb{E}[Z_iZ_j] - \overline{F}_Q(y_i)\overline{F}_Q(y_j)\mathbb{E}[Z_i]\mathbb{E}[Z_j];\end{aligned} \tag{18}$$

and for $i > j$,

$$\begin{aligned}\frac{1}{2}\frac{d^2 Var\Pi^N}{dy_i dy_j} &= \overline{F}_Q(y_i)\left(\ \mathbb{E}[Z_iZ_j] - \mathbb{E}[Z_i]\mathbb{E}[Z_j]\overline{F}_Q(y_j)\ \right)\\ &= \overline{F}_Q(y_i)\mathbb{E}[Z_iZ_j] - \overline{F}_Q(y_i)\overline{F}_Q(y_j)\mathbb{E}[Z_i]\mathbb{E}[Z_j].\end{aligned} \tag{19}$$

We claim that the Hessian without the diagonal terms $-f_Q(y_i)\Phi_i(\mathbf{y})$ is a positive definite matrix.

Define $\alpha_{n+1} = 0$, $\alpha_i = \overline{F}_Q(y_i)$ and $\Delta_i = \alpha_i - \alpha_{i+1}$ for $i = 1, \ldots, n$. Clearly, $\Delta_1 + \ldots + \Delta_n = \alpha_1 \leq 1$, we will use this inequality later. Using Equations (17), (18) and (19), together with (using here Assumption 2) the fact that

$$\mathbb{E}[Z_0Z_i]\frac{f_Q(y_i)\displaystyle\int_{y_i}^{\infty}\overline{F}_Q(u)du}{\overline{F}_Q(y_i)^2} + \sum_{j>i}\mathbb{E}[Z_iZ_j]\frac{f_Q(y_i)\displaystyle\int_{y_i}^{y_j}\overline{F}_Q(u)du}{\overline{F}_Q(y_i)^2} > 0,$$

we obtain that

$$
\begin{aligned}
&2H - 2\begin{pmatrix} \ddots & & 0 \\ & -f_Q(y_i)\Phi_i(\mathbf{y}) & \\ 0 & & \ddots \end{pmatrix} \\
&\succ \begin{pmatrix} \ddots & & & \ddots \\ & \alpha_i\mathbb{E}[Z_i^2] - \alpha_i^2\mathbb{E}[Z_i]^2 & \cdots & \alpha_j\mathbb{E}[Z_iZ_j] - \alpha_i\alpha_j\mathbb{E}[Z_i]\mathbb{E}[Z_j] \\ & \vdots & \ddots & \vdots \\ & \alpha_j\mathbb{E}[Z_iZ_j] - \alpha_i\alpha_j\mathbb{E}[Z_i]\mathbb{E}[Z_j] & \cdots & \alpha_j\mathbb{E}[Z_j^2] - \alpha_j^2\mathbb{E}[Z_j]^2 \\ \ddots & & & \ddots \end{pmatrix} \\
&= \begin{pmatrix} \alpha_1\mathbb{E}[Z_1^2] & \cdots & \alpha_i\mathbb{E}[Z_1Z_i] & \cdots & \alpha_j\mathbb{E}[Z_1Z_j] & \cdots \\ & \ddots & & & \vdots & \\ & & \alpha_i\mathbb{E}[Z_i^2] & \cdots & \alpha_j\mathbb{E}[Z_iZ_j] & \cdots \\ & & & \ddots & \vdots & \\ \vdots & & & & \alpha_j\mathbb{E}[Z_j^2] & \\ & & & & & \ddots \end{pmatrix} - \mathbb{E}\begin{pmatrix} \alpha_1Z_1 \\ \vdots \\ \alpha_iZ_i \\ \vdots \\ \alpha_nZ_n \end{pmatrix} \mathbb{E}\begin{pmatrix} \alpha_1Z_1 \\ \vdots \\ \alpha_iZ_i \\ \vdots \\ \alpha_nZ_n \end{pmatrix}'
\end{aligned}
\tag{20}
$$

Define

$$
U_i = \begin{pmatrix} Z_1 \\ \cdots \\ Z_i \\ 0 \\ \cdots \\ 0 \end{pmatrix}.
$$

