Exploiting Market Fluctuations and Price Volatility Through Feedback Control

Christopher V. Rao*
Chemical and Biomolecular Engineering
University of Illinois at Urbana-Champaign
Urbana, IL 61801

Abstract

In this work, we explore how feedback control can be used to make chemical producers responsive to market forces through dynamic operating policies. Using a toy model of a marginal chemical producer, we examine two different control strategies for dealing with stochastic fluctuations in operating margins. These results provide a basis for exploring more complex control problems that include the effects of market forces.

Keywords

Dynamic Programming, Process Control, Economics, Capacity Utilization

Introduction

Many products of the chemical and petroleum industries are fungible goods. They are highly liquid and in many cases actively traded on centralized exchanges. Examples include not only petroleum distillates but also common plastic such as polyethylene and polypropylene. While this liquidity reduces market friction, it also subjects producers to rapid changes and potential shocks in prices. Even in cases where the products themselves are not explicit commodities, the feedstocks are – petroleum and natural gas being prime examples. While financial instruments such as future contracts and swap can be used to hedge against these potentially disruptive market forces, they may not always exist or their premiums are too high. Additionally, these instruments prevent producers from directly exploiting prevailing market conditions, be they good or bad.

In this work, we explored how feedback control can be used to make a chemical producer responsive to market forces through the use of dynamic operating policies. To explore this problem, we developed a toy model of a marginal chemical producer with variable capacity operating in a dynamic market environment. The problem is admittedly simple and borrows heavily from classic work on inventory control, production scheduling, and portfolio optimization (e.g. [3]). However, our immediate goal was not to solve a realistic problem but rather to explore the general nature of these policies under very simplifying assumptions. In particular, we sought to explore how these policies differ from the ones traditionally used in chemical process control as a first step towards addressing more complex problems.

*To whom all correspondence should be addressed

Problem Statement

Consider the idealized scenario where there are M identical plants, each yielding the same operating profit or loss x_t (e.g. crack spreads) during the time interval $[t, t+\Delta t)$. Our problem is to determine at the current time t the number of plants that should be active or idle at the future time $t + \Delta t$ when $x_{t+\Delta t}$ is an unknown quantity. In other words, if each active plant will yield a profit or loss of x_t during the current interval $[t, t+\Delta t)$, then how many plants should be active at time $t + \Delta t$?

We further assume that there are costs associated with activating idle plants and idling active ones. Specifically, we assign a cost c_A associated with the action of activating an idle plant at time $t+\Delta t$ and a cost c_D with the action of idling an active plant $t + \Delta t$. For simplicity, we take these costs to be additive. In addition, we assume that there is a cost c_I associated with maintaining a plant in the idle state over the interval $[t, t + \Delta t)$. Again, we take this cost to be additive.

This decision is clearly affected by our ability to accurately forecast future operating margins. If these future values are known with certainty, then the problem is straightforward to solve. However, if these values are unknown, then we can at best only base our decision on expected future profits or losses. This in turn requires some measure of our uncertainty. In particular, we need to assign probabilities to our forecasts. To make the problem tractable, we assume that the price dynamics are Markov and that the conditional transition probabilities

$$\mathrm{P}(x_{t+\Delta t} \,|\, x_t, m_t)$$

are known, where the integer-valued variable m_t is used to denote the number of plants active during the time interval $[t, t + \Delta t)$. Equivalently, we can model future profits and losses using a discrete-time dynamic model of the form

$$x_{t+\Delta t} = f(x_t, m_t, w_t),$$

where w_t is a random parameter that is characterized by the conditional distribution $\mathrm{P}(\cdot|x_t, m_t)$.

To mathematically formulate and numerically solve this problem, we employ a dynamic programming framework in discrete time over a discounted horizon. We begin by first providing the stage costs associated with the time interval $[t + k\Delta t, t + (k + 1)\Delta t)$, hereafter denoted as stage k:

$$\begin{aligned} g(x_k, m_k, m_{k+1}) &= m_k x_k - c_I(M - m_{k+1}) \\ &\quad - c_A[m_{k+1} - m_k]_+ \\ &\quad - c_D[m_k - m_{k+1}]_+, \end{aligned}$$

where, with abuse of notion, we employ the following simplifying definitions: $x_k \overset{\Delta}{=} x_{t+k\Delta t}$ and $m_k \overset{\Delta}{=} m_{t+k\Delta t}$. The notation $[\cdot]_+$ is used to denote a function that returns the argument only when it is positive and zero otherwise (i.e. $[x]_+ \overset{\Delta}{=} \max\{x, 0\}$). Note that for the k^{th} stage, the state variables are x_k and m_k and the manipulated variable is m_{k+1}.

