Integrated Production Planning and Risk Hedging

Suggested Citation: Liao Wang and David D. Yao (2017), "Integrated Production Planning and Risk Hedging", : Vol. issue, No. XX, pp 1–12. DOI: XXX.

Liao Wang
lw2489@columbia.edu

David D. Yao
yao@columbia.edu

Contents

Integrated Production Planning and Risk Hedging

Liao Wang[1] and David D. Yao[2]

[1,2] *Department of Industrial Engineering and Operations Research, Columbia University; lw2489@columbia.edu, yao@columbia.edu.*

ABSTRACT

We study production planning integrated with risk hedging. In addition to using a one-time production quantity decision, made at the beginning of a planning horizon, as a way to manage demand uncertainty, we illustrate how to construct and execute a hedging strategy throughout the horizon, as a better and more effective approach to mitigating the risks involved. Furthermore, whereas traditional production planning models focus on the expected net-profit as an objective function, we study two *risk measures*, variance and shortfall. In both cases, we characterize the efficient frontier, and demonstrate the improved risk-return profile over a production-only decision.

1 Motivation

Production planning is essentially a *capacity* decision: how much to produce, so as to meet certain revenue/profit target, based on demand forecast and taking into account its uncertain nature. Once this decision is made, many other actions (and decisions) will follow — ordering raw materials and/or sub-assemblies, organizing the work force, setting up production lines, as well as carrying out the physical process that

Liao Wang and David D. Yao (2017), “Integrated Production Planning and Risk Hedging”, : Vol. issue, No. XX, pp 1–12. DOI: XXX.

produces the goods. The time between making the production quantity decision and the time when all units being produced is usually called the (production) "lead time," which we will also refer to as the "planning horizon," denoted $[0, T]$ below.

Here's the rub. The actual demand, when it materializes at the end of the planning horizon, can differ substantially from the forecast at the beginning of the horizon when the production quantity decision was first made and then implemented, resulting in serious consequences due to either over-production (units that cannot be sold will incur a loss) or under-production (unmet demand means reduced profit).

The classical approach to managing this risk is the newsvendor (NV) model. It tries to find a production quantity Q so as to maximize a payoff function (in expectation), which takes the form of the net profit from sales (the smaller of Q and the realized demand) minus the net cost (cost minus salvage value). The optimal production quantity Q^* is widely known to be determined by the so-called "critical ratio," $p/(p+c)$ (with p being per unit profit and c per unit cost), applied to the inverse distribution function of the demand. The demand distribution is, of course, determined by forecast and typically takes the form of a mean (or rate) plus a Gaussian noise, meaning a normal distribution with mean and variance say, $\tilde{\mu}T$ and $\tilde{\sigma}^2T$, with $\tilde{\mu}$ and $\tilde{\sigma}$ being two parameters (determined by forecast) representing the demand rate and volatility.

In Wang and Yao (2017a,b), we have developed a new demand model that makes the demand rate a function of a financial asset, denoted $\tilde{\mu}(X_t)$, where X_t is the asset price at $t \in [0, T]$. This is motivated by the following business cases:

- A firm that manufactures certain equipment for planting or harvesting corn, a tradable commodity, experienced volatile demand for its product as the corn price fluctuates on the futures market. (*Wall Street Journal*, May 14, 2014, "Deere Needs to Wait to Harvest Its Bounty.")

- Wal-Mart experienced increased demand during the last financial crisis as consumers sought lower-priced goods and its smaller-sized competitors went out of business. (*Wall Street Journal*, Nov 14, 2008, "Wal-Mart Flourishes as Economy Turns Sour.")

- The U.S. automobile industry sharply increased forecast and production when the last recession ended. (*Wall Street Journal*, Jan 14, 2014, "Auto Makers Dare to Boost Capacity: North American Factories Will Build One Million More Cars a Year.")

In the first case ("Deere"), X_t is, naturally, the price of corn futures; in the other two cases, X_t can be taken as the price of a broad market index (e.g., S&P500) serving as a proxy for the general economy.

