OPERATIONS RESEARCH
http://dx.doi.org/10.1287/opre.1110.0959ec

e-companion

ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—"Integrated Optimization of Procurement, Processing, and Trade of Commodities" by Sripad K Devalkar, Ravi Anupindi, and Amitabh Sinha, *Operations Research*, http://dx.doi.org/10.1287/opre.1110.0959.

Heuristic and Upper Bound for Multi-Factor Models

While the analytical results in Section 3 are not dependent on the specific price process, the numerical experiments implemented in Section 4 assumed all the commodity prices follow single factor, mean-reverting processes. In this section, we describe a heuristic to compute policy parameters when the output commodity price dynamics follow multi-factor processes. We also describe a computationally tractable upper bound on the optimal expected profits and perform numerical studies to evaluate the performance of the heuristic.

EC.1. Computation of Heuristic Marginal Values

The primary difficulty in computing the optimal policy with multi-factor price processes is the fact that modeling the joint evolution of more than three factors becomes computationally inefficient. We now describe a tractable heuristic to compute approximate policies for such cases. The heuristic is based on approximating the input spot and output forward price for each maturity as single factor processes. For instance, if the output commodity spot price dynamics follow a multi-factor processes, the output forward price for each maturity, F_n^ℓ, can be modeled as a single factor, driftless geometric Brownian motion with the Brownian motion increments for different maturities ℓ, m correlated with a correlation factor $\rho_{\ell m} \in (-1, 1)$.

In any period n, we only model the joint evolution of the input spot price and the nearest maturing output forward price. More precisely, define

$$\hat{\mathcal{I}}_n = (S_n, F_n^\ell, F_1^{\ell+1}, F_1^{\ell+2}, \ldots, F_1^L) \text{ for } n \text{ such that } N_{\ell-1} \leq n < N_\ell \tag{EC.1}$$

The variable $\hat{\mathcal{I}}_n$ approximates the information available in period n by only considering S_n and F_n^ℓ, while assuming no information other than the initial prices of the remaining contracts is known. Thus, in the interval, $N_{\ell-1} \leq n < N_\ell$, we only consider the joint evolution of (S_n, F_n^ℓ) and take all expectations conditional on $\hat{\mathcal{I}}_n$. This approach is similar to the information approximation used in the approximate dynamic programming model of Lai et al. (2010a).

Next, we approximate the marginal value of output inventory given in equation (4) by conditioning the expectations on $\hat{\mathcal{I}}_n$ as follows.

$$\hat{\Delta}_n = \begin{cases} 0 & \text{if } n \geq N_L \\ \beta \max\left\{F_n^\ell, \mathbb{E}_{\hat{\mathcal{I}}_n}\left[\hat{\Delta}_{n+1}\right]\right\} - h_O & \text{if } n = N_\ell - 1 \text{ for } \ell = 1, \ldots, L \\ \beta \mathbb{E}_{\hat{\mathcal{I}}_n}\left[\hat{\Delta}_{n+1}\right] - h_O & \text{otherwise} \end{cases} \tag{EC.2}$$

We approximate marginal value of input inventory in a similar manner. That is,

$$\hat{\Theta}_n^k = \max\{\hat{\Omega}_n^{(k+b)}, \min\{S_n, \hat{\Omega}_n^{(k)}\}\} \tag{EC.3}$$

where

$$\hat{\Omega}_n^{(k)} = \max\left\{\beta \mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Theta}_{n+1}^k] - h_I, \min\left\{\hat{\Delta}_n - p, \beta \mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Theta}_{n+1}^{k-a}] - h_I\right\}\right\} \tag{EC.4}$$

for $n < N$ and all positive integers k and $\hat{\Theta}_N^k = S_N$ for all positive integers k. For all $n < N$, we set $\hat{\Theta}_n^k \triangleq \infty$ for $k \leq 0$.