The first term to the right of the equality gives the operating profit or loss associated with running m_k active plants, the second the cost associated with keeping $N - m_k$ plants idle, the third the cost associated with increasing the number of active plants at the next time interval ($m_{k+1} > m_k$), and the fourth the cost associated with decreasing the number at the next time interval ($m_{k+1} < m_k$). Overall, the stage cost $g_k(\cdot)$ provides the operating margins given existing utilization m_k and changes to future utilization m_{k+1} during the time interval $[t + k\Delta t, t + (k + 1)\Delta t)$.

The total discounted cost over the horizon $[t, t + N\Delta t)$ is

$$\Phi_N = \sum_{k=0}^{N-1} \alpha^k g(x_k, m_k, m_{k+1}) + \alpha^N m_N x_N,$$

where the parameter α, a real-valued number between zero and one, denotes the discounting factor. The reason that we consider a discounted cost is to account for the time value of money. In general, current profits are more desirable than those in the distant future.

Following Bertsekas [1], we formulate the control problem then as finding the set of policies

$$\pi = \{\mu_0, \mu_1, \ldots, \mu_{N-1}\}$$

that maximizes the expected profits over a discounted time horizon

$$\max_{\pi} \mathrm{E}\{\Phi_N \mid x_0, m_0\},$$

where μ_k is the control law that specifies the number of active plants, m_{k+1}, at time index $k + 1$ as a function of the current profits and number of active plants at time index k (i.e. $m_{k+1} = \mu_k(x_k, m_k)$). Table 1 summarizes the governing problem variables and parameters.

We can solve this optimization problem by applying the dynamic programming algorithm:

$$\begin{aligned} J_k(x_k, m_k) &= \max_{m_{k+1}} \mathrm{E}\{g(x_k, m_k, m_{k+1}) \\ &\quad + \alpha J_{k+1}(x_{k+1}, m_{k+1}) \mid x_k, m_k\}. \end{aligned}$$

Table 1: Variable and parameter definitions

M	Number of plants
N	Horizon length
x_k	Operating margin at time k
m_k	Number of active plants at time k
c_A	Cost to activate idle plant
c_D	Cost to idle active plant
c_I	Cost to keep plant idle
α	Discount factor
γ	Risk factor

We refer to these sorts of policies as risk-neutral control because the governing optimization problem equally balances profits and losses. In the next section, we discuss the solution structure and provide a numerical example. In subsequent sections, we consider the case where the producer is risk sensitive.

From the perspective of control, these sorts of problems are interesting as the setpoint is not defined a priori. Rather, it is determined by producer prices and operating costs, variables that are not necessarily fixed but more often than not fluctuate in an unknown manner in response to market conditions. As a comparison, most process control problems seek to robustly maintain and stabilize a system about some known setpoint in the face of external disturbances. In the problem considered in this work, the setpoint is not known at future times and instead only the relevant conditional probabilities are. In other words, the problem is to design a controller that will enable a system to track a stochastic process. This problem that falls within the general class of stochastic pursuit-evasion games. However, the problem is more naturally formulated using classical approaches based on dynamic programming and Markov decision processes.

Solution Structure: Risk-Neutral Control

The two-state ($N = 1$) problem admits the analytic solution:

$$m_{k+1} = \begin{cases} M & \text{if } \bar{x}_{k+1} \geq (c_A - c_I)/\alpha \\ m_k & \text{if } (c_A - c_I)/\alpha \geq \bar{x}_{k+1} \geq -(c_I + c_D)/\alpha \\ 0 & \text{if } \bar{x}_{k+1} \leq -(c_I + c_D)/\alpha, \end{cases}$$

where $\bar{x}_{k+1} \overset{\Delta}{=} E\{x_{k+1} \mid x_k\}$. In other words, if the profits associated with increasing production exceed the costs associate with increasing capacity utilization then one should operate at full capacity. Likewise, if the costs associated with decreasing production are less than the losses associated with decreasing utilization then one should idle all production capacity. Interestingly, the solution exhibits hysteresis or a deadband. In addition, only two controls are realized: $\mu_k(\cdot) = 0$ or M. Note also that the optimal policy does not depend on the number of active plants at time k: $\mu_k = \mu_k(x_k)$.