This new model not only captures what really drives the demand for a product, more importantly, it opens a new dimension in production planning. In addition to using the production quantity decision as a way to manage demand uncertainty, we can also do *hedging* — by taking a position on the underlying financial asset and adjust it from time to time throughout the planning horizon — to mitigate the risk in demand uncertainty. Specifically, in addition to the quantity decision Q, which is made at $t = 0$, we can now pursue a risk-hedging strategy, denoted $\vartheta := \{\theta_t, t \in [0, T]\}$, with θ_t being a position (number of shares long or short) on the underlying asset at time t. More details about this will be spelled out in §2 below.

Once risk is integrated into production planning, it only makes sense to consider risk measures as objective functions, instead of the mean payoff function as in the NV model. In this regard, we consider two risk objectives in Wang and Yao (2017a,b). The first one is to minimize the *variance* of the total payoff (from both production and hedging), subject to its mean achieving a pre-specified target. The second one is to minimize the *shortfall*, the gap between the total payoff and a pre-specified target. The two models will be overviewed in the rest of this article. Specifically, we present the problem formulations in §2, with solutions detailed in §3, and results and insights highlighted in §4.

2 Problem Formulation

The Demand Model

Let X_t denote the price at time t of a tradable asset (including a broad market index such as S&P 500), and assume it follows a geometric

Brownian motion:

$$dX_t = X_t(\mu dt + \sigma dB_t), \tag{1}$$

where μ and σ are positive constants, and B_t is the standard Brownian motion (BM). Let D_t denote the (cumulative) demand up to t, and assume the following dynamics:

$$dD_t = \tilde{\mu}(X_t)dt + \tilde{\sigma}d\tilde{B}_t, \tag{2}$$

where $\tilde{\sigma}$ is a positive constant; $\tilde{\mu}(x) \geq 0$ is a non-negative function, and $\tilde{B}_t$ is another BM, independent of B_t. Below, we will use $\mathcal{F}_t$ to denote the filtration associated with both BM's, B_t and $\tilde{B}_t$; and when we consider partial information — information on the asset price but not demand forecast, we will use $\mathcal{F}_t^X$ to denote the filtration associated with B_t alone.

Note although B_t and $\tilde{B}_t$ are assumed independent, the demand D_t above does depend on B_t, and does so in a very strong, first-order manner via the rate function $\tilde{\mu}(x)$. (That is, much stronger than statistical dependence via, for instance, adding a dB_t term in (2).) Also note that here D_t is the *forecast* demand; the actual demand will only realize at $t = T$ and beyond, not before. Thus, that D_t need not be increasing in t (due to the $d\tilde{B}_t$ term) is not a handicap, as forecast typically will be adjusted up and down over time.

To implement this new demand model, the additional work appears quite minimal. All one needs is to add a machine learning scheme to the traditional demand forecast routine. Specifically, first machine-learn the functional form of the rate function $\tilde{\mu}(\cdot)$, and then use forecast to fine-tune certain parameters involved in the function. To illustrate, we take monthly sales data from John Deere over the period 2011-15, along with the daily price data for corn futures over the same time window. Applying relevant machine learning techniques, we conclude that a linear function is a good fit for $\tilde{\mu}(\cdot)$. The same analysis is also applied to General Motor's monthly sales data versus S&P500. Refer to Figure 1 for both cases. A usual forecast routine can then be called to fine-tune the parameters involved in the linear function, its intercept and slope.

Figure 1: Monthly sales versus tradable asset prices.