We observe that

$$
\begin{aligned}
\begin{pmatrix} \alpha_1Z_1^2 & \cdots & \alpha_iZ_1Z_i & \cdots & \alpha_jZ_1Z_j & \cdots \\ & \ddots & & & \vdots & \\ & & \alpha_iZ_i^2 & \cdots & \alpha_jZ_iZ_j & \cdots \\ \vdots & & & \ddots & \vdots & \\ & & & & \alpha_jZ_j^2 & \\ & & & & & \ddots \end{pmatrix} &= \sum_i \Delta_i \begin{pmatrix} Z_1^2 & \cdots & Z_1Z_i & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \vdots & & \vdots \\ Z_1Z_i & \cdots & Z_i^2 & 0 & \cdots & 0 \\ 0 & \cdots & 0 & 0 & & \vdots \\ \vdots & & \vdots & & \ddots & \\ 0 & \cdots & 0 & \cdots & & 0 \end{pmatrix} \\
&= \sum_i \Delta_i U_i U_i'
\end{aligned}
$$

and also

$$\begin{pmatrix} \alpha_1 Z_1 \\ \vdots \\ \alpha_i Z_i \\ \vdots \\ \alpha_n Z_n \end{pmatrix} = \sum_i \Delta_i U_i.$$

We can now express the right hand side of Equation (20) as:

$$\sum_i \Delta_i^2 \mathbb{E}\Big[U_i U_i'\Big] + \sum_i \Delta_i (1-\Delta_i) \mathbb{E}\Big[U_i U_i'\Big] - \sum_i \Delta_i^2 \mathbb{E}\Big[U_i\Big]\mathbb{E}\Big[U_i'\Big] - \sum_{i,j\ i\neq j} \Delta_i \Delta_j \mathbb{E}\Big[U_i\Big]\mathbb{E}\Big[U_j'\Big].$$

Since a variance-covariance matrix is positive semi-definite, we have that for all i,

$$\mathbb{E}\Big[U_i U_i'\Big] \succeq \mathbb{E}\Big[U_i\Big]\mathbb{E}\Big[U_i'\Big].$$

This implies that the first term minus the third term is the sum of n positive semi-definite matrices.

Also, as pointed out before,

$$1 - \Delta_i \geq \sum_{j\ j\neq i} \Delta_j.$$

Thus, the second term minus the fourth term is greater than (in the positive semi-definite ordering sense)

$$\sum_{i,j\ i<j} \Delta_i \Delta_j \Big\{ \mathbb{E}\Big[U_i U_i'\Big] + \mathbb{E}\Big[U_j U_j'\Big] - \mathbb{E}\Big[U_i\Big]\mathbb{E}\Big[U_j'\Big] - \mathbb{E}\Big[U_j\Big]\mathbb{E}\Big[U_i'\Big] \Big\}. \tag{21}$$

Observe that for all $i < j$, we have

$$\begin{aligned} & \mathbb{E}\Big[U_i U_i'\Big] + \mathbb{E}\Big[U_j U_j'\Big] - \mathbb{E}\Big[U_i\Big]\mathbb{E}\Big[U_j'\Big] - \mathbb{E}\Big[U_j\Big]\mathbb{E}\Big[U_i'\Big] \\ \succeq\ & \mathbb{E}\Big[U_i\Big]\mathbb{E}\Big[U_i'\Big] + \mathbb{E}\Big[U_j\Big]\mathbb{E}\Big[U_j'\Big] - \mathbb{E}\Big[U_i\Big]\mathbb{E}\Big[U_j'\Big] - \mathbb{E}\Big[U_j\Big]\mathbb{E}\Big[U_i'\Big] \\ =\ & \mathbb{E}\Big[U_i - U_j\Big]\mathbb{E}\Big[U_i' - U_j'\Big] \\ \succeq\ & 0. \end{aligned}$$

In other words, all the terms in the sum in Equation (21) are positive semi-definite, and thus the following matrix is positive definite,

$$H \succ \begin{pmatrix} \ddots & & 0 \\ & -f_Q(y_i)\Phi_i(\mathbf{y}) & \\ 0 & & \ddots \end{pmatrix}.$$

■

B.2 Proposition 4

Proof. We use the results from Proposition 3. Without loss of generality, we can assume that $I = \{1, \ldots, n\}$, since when I is smaller, we can conduct the same calculations of Proposition 3 with the subset I of contracts.