The linear structure of the cost function means that the solution exhibits certainty equivalence. Similarly, one can show that the N-stage problem exhibits the same bang-bang type control with a deadband hysteresis (this same type of bang-bang type control also arises in optimal control problems involving linear objective functions [7]). In particular, the N-stage problem admits a solution of the form

$$m_{k+1} = \begin{cases} M & \text{if } \bar{x}_{k+1} \geq u_N \\ m_k & \text{if } l_N \geq \bar{x}_{k+1} \leq u_N \\ 0 & \text{if } \bar{x}_{k+1} \leq l_N, \end{cases}$$

where the bounds u_N and l_N are functions of the parameters and horizon length. As will be shown in the numerical example below, the only change as the horizon length N increases is that the size of the deadband $u_N - l_N$ deceases as the longer horizon affords more aggressive control.

Example: Risk-Neutral Control

As an illustrative numerical example, consider the case where $\bar{x}_{k+n} = x_k$ for all $n \geq 0$ (i.e. x_k is a discrete-time martingale). The results are shown in Figures 1 and 2 for the parameter values: $M = 10, c_A = 1.0, c_D = 1.0, c_I = 0.5, \alpha = 0.95$). The key observation, as noted previously, is that the size of the threshold, $u_N - l_N$, decreases as the horizon length N increases. Basically, the control becomes more aggressive as longer horizons afford more opportunities for future recourse. Not unexpectedly, the benefits are greatest when there are many active plants when margins are negative or many idle plants when margins are positive because a dynamic policy enables corrective action.

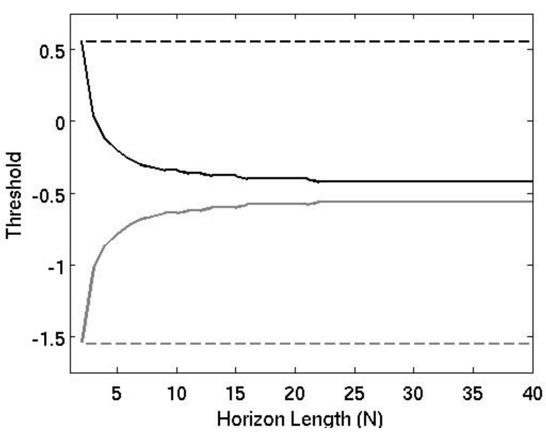

Figure 1: Switching threshold as function of horizon length. The black line denotes u_N and the gray line l_N. The dashed lines denote the respective thresholds for $N = 1$ as determined from the analytical solution.

Risk-Sensitive Control

One limitation of the preceding formulation is that the optimal policy does not directly account for stochastic fluctuations in the operating margins. The policy is the same if the volatility is large or small. In practice, one often desires a more cautious policy if the expected margins x_k are volatile versus a more aggressive one if they are not. This choice ultimately reduces to a question of risk tolerance: how much

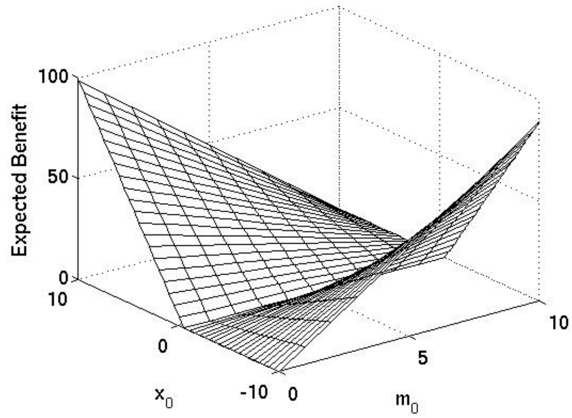

Figure 2: Expected benefit per time associated with employing a stationary dynamic policy ($N = \infty$) as opposed to static one ($\mu_k(x_k, m_k) = m_k$) as a function of the initial conditions x_0 and m_0.

risk is the producer willing to incur in order to maximize future profits? In the preceding example which we termed risk neutral, a large profit followed by an equally large loss is equally preferably to no profit at all. When the producer is risk sensitive, the later scenario is more preferable than the former.

