MICHIGAN
ROSS SCHOOL OF BUSINESS

Working Paper

Integrated Optimization of Procurement, Processing and Trade of Commodities in a Network Environment

Sripad K. Devalkar
Stephen M. Ross School of Business
University of Michigan

Ravi Anupindi
Stephen M. Ross School of Business
University of Michigan

Amitabh Sinha
Stephen M. Ross School of Business
University of Michigan

Ross School of Business Working Paper
Working Paper No. 1095
April 2010

UNIVERSITY OF MICHIGAN

Integrated Optimization of Procurement, Processing and Trade of Commodities in a Network Environment

Sripad K Devalkar, Ravi Anupindi, Amitabh Sinha
Stephen M. Ross School of Business, University of Michigan, Ann Arbor, Michigan 48109, devalkar@umich.edu, anupindi@umich.edu, amitabh@umich.edu

We consider the integrated optimization problem of procurement, processing and trade of commodities over a network in a multiperiod setting. Motivated by the operations of a prominent commodity processing firm, we model a firm that operates a star network with multiple locations at which it can procure an input commodity and has processing capacity at a central location to convert the input into a processed commodity. The processed commodity is sold using forward contracts, while the input itself can be traded at the end of the horizon. We show that the single-node version of this problem can be solved optimally when the procurement cost for the input is piecewise linear and convex, and derive closed form expressions for the marginal value of input and output inventory. However, these marginal values are hard to compute because of high dimensionality of the state space and we develop an efficient heuristic to compute approximate marginal values. We also show that the star network problem can be approximated as an equivalent single node problem and propose heuristics for solving the network problem. We conduct numerical studies to evaluate the performance of both the single node and network heuristics. We find that the single node heuristics are near-optimal, capturing close to 90% of the value of an upper bound on the optimal expected profits. Approximating the star network by a single node is effective, with the gap between the heuristic and upper bound ranging from 7% to 14% for longer planning horizons.

Key words: Integrated Optimization, Commodities, Network

History: December 2008, Revised: August 2009, April 2010

1. Introduction

The motivation for our work comes from the innovative practices of one of India's largest private sector companies, The ITC Group (`www.itcportal.com`). The International Business Division (IBD) of ITC, started in 1990, exports agricultural commodities such as soybean meal, rice, wheat and wheat products, lentils, shrimp, fruit pulps, and coffee. Increased competition, along with an inefficient farm-to-market supply chain made it imperative for ITC-IBD to re-engineer the procurement process for commodities in rural India. Specifically, in the year 2000 ITC-IBD (hereafter referred to as ITC) embarked on the *e-Choupal* initiative to deploy information and communication technology (ICT) to reengineer the procurement of commodities from rural India. By purchasing directly from the farmers, and not just the local spot markets, ITC significantly improved the efficiency of the channel and created value for both the farmer and itself. The initiative has been hailed as an outstanding example of the use of ICT by a private enterprise to streamline supply chains, alleviate poverty and bring about social transformation. The e-Choupal platform has been extremely successful for ITC and has been well documented by Prahalad (2005) and Anupindi and Sivakumar (2006).

The e-Choupal platform for commodity procurement consists of a hub-and-spoke network where spokes correspond to village level ICT kiosks (called *e-Choupals*) consisting of a personal computer with internet access and the hubs are procurement centers or processing plants where direct deliveries occur (called the *direct-channel*). ITC creates a one-day forward market for procurement of commodities by announcing an offered price at each of its hubs. Typically, the forward price offered for the next period is the realized spot price in the current period. Farmers can access the e-Choupal kiosks for various information including ITC's prices, but have the option to sell their produce in the local spot market or directly to ITC at their hub location. One of the benefits to the farmers of selling directly to ITC is that the farmers are guaranteed same day service, which is not usually the case when they sell in the spot market. In order to satisfy the same day service guarantee, ITC places an upper limit on the total quantity that it will purchase through the direct channel in any period. In addition to the direct channel, ITC can also procure in the local spot market, if necessary. By 2007, there were close to 6000 e-Choupals and 140 procurement hubs in the network, with soybean being one of the largest commodities procured by ITC using the e-Choupal network. A schematic of the eChoupal network for soybean is shown in Figure 1.