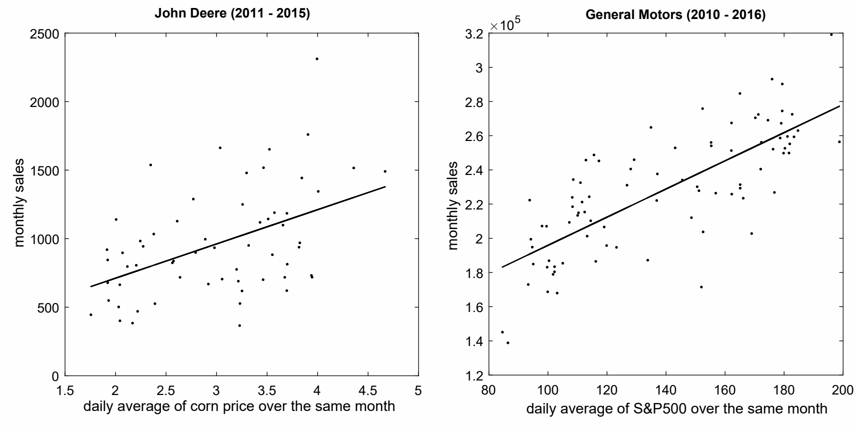

Mean-Variance Hedging

There are two components in the payoff function — the terminal wealth at $t = T$, respectively, from production and from hedging:

$$H_T(Q) := pQ - (p + c)(Q - D_T^+)^+, \quad \chi_T(\vartheta) := \int_0^T \theta_t dX_t$$

where Q denotes the production quantity; p is the unit profit (selling price minus cost), and c the unit net cost (cost minus any salvage value); θ_t is the position (number of shares) taken on the asset with per share price X_t. The decisions variables are Q and $\vartheta := \{\theta_t, t \in [0, T]\}$. Note, D_T^+ enforces the non-negativity of the realized demand at T.

For a mean-variance formulation, the problem we want to solve is:

$$B(m) := \inf_{Q \geq 0, \vartheta} \{\mathsf{Var}[H_T(Q) + \chi_T(\vartheta)] \,|\, \mathsf{E}[H_T(Q) + \chi_T(\vartheta)] = m\}, \quad (3)$$

i.e., we want to minimize the variance of the total payoff, subject to its mean meeting a given target level m. This formulation essentially follows the same spirit as Markowitz mean-variance portfolio optimization model (Markowitz (1987)). Another related study that optimizes a mean-variance utility objective is Caldentey and Haugh (2006).

Shortfall Hedging

Here we are given a budget (or credit limit) $C > 0$, to carry out the hedging strategy; and we want to minimize the *shortfall* from a given wealth level m (which can represent, for example, an earnings target):

$$\inf_{Q\geq 0,\vartheta} \mathsf{E}\Big\{\big[m - H_T(Q) - \chi_T(\vartheta)\big]^+\Big\} \tag{4}$$
$$\text{s.t.} \quad \chi_t := \int_0^t \theta_s dX_s \geq -C, \quad \theta_t \in \mathcal{F}_t^X, \qquad t \in [0,T].$$

Note, by restricting the hedging strategy to $\{\mathcal{F}_t^X\}$, the filtration generated by $\{X_t\}$ (or, equivalently, by $\{B_t\}$) alone, we are pursuing a formulation with "partial information". This is often a good representation of reality: in practice the hedging/trading decision is often a real-time decision taking input from the financial market, as embodied in the filtration $\{\mathcal{F}_t^X\}$. It would be unrealistic to assume that one could simultaneously also keep track of demand projection, which typically involves piecing together disperse information garnered from polling the sales force, and is hence updated much less frequently, at much longer time scales. Thus, in the research literature partial information is an important model in its own right; see, e.g., Caldentey and Haugh (2006).

3 Approaches and Solutions

Mean-Variance Hedging

To solve the problem in (3), we consider its *conjugate dual*:

$$A(\lambda) := \inf_{\vartheta} \mathsf{E}\Big\{\Big[\lambda - H_T(Q) - \chi_T(\vartheta)\Big]^2\Big\}, \tag{5}$$

where λ is a given parameter, in parallel to m; and the two are related via (11) below. Let's first motivate this conjugate duality. Write $Y := H_T + \chi_T$, and note

$$\begin{aligned} B(m) &= \mathsf{Var}(Y) = \mathsf{Var}(Y-\lambda) = \mathsf{E}[(Y-\lambda)^2] - [\mathsf{E}(Y-\lambda)]^2 \\ &= \mathsf{E}[(\lambda - Y)^2] - (m-\lambda)^2 = A(\lambda) - (m-\lambda)^2. \end{aligned} \tag{6}$$