We assume that $\mathbf{y}$ is a critical point of the variance, i.e. $\dfrac{dVar\Pi^N}{y_i}(\mathbf{y}) = 0$ for $i = 1, \ldots, n$. Proposition 3 implies that $\Phi_i(\mathbf{y}) = 0$ for all i, and thus the Hessian is definite positive. Since the variance is twice continuously differentiable, the critical point is a strict minimum. ∎

B.3 Proposition 5

Proof. Before starting this proof, define the norm $\|\cdot\|_\infty$ in $\mathbb{R}^n_+$, i.e.

$$\|\mathbf{y}\|_\infty = \max_i |y_i|.$$

Consider for every $k = 1, \ldots, n$ the portfolio where $y_{k-1} < y_k = y_{k+1} = \ldots = y_n = y_{(k)}$. In other words, the $n-k$ highest inequality constraints of the feasible set are tight. We claim that there exists an M_k such that, for $y_{(k)} > M_k$, the variance is non-decreasing as we increase $y_{(k)}$. To see this, we look at the derivative of the variance as a function of $y_{(k)}$. Since $Covar[Z_i, Z_j] \geq 0$ for all pairs (i,j) and $y_j \geq \int_0^{y_j} \overline{F}_Q(u)du$, we have

$$\begin{aligned}
\frac{dVar\Pi^N}{dy_{(k)}} &= \overline{F}_Q(y_{(k)}) \left(\begin{array}{l}
\mathbb{E}[(Z_k + \ldots + Z_n)^2]y_{(k)} - \mathbb{E}[Z_k + \ldots + Z_n]^2 \int_0^{y_{(k)}} \overline{F}_Q(u)du \\[2ex]
+\mathbb{E}[Z_0(Z_k + \ldots + Z_n)](y_{(k)} + \dfrac{\int_{y_{(k)}}^\infty \overline{F}_Q(u)du}{\overline{F}_Q(y_{(k)})}) - \mathbb{E}[Z_0]\mathbb{E}[Z_k + \ldots + Z_n]\mathbb{E}[Q] \\[2ex]
+ \displaystyle\sum_{j \leq k-1} \Big\{ \mathbb{E}[(Z_k + \ldots + Z_n)Z_j]y_j - \mathbb{E}[Z_k + \ldots + Z_n]\mathbb{E}[Z_j] \int_0^{y_j} \overline{F}_Q(u)du \Big\}
\end{array} \right) \\
&\geq \overline{F}_Q(y_{(k)}) \left(\begin{array}{l}
\mathbb{E}[(Z_k + \ldots + Z_n)^2]y_{(k)} - \mathbb{E}[Z_k + \ldots + Z_n]^2 \int_0^{y_{(k)}} \overline{F}_Q(u)du \\[2ex]
+\mathbb{E}[Z_0(Z_k + \ldots + Z_n)](y_{(k)} + \dfrac{\int_{y_{(k)}}^\infty \overline{F}_Q(u)du}{\overline{F}_Q(y_{(k)})}) - \mathbb{E}[Z_0]\mathbb{E}[Z_k + \ldots + Z_n]\mathbb{E}[Q]
\end{array} \right)
\end{aligned}$$

The expression

$$\mathbb{E}[(Z_k + \ldots + Z_n)^2]y_{(k)} - \mathbb{E}[Z_k + \ldots + Z_n]^2 \int_0^{y_{(k)}} \overline{F}_Q(u)du$$
$$+\mathbb{E}[Z_0(Z_k + \ldots + Z_n)](y_{(k)} + \frac{\int_{y_{(k)}}^{\infty} \overline{F}_Q(u)du}{\overline{F}_Q(y_{(k)})}) - \mathbb{E}[Z_0]\mathbb{E}[Z_k + \ldots + Z_n]\mathbb{E}[Q]$$

is strictly increasing and tends to $+\infty$ when $y_{(k)} \to +\infty$. Thus, there is $M_k > 0$ such that $\frac{dVar\Pi^N}{dy_{(k)}} \geq 0$ for $y_{(k)} > M_k$.

By writing $M = \max\{M_1, \ldots, M_n, \mathbf{y^0}_n, \mathbf{y^1}_n\}$, define the compact set

$$B = \left\{\mathbf{y}\middle| 0 \leq y_1 \leq \ldots \leq y_n \leq M\right\}.$$

Observe that the set is compact because it is closed and bounded in a finite-dimensional space.

By construction, for every point $\mathbf{y} \in F \backslash B$, we can find a point $\mathbf{y'} \in B$ with $Var\Pi^N(\mathbf{y}) \geq Var\Pi^N(\mathbf{y'})$. This can be done by defining $\mathbf{y'}_i = M$ for any i such that $\mathbf{y}_i > M$.

The function $Var\Pi^N(\cdot)$ is twice differentiable in the bounded set B, and this implies that the sets

$$L_c = \left\{\mathbf{y} \in B \middle| Var\Pi^N(\mathbf{y}) \leq c\right\}$$

are compact sets, since the variance is a continuous function. By increasing c, the level sets L_c increase in the inclusion sense.