In order to devise a risk-sensitive operating policy, we need to numerically represent the producer's tolerance for risk. The standard approach is to employ a utility function $U(x)$ that assigns numerical values to different possible outcomes based on the producer's tolerance to risk [4]. In the associated control problem, the optimal policy is the one that maximizes the expected utility, itself a function future profits, as opposed to simply maximizing the expected profits themselves.

In risk-sensitive decision problems, the utility function is monotonically increasing and concave (see Figure 3). Here, we consider a utility function of the form:

$$U(x) = -\exp(-\gamma x),$$

where the parameter $\gamma > 0$ is used to quantify our degree of risk sensitivity: the greater the value of γ the less tolerant one is to risk. The key attribute of this utility function is that it penalizes losses more than it rewards profits. Utility functions of this form are commonly employed in risk-sensitive control problems as they admit tractable dynamic programming solutions that preserve the staged structure of the problem. We note that risk-sensitivity policies based on utility functions of this form have parallels to $\mathcal{H}_\infty$ control [2].

We formulate the risk sensitive problem as finding the optimal policy that maximizes the expect utility over a discounted horizon

$$\max_{\pi} \mathrm{E} U(\Phi_N).$$

The solution to this problem can be obtained by applying a modified dynamic programming algorithm of the form

$$J_k(x_k, m_k) = \max_{m_{k+1}} \qquad \exp(-\gamma g(x_k, m_k, m_{k+1}) \times$$

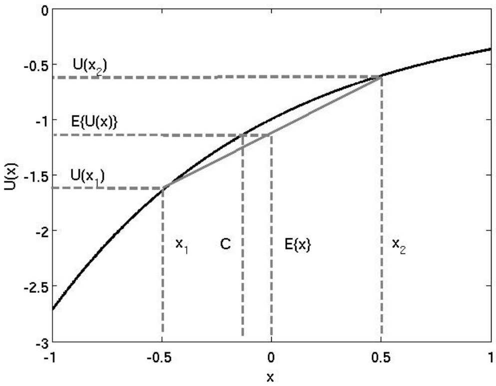

Figure 3: Graphical illustration of risk tolerance for the utility function $U(x) = -\exp(-x)$ (denoted by the black line). Following Luenberger [5], consider a proposition where there are two equally likely payoffs x_1 and x_2. Utility theory provides a framework for determining the value of such a proposition based on an individuals risk preferences. If the preferences are risk neutral, then the value of the proposition is the expected payoff $E\{x\}$, an even-money bet. However, if the preferences are governed by a utility function, then the assigned value is the expected utility $C = \mathrm{E}\{U(x)\}$. As the utility function is strictly concave when the preferences are risk sensitive, $C \leq \mathrm{E}\{x\}$ irrespective of the governing probabilities. In other words, a risk sensitive individual would equally value the known outcome $C < \mathrm{E}\{x\}$ to the random outcome of x_1 or x_2.

$$\mathrm{E}\{J_{k+1}(x_{k+1}, m_{k+1}) \mid x_k\}.$$

As before, we can easily solve this problem using a naive value iteration approach. The key difference with the previous formulation is that certainty equivalence does not apply in the risk-sensitive control problem as we need to evaluate the expectation of a exponential function: $\mathrm{E}\{\exp(y)\} \neq \exp(\mathrm{E}\{y\})$. In other words, we need to explicitly account for stochastic fluctuations in the operating margins x_k. This in turn requires an expression for the conditional transition probabilities.

Numerical Solution Procedure

In the remainder, we assume that the operating margins can be modeled using a simple, mean-reverting stochastic process in continuous time

$$dX_t = -\theta X_t dt + \sigma W_t,$$

where W_t is a standard Wiener process (see Figure 4). Here we use X_t to denote the random variable and x_t its numerical realization. The strength of this model is that parameters are easy to interpret in terms of the statistical properties of X_t: the stationary mean is zero and the correlation function is given by the expression $\frac{\sigma^2}{2\theta}\exp(-\theta\tau)$, where τ is the correlation time. From a pragmatic viewpoint, a mean-reverting process provides a simple mechanism for instantiating our prior beliefs regarding intrinsic volatility and temporal correlations in prices. Such processes provide simple models for commodity price dynamics [10].