The heuristic procurement, processing and commitment quantities $(\hat{x}_n, \hat{m}_n, \hat{q}_n)$ are then given by equations (16), (17) and (5) respectively, with the approximate marginal values replacing the true marginal values. Notice that this heuristic requires only modeling the joint evolution of single factor price processes in any given period. Thus, the binomial discretization approaches mentioned in section 4.1 can be used to compute the approximate marginal values efficiently. Further, the heuristic is exact in the case when the input and output commodity prices truly follow single factor price processes. We now describe a computationally tractable upper bound on the optimal expected profits, which can be used to evaluate the performance of the heuristic.

EC.2. Upper Bound on Optimal Expected Profits

We construct an upper bound for the optimal expected profits using the approach of information relaxation and dual penalties described in Brown et al. (2010). The key idea is that when information constraints are relaxed, i.e., more information is available at the time of decision than in the original problem, the solution to the relaxed problem will be an upper bound on the solution to the original problem. This is similar to relaxing the constraints in a linear program. Analogous to the dual variables corresponding to the constraints in a linear program which penalize violations of the constraints in the original problem, Brown et al. (2010) define feasible dual penalties for information relaxations, such that for any appropriately defined feasible dual penalty, the solution to the relaxed problem provides an upper bound to the optimal solution of the original problem. We use this technique to compute an upper bound on the optimal expected profits of the original problem.

We consider the perfect information relaxation for developing an upper bound on the optimal expected profits; that is, we consider a information structure where the input spot prices and output forward prices for all periods are known at the beginning of the horizon. Let $\Gamma_N = (\mathcal{I}_n)_{n=1}^N$ be a particular sample path of prices over the entire horizon. In period n, let $z_n(e_n, q_n, x_n, m_n, \Gamma_N)$ be a feasible dual penalty. For a specific Γ_N, let $H_n^{UB}(e_n, Q_n; \Gamma_N)$ be defined as

$$H_N^{UB}(e_N, Q_N; \Gamma_N) = S_N e_N \tag{EC.5}$$

$$\begin{aligned} H_n^{UB}(e_n, Q_n; \Gamma_N) = \max_{q_n, x_n, m_n \in \mathcal{B}_n} \Bigg\{ \left[\beta^{N_\ell - n} F_n^\ell - h_O \sum_{t=0}^{n_\ell - n - 1} \beta^t \right] q_n - p m_n - S_n \times x_n - h_I e_{n+1} \\ - z_n(e_n, q_n, x_n, m_n, \Gamma_N) + \beta H_{n+1}^{UB}(e_{n+1}, Q_{n+1}; \Gamma_N) \Bigg\} \\ \text{for } n = 1, 2, \ldots, N-1 \end{aligned} \tag{EC.6}$$

where $e_{n+1} = e_n + x_n - m_n$ and $Q_{n+1} = Q_n + m_n - q_n$. The set of feasible decisions, $\mathcal{B}_n$, is given by

$$\mathcal{B}_n = \left\{ (q_n, x_n, m_n) : \begin{array}{ll} 0 \le x_n \le K & \\ 0 \le m_n \le \min\{e_n + x_n, C\} & \\ q_n = 0 & \text{if } n \neq N_\ell - 1 \text{ for } \ell \in \{1, 2, \ldots, L\} \\ 0 \le q_n \le Q_n + m_n & \text{if } n = N_\ell - 1 \text{ for } \ell \in \{1, 2, \ldots, L\} \end{array} \right\}$$

Notice that H_n^{UB} is the same as V_n given by equations (2)–(3), except for the penalty term z_n and the fact that decisions involved in evaluating H_n^{UB} are made under perfect information. Define $V_1^{UB}(e_1, Q_1, \mathcal{I}_1)$ as

$$V_1^{UB}(e_1, Q_1, \mathcal{I}_1) = E_{\mathcal{I}_1}[H_1^{UB}(e_1, Q_1; \Gamma_N)] \tag{EC.7}$$

where the expectation is taken over all Γ_N.