By contradiction, assume that the level sets L_c of the variance are not connected, for some $c = a \geq 0$, i.e. L_a has at least two connected components. Let C_0 and C_1 be two unconnected components of L_a, and take $\mathbf{y^0} \in C_0$ and $\mathbf{y^1} \in C_1$. The variance on B is a continuous function on a closed finite-dimensional set, and hence attains a maximum over B. This implies that for some c, big enough, $L_c = B$, which connects $\mathbf{y^0}$ and $\mathbf{y^1}$.

We can now define b the largest c such that $\mathbf{y^0}$ and $\mathbf{y^1}$ cannot be connected through a continuous path in L_c, i.e., they belong to two different connected parts of L_c. In this definition, b is a supremum. We can define for all $n \geq 0$, two sequences in B, p_n and q_n that solve

$$\min \|p - q\|_\infty$$
$$\text{subject to } \begin{cases} p \text{ and } \mathbf{y^0} \text{ are connected in } L_{b-1/n} \\ q \text{ and } \mathbf{y^1} \text{ are connected in } L_{b-1/n} \end{cases}$$

This is possible because $L_{b-1/n}$ is a closed set. These sequences belong to the compact set L_b, thus, by the Bolzano-Weierstrass theorem, they have adherence points. We can then extract a subsequence of (p_n, q_n) such that both p_n and q_n converge to a point $\bar{\mathbf{y}}$ that belongs to L_b, but not to any $L_{b-1/n}$. This implies that $Var(\bar{\mathbf{y}}) = b$.

Define I as the set of indexes such that $\bar{y}_i > \bar{y}_{i-1}$ and A_I as in Equation (15). We claim that $\bar{\mathbf{y}}$ is a I-unconstrained critical point. If this was not the case, then, since the variance is twice differentiable, we could redirect the path that connects $\mathbf{y^0}$ to $\mathbf{y^1}$ in L_b avoiding the variance level b. That is, for some n, find a path that connects $\mathbf{y^0}$ to $\mathbf{y^1}$ in $L_{b-1/n}$. Such a path exists when the first derivative of the variance (as a function of the variables z_j used in Definition 1) at that point is non-zero. We can thus apply Proposition 4 to establish that $\bar{\mathbf{y}}$ is a local minimum when the feasible set is A_I.

In addition, for $i \notin I$, it must be that $\dfrac{dVar\Pi^N}{dy_i} \geq 0$, since otherwise we could, again, redirect the path that connects $\mathbf{y^0}$ to $\mathbf{y^1}$ in L_b avoiding the variance level b. This implies that $\bar{\mathbf{y}}$ is a local minimum for the variance.

However, the sequence p_n tends to $\bar{\mathbf{y}}$ with a variance always strictly smaller that b. This is a contradiction, and thus the level sets L_c must be connected. ■

B.4 Proposition 6

Proof. We have that for all $i = 1, \ldots, n$, from Equation (10)

$$\frac{d\mathbb{E}\Pi^N}{dy_i} = -a_i + \overline{F}_Q(y_i)\mathbb{E}[Z_i],$$

where $a_i = v_i - v_{i+1}$.

Using Proposition 3, we also know that

$$\frac{dVar\Pi^N}{dy_i} = \overline{F}_Q(y_i)\Phi_i(\mathbf{y}).$$

Thus, by writing $\Psi_i(\mathbf{y}) = \mathbb{E}[Z_i] - \lambda\Phi_i(\mathbf{y})$, we have that

$$\frac{dU}{dy_i} = -a_i + \overline{F}_Q(y_i)\Psi_i(\mathbf{y}).$$

Moreover, using again Proposition 3, we know that

$$\left(\frac{d^2U}{dy_i dy_j}\right)_{(i,j)} \prec \begin{pmatrix} \ddots & & 0 \\ & -f_Q(y_i)\Psi_i(\mathbf{y}) & \\ 0 & & \ddots \end{pmatrix}.$$

In addition, when for all i, $i = 1, \ldots, n$, $a_i \geq 0$, we can use the arguments of Proposition 4 to show that, for any set I, whenever some portfolio $\mathbf{y}$ is an I-unconstrained critical point of the utility, it is a strict local maximum of the utility over A_I. Intuitively, for $I = \{1, \ldots, n\}$, $\dfrac{dU}{dy_i} = 0$ implies $\Psi_i(\mathbf{y}) \geq 0$, and hence the Hessian is negative definite.

The proof is completed using the same arguments as that of Proposition 5, and Corollary 1. ■