Close to seventy percent of the soybean procured is processed at several processing plants; the rest is traded. Beans are processed to produce soybean oil and soybean meal, both of which are

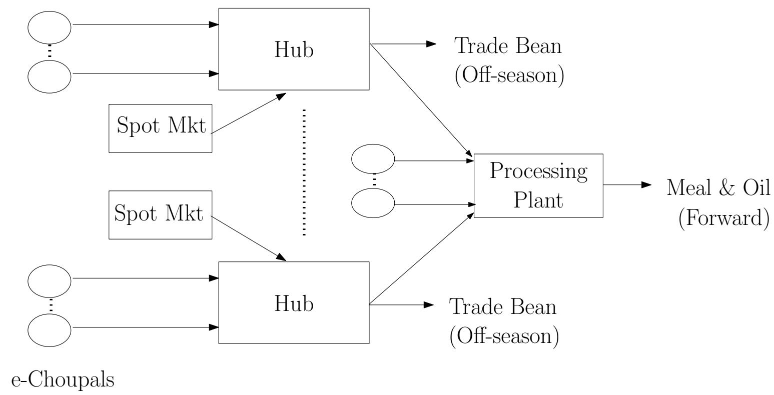

Figure 1 ITC e-Choupal Network.

traded through various channels. Managing this network requires decisions regarding procurement and trading of different commodities to maximize profits. Procurement decisions, which include price and quantity decisions for each hub, need to be integrated with the sales decision in terms of the form of output commodity and channels to trade in; that is, for the soybean procured, ITC needs to make decisions regarding whether to trade the bean or process it and trade the oil and meal. Trade options for the various commodities include trading in open markets and with other processors. We wish to determine the optimal policy for managing a commodity storage and processing network such as the *e-Choupal* network. Specifically, we are interested in the relationship between procurement, processing and trade decisions for the various commodities and the impact of operational constraints such as procurement and processing capacities on these decisions. While ITC's operations provide the basic context for our research, the problem considered in this paper is applicable in a more general context to firms in the commodities processing business. Profits for such firms is affected by both input and output commodity prices in globally traded exchanges and local spot markets. The procurement, processing and trade decisions for such firms are interdependent because of operational constraints and ignoring these dependencies can result in significant loss of value.

We consider a multiperiod optimization problem, in which a firm procures an input commodity across multiple locations, with the marginal cost of procurement dependent on prices realized in spot markets for the commodity. The firm earns revenues by processing the input commodity at a central processing location and committing to sell the processed output using forward contracts in every period. In addition, the firm can also trade the input inventory with other processors at the end of the horizon. In the single node version of this problem, we show that there exist two inventory and price dependent thresholds such that it is optimal to process up to capacity if the expected revenue from trading the output commodity is greater than the higher threshold. Similarly, it is optimal to procure the input commodity up to capacity if the spot price is less than the lower threshold. For values in between the two thresholds, the procurement and processing quantities are interdependent, but can be quantified. We derive recursive expressions to determine these thresholds in closed form when the procurement cost is linear or piecewise linear and convex. These expressions can be computed efficiently when a single forward contract is available to trade the output commodity over the entire planning horizon. However, when multiple forwards with different maturities are available to sell the output, these expressions are computationally intractable because of the high dimensionality of the state space required to model the dynamics of the various price processes. We develop efficient heuristics to compute approximate values of these thresholds and near-optimal policies. We analyze the network problem when the firm operates a star network; i.e., a network with a central processing and trade location connected to multiple procurement locations. We show that the star network problem is equivalent to the single node problem with convex cost of procurement when the transshipment costs between the nodes is negligible compared to the commodity prices and develop a heuristic to solve the network problem. Our numerical studies show that the heuristics for the single node and network problem perform quite well, capturing more than 90% of the value of an upper bound on the optimal expected profits in many cases.

The problem considered in this paper is related to the warehouse management problem studied by Bellman (1956) and Dreyfus (1957), and later extended by Charnes et al. (1966). The warehouse problem is one of determining the optimal trading policy for a commodity with constraints on the total inventory of the commodity that can be stored. Charnes et al. (1966) show that the value function is linear in the starting inventory level and derive expressions for the marginal value of inventory. These papers, however, do not consider constraints on the procurement and sales; i.e., it is assumed that any desired quantity of the commodity can be procured or sold in a period. More recently, Secomandi (2009b) considers a similar problem in the context of managing a natural gas storage asset. In addition to storage constraints, the paper also incorporates injection and withdrawal constraints and establishes the optimality of a price dependent double base-stock policy. While there are similarities, the problem addressed in the current paper has some significant differences, namely: a) we consider multiple commodities, in contrast to the single commodity trading decisions addressed in the warehouse management problem, b) in addition to procurement and trade of commodities, we also consider the additional decision of irreversibly transforming some of the commodities and c) our analysis includes operations over a network.

The single node problem considered here has similarities to the firm level production and inventory control problem studied in Wu and Chen (2009) for a storable input-output commodity pair. While Wu and Chen (2009) consider the optimal procurement and sales policy for the individual firm, their main focus is on analyzing the propagation of demand and supply shocks across production stages and the price-inventory relationship across input-output commodities using a rational expectations equilibrium model. Routledge et al. (2001) also consider a multi-commodity processing and storage network, but focus on deriving a rational expectations equilibrium model that can be used to extend the theory of storage to non-storable commodities like electricity and explain some of the empirically observed features of electricity prices. In contrast, we are interested in characterizing the optimal policy and deriving managerial insights for a firm operating a commodity processing business. As such, we do not adopt an equilibrium approach and instead model the evolution of the various commodity prices as exogenously given.