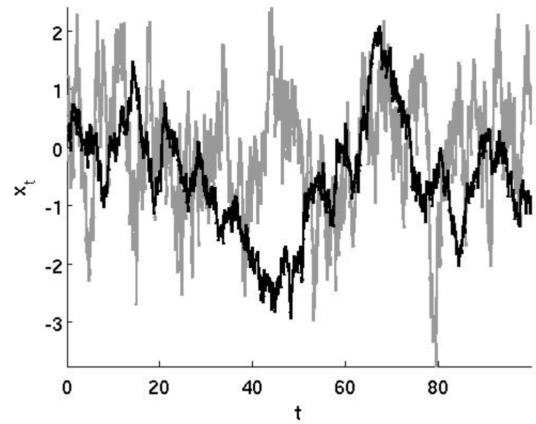

Figure 4: Illustration of two mean reverting processes with the same variance. The black curve ($\theta = 1$ and $\sigma = 0.71$) shows a realization of a process with a slow correlation time whereas the gray curve ($\theta = 0.1$ and $\sigma = 2.2$) shows one with a fast correlation time.

We can also view the operating margins as the aggregate of multiple, not necessarily stationary stochastic processes. However, the co-integration of these processes is most likely stationary with mean-reverting dynamics [6]. For example, while the prices of crude oil and gasoline individually cannot be accurately modeled using a simple mean-reverting process, their difference, the so-called crack spread, can to a first approximation.

To determine the governing transition probability, we need to solve the Fokker-Planck equation

$$\partial_t p(x,t) = \partial_x \theta x p(x,t) + \frac{\sigma^2}{2}\partial_x^2 p(x,t),$$

subject to the boundary conditions

$$p(x,t) \to 0 \text{ as } x \to \pm\infty$$

and

$$\int_{-\infty}^{\infty} p(x,t)dx = 1.$$

Note that the stationary solution to the Fokker-Planck equation is given by the Gaussian distribution

$$p_s(x) = \frac{\theta}{\pi\sigma^2}\exp\left(-\frac{\theta x^2}{\sigma^2}\right).$$

In order to solve the risk-control problem, we need to reformulate it as Markov decision process on a finite state space of size n. Here we assume that the operating margins x_k take on only fixed number of discrete values in the set

$$S = \{x_{\min}, \ldots, -2\delta, -\delta, 0, \delta, 2\delta, \ldots, x_{\max}\}$$

where $\delta = (x_{\max} - x_{\min})/(n+1)$. To adequately capture the range of the stochastic process, we chose $x_{max} = 3\sigma/\sqrt{\theta}$ and $x_{min} = 3\sigma/\sqrt{\theta}$. The probability that X_t remains within these bounds is greater than 99%. To discretize the Fokker-Plank equation, we can employ a first-order upwind method to discretize the convection operator and finite differences to discretize the Laplacian, both on a fixed domain with reflecting boundary conditions. This spatial discretization procedure approximates the Fokker-Planck equation with a finite-dimensional, linear differential equation of the form $\dot{p}(t) = Ap(t)$. We then can integrate this equation using a fixed time step Δt, yielding a Markov chain of the form $p_{k+1} = Pp_k$. Here, the i$^{\text{th}}$ element of the vector p_k gives

$$\mathrm{P}(x_k = s_i),$$

where s_i is used to denote the i$^{\text{th}}$ element in the set S ($s_i = x_{\max} - (i-1)\delta$). Similarly, the elements of the the matrix P define the transition probabilities for the Markov chain:

$$P_{ij} = \mathrm{P}(x_{k+1} = s_j \,|\, x_k = s_i).$$

By approximating the Fokker-Planck equation as a Markov chain, the resulting risk-sensitive dynamic program reduces to following simplified form

$$J_k(s_i, m_k) = \max_{m_{k+1}} \quad \exp(-\gamma g(s_i, m_k, m_{k+1}) \times \sum_{j=1}^{N} P_{ij} J_{k+1}(s_j, m_{k+1}).$$

Note that the expectation is replaced by a simple summation and that J_k is now an $n \times (N+1)$ matrix. As before, this problem can easily be solved using a naive value iteration approach.