Thus, finding the optimal hedging strategy ϑ via solving the $B(m)$ problem in (3) is the same as via solving the $A(\lambda)$ problem in (5), since m and λ and given parameters. Furthermore, the relation in (6) above leads to (see Proposition 6.6.5 in Pham (2009)):

$$A(\lambda) = \min_m \left[B(m) + (m-\lambda)^2\right], \quad B(m) = \max_\lambda \left[A(\lambda) - (m-\lambda)^2\right]. \quad (7)$$

Note, the problem in (5) is known as a mean-square error (MSE) problem, from which we can derive the the optimal hedging strategy (given Q) using a numeraire-based approach (Gourieroux *et al.* (1998)), and express it quite explicitly as follows:

$$\theta_t^* = -\xi_t(Q) + \frac{\eta}{\sigma X_t}[\lambda - V_t(Q) - \chi_t^*], \quad (8)$$

where $\eta := \mu/\sigma$, and $V_t(Q) := \mathsf{E}^M[H_T(Q)|\mathcal{F}_t]$. Here E^M denotes expectation with respect to the risk-neutral measure P^M (as opposed to the original, physical measure P). Note that V_t is a P^M-martingale, and admits the following representation:

$$V_t(Q) = V_0(Q) + \int_0^t \xi_s(Q)dX_s + \int_0^t \delta_s(Q)d\tilde{B}_s, \quad (9)$$

and the two processes ξ_t and δ_t can be derived explicitly via Itô's formula.

Having solved the $A(\lambda)$ problem, we can solve $B(m)$ (which we write as $B(m, Q)$ below to emphasize that Q is given) from the second equality in (7), and express the solution as follows:

$$B(m, Q) = \frac{[m - V_0(Q)]^2}{e^{\eta^2 T} - 1} + \int_0^T e^{-\eta^2(T-t)}\mathsf{E}[\delta_t^2(Q)]dt. \quad (10)$$

Finally, the λ and m relation:

$$\frac{\lambda - V_0}{m - V_0} = \frac{e^{\eta^2 t}}{e^{\eta^2 t} - 1}. \quad (11)$$

Shortfall Hedging

The hedging problem, *given* Q, is solved through a convex duality approach (refer to, e.g., Pham (2009)). Applying Jensen's inequality, along with conditional expectation on X_T and $A_T := \int_0^T \tilde{\mu}(X_t)dt$, we first turn the real-time hedging problem into a *static* optimization problem:

$$\min_{V_T} \mathsf{E}\big[(m - H_T - V_T)^+\big] \quad \text{s.t.} \quad V_T \geq -C, \quad \mathsf{E}^M(V_T) \leq 0\,. \qquad (12)$$

Note the objective function in the above optimization problem can be shown to be a lower bound of the original objective in (4). On the other hand, the constraint $\mathsf{E}^M(V_T) \leq 0$ was not present in the original problem. It follows from χ_t being a P^M-supermartingale; and it serves the purpose of closing the duality gap: the problem in (12) can be shown to be equivalent to the original problem in (4).

The dual (lower-bound) problem is solved by a standard Lagrangian multiplier approach, and the solution is:

$$V_T^* = (p + c)(Q - \hat{D}_T^+)^+ + (m - pQ + C)\mathbf{1}\{\lambda^* Z_T \leq 1\} - C, \qquad (13)$$

where λ^* is the positive Lagrangian multiplier, $Z_T := \frac{d\mathsf{P}^M}{d\mathsf{P}} = e^{\eta B_T - \eta^2 T/2}$ is the Radon-Nikodym derivative, and $\hat{D}_T := A_T + \tilde{\sigma}\sqrt{T}\Phi^{-1}(\lambda^* Z_T)$ is the "proxy" for D_T (as the latter is not accessible due to partial information) with Φ^{-1} being the inverse distribution function of the standard normal random variable. Once $V_T^*(= \chi_T^*)$ is derived, the optimal hedging strategy θ_t^* follows from Itô's Lemma along with martingale representation theorem.