The procurement, processing and trade decisions considered in this paper are related to the valuation of real options. In the current problem, a unit of output inventory can be committed for sale against any of the forward contracts that are yet to expire. Thus, the marginal value of a unit of output in any period is similar to valuing a compound exchange option (cf. Carr (1988)) with the underlying assets being the various forward prices. While not the focus, the heuristics we develop in this paper can be used to approximate the value of such compound exchange options.

The concept of spread options is also closely related to the problem considered here, especially the processing decision. Spread options are call or put options on the spread between the prices of two commodities and arise naturally in the context of commodity industries. The valuation of spread options has typically been considered in single period setting; i.e., the valuation of a spread option with specific maturity date or situations where the exercise of the spread option maturing on one date does not affect the value or optimal exercise policy of a spread option maturing at a later date. Geman (2005) provides a discussion of different spread options in the commodity industries; e.g., crush spreads for agricultural commodities (soybean, for instance), crack spread (crude oil and refined petroleum products), location spreads (natural gas prices at different locations), calendar spreads (difference in natural gas forward prices for different maturities). Secomandi (2009a) considers the valuation of pipeline capacity used to transport natural gas across two locations using spread option valuation models on the spread between the natural gas prices at the two locations. Similarly, Deng et al. (2001) use spark spread options on the spread between electricity and generating fuel prices and location spread options on the spread between electricity prices at different locations to value generation and transmission assets. In a closely related context, Plato (2001) examines the decision of US soybean processors to commit processing capacity to crush soybeans and produce soybean meal and oil. This decision is similar to the exercise of a spread option on the gross processing margin at a future date, i.e., the spread between the futures price of soybean meal and oil and soybean, with the exercise price being equal to the variable cost of processing. In this paper, the decision to process a unit of input is akin to exercising a spread option on the difference between the values of a unit of output and input, with the processing cost as the exercise price. Processing (and procurement) decisions across periods are, however, linked through the storage of input inventory and operational capacity constraints. This crucial difference makes the processing decision considered here different from the exercise of a simple spread option considered in the extant literature.

The star network analyzed in this paper is based on the features of the *e-Choupal* network with different procurement hubs serving a single processing plant that also has an associated procurement facility. While more complex commodity production and distribution networks have been considered in literature, (cf., Markland (1975), Markland and Newett (1976)), these papers assume deterministic commodity prices and have no capacity constraints. In contrast, we consider stochastic commodity prices and capacity constraints on procurement and processing which make the problem non-trivial.

The rest of the paper is organized as follows. In Section 2, we solve the integrated procurement, processing and trade decisions for a risk-neutral firm operating a single node. We solve the single node problem completely and obtain expressions for the marginal value of inventory. Section 2.1 presents the analysis for the linear procurement cost case, while Section 2.2 describes the convex, piecewise linear procurement cost situation. We describe computation of the optimal policy when a single forward contract is available to make output sale commitments and develop computationally tractable heuristics for the more general case with multiple forward contracts with different maturities in Section 2.3.1. We provide numerical examples of the computations in Section 2.4. We analyze the network problem in Section 3 and develop heuristics to solve the star network problem in Section 3.1. Section 4 concludes the paper with directions for future research.

2. The Single Node Problem

2.1. Linear Procurement Cost

We consider a finite horizon problem with the time periods indexed by $n = 1, 2, \ldots, N-1, N$ where $n = 1$ is the first decision period. In any period n, let S_n denote the price for the input in the spot market. The procurement season for the input commodity may span multiple output forward maturities. For instance, the soybean meal and oil forward contracts traded on the Chicago Mercantile Exchange (CME) have maturity months of January, March, May, July, August, September, October and December - implying multiple forward contracts expiring during the procurement season (September–March/April). We consider L forward contracts available for selling the output during the planning horizon. The forward contracts are indexed by ℓ, with $\ell \in \{1, 2, \ldots, L\}$ and maturity N_ℓ. We assume $N_\ell - 1$ is the last possible period in which the firm can sell the output using forward contract ℓ. Without loss of generality, we assume $N_\ell < N_{\ell+1}$ for all $\ell < L$. Let F_n^ℓ denote the period n forward price on contract ℓ, for $n < N_\ell \le N$. The firm sells all the output using forward contracts. In addition, the firm can also trade the input itself with other processors over the horizon. For ease of exposition, we assume that all, if any, input sales happen at the end of the horizon with a per-unit trade (or salvage) value of S_N. Agricultural commodities exhibit seasonality with increased supply soon after harvest periods. With enough supply available, bulk of the processors' procurement from spot markets happen during this period, termed the 'procurement season'. Trade between processors is typically low during these periods. In our context, the planning horizon can be considered as the procurement season, when bulk of the procurement happens. End of the horizon can be thought of representing the off-season, when most of the trading of the input (soybean) between processing firms happens.