Example: Risk-Sensitive Control

Figures 5 and 6 show the optimal stationary policies ($N = \infty$) for the problem with two different tolerances for risk. Aside from the risk parameter γ, all other parameter values were the same: $c_A = 2$, $c_D = 2$, $c_I = 0.5$, $\alpha = 0.95, \theta = 0.1, \sigma = 1$, and $\Delta t = 1$. When the sensitivity to risk is low (Figure 5), the optimal policy involves bang-bang type control with a threshold, the same as the risk-neutral scenario. However, as the risk tolerance increases (Figure 6), the optimal policy is no longer discontinuous. Reflecting the sensitivity to risk, the optimal policy leads to more cautious changes in capacity utilization. Idle plants are activated in proportion to the increase in expected margins. Interestingly, the threshold is still present, indicating that the same general solution structure is present. We also note that the control action is not symmetric on either side of the threshold. The reason is that the utility function itself is not symmetric.

When we evaluate the expected benefits associated with a risk-sensitive operating policy (Figures 7 and 8), we observe that it outperforms the no-control scenario only when existing conditions are unfavorable. The optimal policy is geared towards minimizing risk as opposed to maximizing profits. In these regards, a risk-sensitive policy functions like an insurance policy, hedging against potential disasters at the cost of a fixed premium.

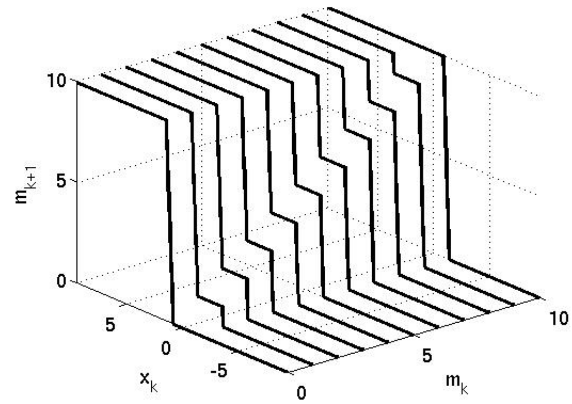

Figure 5: Optimal stationary policy for $\gamma = 0.01$.

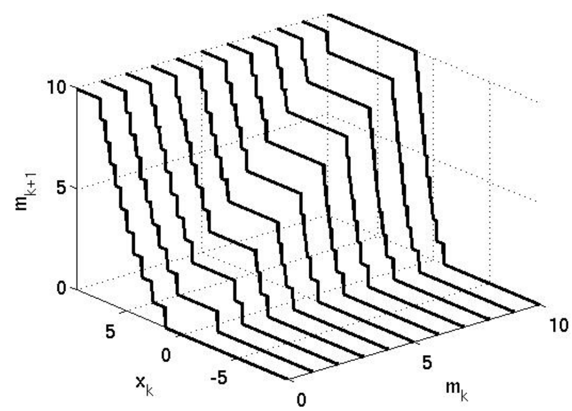

Figure 6: Optimal stationary policy for $\gamma = 0.2$.

Conclusions

In this paper, we explored a toy model of a marginal chemical producer with variable operating capacity subject to stochastic fluctuations in the operating margins. This problem was motivated by the goal of incorporating economics into process control problems. In particular, we wished to explore a problem where the setpoint is not a fixed or known value but rather a stochastic one reflecting market forces. Such a problem arises when producers vary production rates in response to variations in market prices and operating costs. From a theoretical point of view, this control problem alters the structure of the associated optimization problem such that the standard quadratic objective functions no longer apply. Moreover, the resulting policies are not longer continuous but rather involve discrete thresholds and hysteresis. In these regards, the resulting control laws are analogous to those used in hybrid systems and suggest that a new paradigm may be appropriate.

We note that analogous problems have previously been explored in the classic dynamic programming and operations research literature. Also, our formulation is somewhat naive given the complexity of problems currently tackled using stochastic programming [9]. In these regards, we do not claim any novel discoveries or algorithms. Rather, our goal was to reappraise a classic problem in light of current market trends, namely the commoditization of the chemical industries, the influence of dynamic markets, and the associated

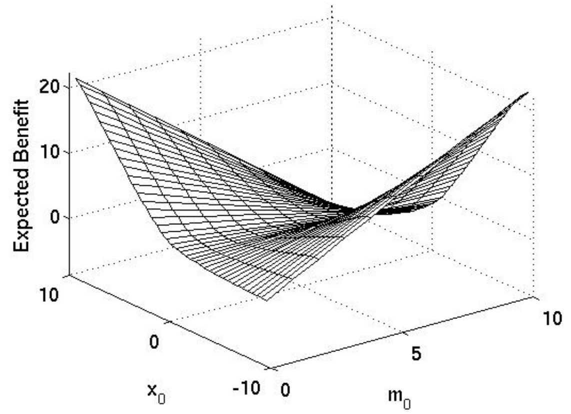

Figure 7: Expected benefit per time in terms of operating margins associated with employing risk-sensitive dynamic policy with $\gamma = 0.01$ as opposed to static one ($\mu_k(x_k, m_k) = m_k$) as a function of the initial conditions x_0 and m_0.