Optimal Production Quantity

Under the optimal hedging strategy corresponding to V_T^* in (13), the minimized shortfall (with Q given) is

$$s(m, Q) = (p + c)\mathsf{E}\big[(Q \wedge \hat{D}_T^+ - D_T^+)^+\big] + (m - pQ + C)\mathsf{P}(\lambda^* Z_T \geq 1). (14)$$

It can be shown (with considerable effort, as both λ^* and $\hat{D}_T$ also depend on Q) that the above is a convex function of Q. Thus, finding the optimal solution $Q^*(m)$, jointly with the optimal hedging, is a readily solved convex minimization problem. To make it even better, a universal upper bound (i.e., independent of m) on the optimal Q can be explicitly identified, and this further facilitates the (line) search for

$Q^*(m)$. Furthermore, it can be shown that $s(m, Q^*(m))$ is increasing in m, hence constitutes an efficient frontier — setting a higher target will lead to a higher shortfall.

Similarly, the optimal production quantity in the mean-variance model can also be obtained from a line search on the $B(m, Q)$ expression in (10), which, however, need not be convex in Q (its first term is, but not the second one). Also, $B(m, Q^*(m))$ is increasing in m, forming an efficient frontier — a higher mean return corresponds to a higher variance.

4 Results and Insights

As a visualization of the efficient frontiers mentioned at the end of the last section, refer to Figure 2, which illustrates the frontiers for the two hedging models, corresponding to an instance of the rate function $\tilde{\mu}(\cdot)$ in the demand model. The curve labeled "NV" illustrates the frontier of the production-only newsvendor solution; not surprisingly, it yields the highest risk (standard deviation or shortfall) for any given mean/target. The curve labeled "NV+" uses the NV production quantity along with a hedging strategy that is optimized given the NV quantity; whereas the "optimal" curve jointly optimizes both production and hedging decisions. Both curves lie below "NV", with "optimal" outperforming "NV+". For both models, the contribution of the hedging strategy is quite substantial, in particular at the higher end of the mean/target, which is where the risk increases more steeply. (Note all the frontier curves are increasing and *convex*.)

Next, we take a closer look at the hedging strategies. The optimal mean-variance hedging in (8), θ_t^*, maintains two positions in the asset (X_t) at any time t:

- a position to "cancel" out the ξ_t in V_t;
- a position equal to the gap between the current wealth $(V_t + \chi_t^*)$ and the target (λ), weighted by the ratio η/σ, so as to "catch up."

For the shortfall hedging, from (13), we have the following intuitive interpretations:

Figure 2: Efficient Frontiers.

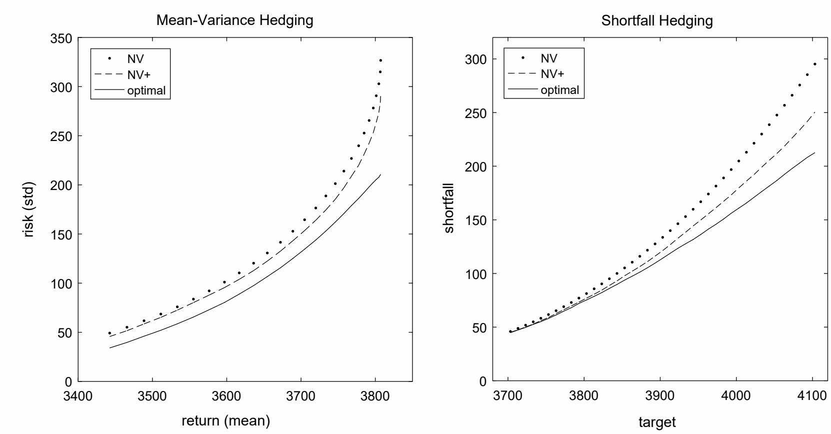

- $m - H_T(Q) = (p+c)(Q - D_T^+)^+ + (m - pQ)$ is the remaining gap (from the target) after the payoff from production.