Using different dual feasible penalties gives different values of V_1^{UB}. For instance, by setting the dual penalty $z_n = 0$ identically for all n, we get the perfect information upper bound equal to the optimal profit when the decision maker has perfect foresight. Using a feasible dual penalty that is easy to compute and approximates the ideal penalty closely can be expected to provide a close upper bound on the optimal expected profits. Consequently, we consider dual penalties derived from the approximate value-to-go function

$$\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1}) = \hat{\Delta}_{n+1} Q_{n+1} + \hat{\Theta}^k_{n+1} e_{n+1} + \hat{\lambda}^k_n \text{ for } e_{n+1} \in [(k-1)D, kD)$$

where the marginal values, $\hat{\Delta}_{n+1}$ and $\hat{\Theta}^k_{n+1}$, are given by equations (EC.2) and (EC.3) and $\hat{\lambda}^k_{n+1}$ are constants such that $\hat{V}_{n+1}$ is continuous in e_{n+1} and $\hat{\lambda}^1_n = 0$ for all n.

We then have

Proposition EC.1. *$V_1^{UB}(e_1, Q_1, \mathcal{I}_1)$ as defined in equation (EC.7), with dual penalties given by*

$$z_n(e_n, q_n, x_n, m_n, \Gamma_N) = \beta \left[\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1}) - \mathbb{E}_{\hat{\mathcal{I}}_n}\left[\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1}) \right] \right] \tag{EC.8}$$

is an upper bound on the optimal value function $V_1(e_1, Q_1, \mathcal{I}_1)$.

Proof*:* The dual penalty in equation (EC.8) is a feasible penalty and hence, by Proposition 3.1 in Brown et al. (2010), $V_1^{UB}(e_1, Q_1, \mathcal{I}_1) \geq V_1(e_1, Q_1, \mathcal{I}_1)$. □

Notice that the DP given by (EC.6) is a deterministic DP for each Γ_N. Thus the upper bound V_1^{UB} can be computed using Monte Carlo simulation by solving a deterministic optimization problem for each sample path, and averaging over sample paths.

EC.3. Numerical Study

We study the performance of the heuristic described in section EC.1 by comparing the expected profits using the heuristic with expected profits under the full commitment policy, and the upper bound on optimal expected profits. We consider a single composite output for the purposes of illustrating the heuristic (as seen from previous results, the composite output approximation itself does not lead to significant loss of optimality). The performance of the heuristic is quantified in Section EC.3.2. In the next section, we describe the implementation of the heuristics. Our implementation of the heuristic follows Lai et al. (2010a) closely.

EC.3.1. Implementation

We model the risk-neutral dynamics of the output forward price with maturity at time T_ℓ, $F(t,T_\ell)$, by a driftless geometric Brownian motion and with constant volatility $\sigma_\ell > 0$ as $\frac{dF(t,T_\ell)}{F(t,T_\ell)} = \sigma_\ell dW_\ell(t)$ where $dW_\ell(t)$ is the increment of a standard Brownian motion. The Brownian motion increments corresponding to forward prices with maturities T_ℓ and T_k have a constant correlation coefficient $\rho_{\ell k} \in (-1,1)$. The Brownian motion increment corresponding to forward price with maturity T_ℓ has a constant correlation coefficient $\rho_{\ell s} \in (-1,1)$ with the Brownian motion increment corresponding to the input spot price. We continue to model the input spot price as a mean-reverting process, as described in section 4.1.

The volatilities of the forward prices and correlation coefficients were estimated using historical data for futures contracts traded on CBOT. These parameters for the four immediately maturing contracts were estimated using the closing futures price on each trading day in the months of June and July, for the years 2001 to 2010. We use only June and July trading dates because futures contracts with maturities every month are not available during other calendar months of the year (soybean meal and oil futures contracts traded on CBOT have maturities in Jan., Mar., May, Jul., Aug., Sept., Oct., and Dec.). The volatilities of the four output forward contracts and the correlation matrix are given in Table EC.1.