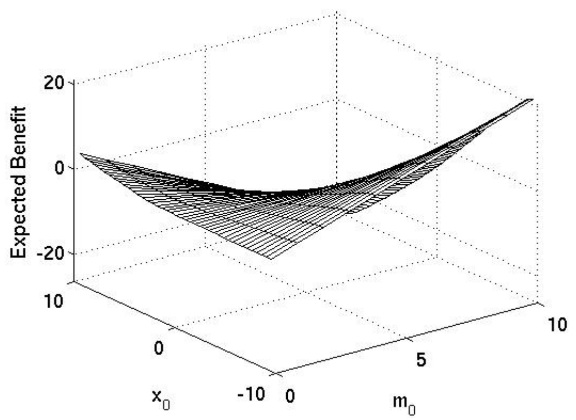

Figure 8: Expected benefit per time in terms of operating margins associated with employing risk-sensitive dynamic policy with $\gamma = 0.2$ as opposed to static one ($\mu_k(x_k, m_k) = m_k$) as a function of the initial conditions x_0 and m_0.

goal of developing control strategies for flexible manufacturing processes. Future efforts are directed towards more realistic problems involving multiple market factors and the dynamics of the governing processes themselves. The goal of current work was to explore what these controllers may look like and how to formulate the governing optimization problems.

Clearly the major limitation of this work is the use of dynamic programming. One cannot mention dynamic programming without acknowledging the "curse of dimensionality". While this admittedly simple problem can easily be solved using naive approaches, the direct application of dynamic programming would be intractable on more complex and realistic ones. That said, the dynamic programming framework is powerful one. In addition, there have been a number of significant advances over the years in developing approximate strategies for solving previously intractable dynamic programs, model predictive control being the most notable.

We conclude by commenting briefly on the parallels to economic model predictive control (EMPC) [8]. Our view is that present work and EMPC take complementary approaches to the same problem, namely of incorporating economics directly into the control problem. The notable attribute of the present work is that it addresses price volatility. That said, process dynamics are ignored, which is not the case in EMPC. Another key difference is that the present work discounts profits and losses over the prediction horizon whereas EMPC does not. The choice depends on the governing timescales of the process, which are considered long relative to the time value of money in the present work and short in EMPC.

Acknowledgments

I wish to thank Donald Chmielewski and Jeffrey Kantor for motivating this work and providing critical feedback.

References

[1] D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, 1995.

[2] K. Glover and J. C. Doyle. State-space formulae for all stabilizing controllers that satisfy an $\mathcal{H}_\infty$-norm bound and relations to risk sensitivity. *System & Control Letters*, 11:167–172, 1988.

[3] C. C. Holt, F. Modigliani, and H. A. Simon. A linear decision rule for production and employment scheduling. *Management Science*, 2:1–30, 1955.

[4] R. A. Howard and J. E. Matheson. Risk-sensitive Markov decision prcesses. *Management Science*, 18:356–369, 1994.

[5] D. G. Luenberger. *Investment Science*. Oxford University Press, New York, 1998.

[6] M. P. Murray. A drunk and her dog: an illustration of cointegraton and error-correction. *The American Statistician*, 48:37–39, 1994.

[7] C. V. Rao and J. B. Rawlings. Linear programming and model predictive control. *Journal of Process Control*, 10:283–298, 2000.

[8] J. B. Rawlings and R. Amrit. Optimizing process economic performances using model predictive control. In L. M. et al, editor, *Nonlinear Model Predictive Control*, pages 199–138. Springer-Verlang, 2009.

[9] N. V. Sahinidis. Optimization under uncertainty: state-of-the-art and opportunities. *Computers & Chemical Engineering*, 28:971–983, 2003.

[10] E. S. Schwartz. The stochastic behavior of commodity prices: implications for valuation and hedging. *The Journal of Finance*, 52:923–973, 1997.