- The "put option," $(p+c)(Q - \hat{D}_T^+)^+$, tries to close the first part of the gap, $(p+c)(Q - D_T^+)^+$, but needs to use $\hat{D}_T$ as a surrogate for D_T due to partial information.

- The "digital option," $(m - pQ + C)\mathbf{1}\{\lambda^* Z_T \le 1\}$, aims to close the other part of the gap (after subtracting C).

In summary, the hedging strategy in the mean-variance model is to dynamically maintain a portfolio of two positions on the underlying financial asset. The first position cancels out the tradable component of the projected wealth from production, i.e., revenue from supplying demand; thus, its function is one of pure risk mitigation. The second position uses trading gains to catch up with the target mean return. In contrast, the shortfall hedging strategy takes the form of two options, a digital option and a put option, and the underlying for both is the "surrogate demand" $\hat{D}_T$, necessitated by the partial information on the real demand. Since the risk measure is shortfall, both options are designed to contribute more terminal wealth, so as to supplement production and help to close the gap from the target. In other words, there is no "cancelling" component in the hedging strategy; and this is only natural as the shortfall measure does not penalize any upside risk.

5 Further Research

Many real-world applications will call for an extension of the models highlighted above to allow multiple products, each with a demand that depends on multiple financial assets. To motivate, consider several examples widely reported in the business news. During the (quite recent) period when oil price was plunging, many car buyers switched out of smaller models into SUV's and other gas guzzlers. Yet, due to relevant standards and regulations (such as Corporate Average Fuel Economy), car producers had already increased the production of more fuel-efficient models only to see them suffering from reduced demand. (*Wall Street Journal*, November 19, 2014, "Ford Presses Ahead With Developing Fuel-Efficient Vehicles."; *Wall Street Journal*, January 13, 2015, "Clash Looms Over Fuel Economy Standard.") Furthermore, demand for cars can also depend on multiple financial assets. For instance, in addition to its dependence on fuel price, the demand can also be a function of the general economy. (*Wall Street Journal*, January 14, 2014, "Automakers Dare to Boost Capacity: North American Factories Will Build One Million More Cars a Year.")

Indeed, the above business cases point to the fact that customer demands on a firm's various product lines are not just correlated in a statistical manner, they often exhibit certain *functional* relationship: a demand surge on one product typically leads to decreases in demands on other products. And this functional inter-dependence may also originate from these demands depending on overlapping sets of financial assets (such as certain commodities *and* the general economy). Thus, it will be useful to have a model for multiple products with the demand rate functions depending on multiple asset prices.

Another extension is to consider asset price models that are different or more general than the geometric BM model highlighted above, for instance, models with a mean-reverting feature (such as the Ornstein-Uhlenbeck process) so as to accommodate broader asset classes.

Acknowledgment This research is supported in part by NSF grant CMMI-1462495.

References

CALDENTEY, R. AND HAUGH, M., Optimal Control and Hedging of Operations in the Presence of Financial Markets. *Math. Opns. Res.* **31** (2006), 285-304.

GOURIEROUX, C., LAURENT, J. AND PHAM, H., Mean-variance hedging and numeraire. *Math. Financ.* **8** (1998), 179-200.

MARKOWITZ, H.M., *Mean-Variance Analysis in Portfolio Choice and Capital Markets* (1987) Basil Blackwell, New York.

PHAM, H., *Continuous-time Stochastic Control and Optimization with Financial Applications* (2009), Springer-Verlag, Berlin/Heidelberg.

WANG, L. AND YAO, D.D., Production with Risk Hedging: Optimal Policy and Efficient Frontier. *Opns. Res.* (2017a), forthcoming.

WANG, L. AND YAO, D.D., Production Planning with Shortfall Hedging under Partial Information and Budget Constraint. *Preprint* (2017b).