The dynamics of the input spot and output forward prices are discretized using the same procedure described in section 4.1. We construct recombining binomial trees to represent the joint evolution of $(S(t), F(t,T_\ell))$ for each $\ell \in \{1,2,\ldots,L\}$, conditional on F_0^k for $k > \ell$. We also generate binomial trees to represent the evolution of $(F(t,T_\ell), F(t,T_{\ell,\ell+1}))$ for each $\ell \in \{1,2,\ldots,L-1\}$.

We obtain a probability mass function $G_n^\ell(S_{n+1}, F_{n+1}^\ell | S_n, F_n^\ell)$ for each $n \le n_\ell - 1$, for each node in the (S, F^ℓ) tree at time n. From the $(F^\ell, F^{\ell+1})$ tree, we obtain a probability mass function $\hat{H}_{n_\ell-1}^\ell(F_{n_\ell-1}^{\ell+1} | F_{n_\ell-1}^\ell)$ which denotes the probability that the next immediately maturing forward price is equal to $F^{\ell+1}$, conditional on the immediately maturing forward price being equal to F^ℓ. The probability mass function $G_n^\ell(\cdot)$ is used to compute expectations, conditional on $\hat{\mathcal{I}}_n$ for $n_{\ell-1} \le$

Table EC.1 Multi-factor Price Parameters

(a) Output Forward Price Volatility

Maturity (Month)	Volatility
1	0.269
2	0.252
3	0.249
4	0.257

(b) Correlation Matrix

	Input	Output Forward Maturity			
		1	2	3	4
Input	1	0.921	0.914	0.885	0.823
1		1	0.946	0.894	0.825
2			1	0.976	0.927
3				1	0.976
4					1

$n < n_\ell - 1$. The probability mass function $\hat{H}_n^\ell(\cdot)$, along with $G^\ell(\cdot)$ is used to compute expectations at the boundary of contracts ℓ and $\ell+1$. Specifically, we use

$$\hat{G}^\ell_{n_\ell-1}\left(S_{n+1}, F^{\ell+1}_{n+1} | S_{n_\ell-1}, F^\ell_{n_\ell-1}\right) = G^{\ell+1}_{n_\ell-1}\left(S_{n+1}, F^{\ell+1}_{n+1} | S_{n_\ell-1}, F^{\ell+1}_{n_\ell-1}\right) \times \hat{H}^\ell_{n_\ell-1}\left(F^{\ell+1}_{n_\ell-1} | F^\ell_{n_\ell-1}\right)$$

to approximate the transition probabilities at the expiration of forward contract ℓ for $\ell < L$.

We compute the heuristic marginal values $\hat{\Delta}_n$ and $\hat{\Theta}_n^k$ for each period n at each node in the binomial tree by using the input spot price value S_n at the node, forward price F_n^ℓ for n such that $n_{\ell-1} \le n < n_\ell$, and the probability mass functions G_n^ℓ for $n_{\ell-1} \le n < n_\ell - 1$ and $\hat{G}^\ell_{n_\ell-1}$ for $n = n_\ell - 1$. The heuristic policy parameters $(\hat{x}_n, \hat{m}_n, \hat{q}_n)$ are computed based on the values of $\hat{\Delta}_{n+1}$ and $\hat{\Theta}^k_{n+1}$ stored at each node. We evaluate the policies using Monte Carlo simulation and compute the expected profits from using the heuristic by averaging the performance on the generated sample paths. We also compute the upper bound for each sample path by solving the mixed-integer program described in Section EC.2.

The operational parameters are the same as in section 4.1; $p = 72$, the physical holding costs are zero, $\beta = 1$ and the procurement and processing capacities are set to 5 and 3 units.

Table EC.2 Performance of Heuristic for Different Horizon Lengths

Horizon Length (N) (Maturities N_ℓ)	Expected Profits FC	Expected Profits Heuristic	Upper Bound (UB)	Gap (% of UB) FC	Gap (% of UB) Heuristic
10 (5,10)	1353.46	1417.35	1491.23	9.24%	4.95%
15 (5,10,15)	1950.31	1981.83	2117.66	7.90%	6.41%
20 (5,10,15,20)	2594.98	2565.21	2868.54	9.55%	10.58%

EC.3.2. Performance of Heuristic

We consider the procurement, processing and trade operations for the firm over the period June to October. We initialize the input spot price to its long run average value, while the output forward prices are set to the average closing price over June 2010. Table EC.2 gives the optimal expected profits for different horizon lengths when using the full commitment and heuristic policies, along with the gap with respect to the upper bound.

We find that the heuristic is able to capture the option value inherent in the commitment decision, especially for short horizon lengths, and outperforms the full commitment policy.

As seen in section 4.2, the processing capacity and price volatilities had a significant impact on the performance of the full commitment policy. We investigate the impact of varying these parameters on the performance of the heuristic.

Impact of processing capacity. As seen from section 4.2, the performance of the full commitment policy deteriorates as the processing capacity becomes tight. Table EC.3 shows the expected profits under the full commitment and heuristic policies for different processing capacities. These numerical results are for a horizon length of 10 periods, with two forward contracts available for the output commodity.

For the same procurement capacity, the fraction of the total profits contributed by the output sales are lower for tighter processing capacities. As a result, it is important to extract the full value of the option to postpone commitment. As the results in Table EC.3 show, the heuristic captures this option value and performs well even for tight processing capacities. In comparison, the full

Table EC.3 Performance of Heuristic for Different Processing Capacities ($N = 0$, $L = 2$, $N_\ell = \{5, 10\}$)

Processing Capacity (C) (as % of K)	Expected Profits FC	Expected Profits Heuristic	Upper Bound (UB)	Gap (% of UB) FC	Gap (% of UB) Heuristic
20%	451.15	556.31	595.70	24.26%	6.61%
40%	902.30	1000.90	1050.05	14.07%	4.68%
60%	1353.46	1417.35	1491.23	9.24%	4.95%
80%	1804.61	1833.51	1914.71	5.75%	4.24%
100%	2255.76	2248.66	2324.93	2.98%	3.28%

Table EC.4 Impact of Price Volatility ($N = 10$, $L = 2$, $N_\ell = \{5, 10\}$)

Volatility (σ_s, σ_1, σ_2)	Expected Profits FC	Expected Profits Heuristic	Upper Bound (UB)	Gap (% of UB) FC	Gap (% of UB) Heuristic
0.30	1338.45	1404.70	1529.48	12.49%	8.16%
0.35	1345.89	1421.64	1593.85	15.56%	10.80%
0.40	1355.89	1444.18	1672.86	18.95%	13.67%
0.45	1367.55	1457.94	1763.00	22.43%	17.30%
0.50	1380.21	1486.18	1864.14	25.96%	20.28%
0.55	1393.52	1556.54	1974.13	29.41%	21.15%

commitment policy has a gap as high as 24.26% with respect to the upper bound for very tight capacities.

Impact of price volatilities. As price volatilities increase, we expect the value of the options inherent in the various decisions to increase. As the results in Table EC.4 show, the value of the upper bound increases more than the expected profits under both the policies as the price volatilities increase. The heuristic is able to capture the option value inherent in the various decisions better than the full commitment policy, as seen from the results in Table EC.4.

In summary, the numerical results for both, single factor as well as multi-factor, price processes illustrate the advantage of using integrated decision making. Taking the option like properties of the various decisions into account, even if it be through approximations, provides a significant improvement in profits compared to myopic heuristics such as the full-commitment policy.