Due to physical or other operational limitations, the firm has a per-period procurement capacity restriction of K units and a processing capacity of C units per period to convert the input into a processed product (also referred to as 'output'). The unit cost of processing one unit of the input commodity into the output commodity is p. The firm incurs a per period holding cost of h_I and h_O per unit of input and output inventory respectively. We assume $h_O \ge h_I$. Initially, we consider a linear cost of procurement, i.e., the cost of procuring x units of input is equal to $S_n \times x$ when the input spot price is equal to S_n. Later, in Section 2.2, we extend the analysis to include convex, piecewise linear cost of procurement.

The relevant information available to the firm at the beginning of period n regarding the spot market prices, output forward prices and trade prices for the input is given by $\mathcal{I}_n$ and all expectations are taken under the risk-neutral measure. We assume interest rates are constant and there is no counter-party risk associated with the forward contracts. As a result, the discount factor per period, β, is the risk-free discount factor. It is a well known result that under these conditions, the forward prices for the output are a martingale process (see Hull (1997), Section 3.9 or Bjork (2004), Section 7.6 for details). We thus have

$$\mathbb{E}_{\mathcal{I}_n}[F_{n+1}^{\ell}] = F_n^{\ell} \text{ for } n < N_{\ell}, \forall\, \ell \tag{1}$$

where $\mathbb{E}_{\mathcal{I}_n}[\cdot]$ denotes expectation, conditional on $\mathcal{I}_n$.

The variable $\mathcal{I}_n$ can include the realized spot market price, forward prices and other state variables which impact the commodity prices; e.g., aggregate inventory levels of the commodities. Our formulation of the integrated procurement, processing and trade problem does not depend on the specific model used to represent the dynamics of the various input and output prices.

In each time period $n \leq N-1$, the firm makes the following sequence of decisions: a) the quantity of the input commodity to be procured: x_n, b) the quantity of the output commodity to be committed for sale against forward contract ℓ in period n: q_n^{ℓ} for all ℓ such that $N_{\ell} > n$ and c) the quantity of input to be processed into output in period n: m_n. In the last period, N, the firm trades any remaining input inventory. Optimal values of these decisions will be denoted by a '*' superscript. Let Q_n (respectively, e_n) denote the total output (respectively, input) inventory available at the beginning of period n.

It is easy to see that in any given period it is optimal to commit against at most one forward contract. Thus, let $\hat{\ell}$ be the forward contract that the firm commits against in period n, if a commitment is made. Notice that the firm can potentially commit to sell more output than is currently available; i.e., 'over-commit' such that $q_n^{\hat{\ell}} > Q_n + m_n$. This is possible because the output needs to be delivered only in period $N_{\hat{\ell}}$ and the firm can process in some future period(s) t between n and $N_{\hat{\ell}}$ to meet the shortfall $q_n^{\hat{\ell}} - (Q_n + m_n)$, which would require that we keep track of the shortfall against each forward contract. However, in light of the martingale property (equation (1)), we can see that such a 'anticipatory commitment' strategy would never be optimal and thus the firm will never over-commit. Therefore, we do not need to keep track of the shortfall against each forward contract and $(e_n, Q_n, \mathcal{I}_n)$ is sufficient to describe the state of the system at the beginning of period n. Further, because commitments once made cannot be reversed, we can recognize the revenues associated with output sales at the time of making the commitment rather than at the time of delivery without loss of generality. Thus, if a commitment is made in period n, it would be against forward contract $\hat{\ell}$ where $\hat{\ell} = \arg\max_{\ell} \left\{ \beta^{N_\ell - n} F_n^\ell - h_O \sum_{t=0}^{N_\ell - n - 1} \beta^t \right\}$ where the term inside the parenthesis is the discounted forward price minus the total discounted holding costs incurred from the current period till delivery at the maturity of the foward contract. We can formulate the firm's problem as a stochastic dynamic program (SDP) in the following manner.

$$V_n(e_n, Q_n, \mathcal{I}_n) = \max_{\substack{0 \le x_n \le K, \\ 0 \le m_n \le \min\{C, e_n + x_n\}, \\ 0 \le q_n^{\hat{\ell}} \le Q_n + m_n}} \Bigg\{ \left[\beta^{N_{\hat{\ell}} - n} F_n^{\hat{\ell}} - h_O \sum_{t=0}^{N_{\hat{\ell}} - n - 1} \beta^t \right] q_n^{\hat{\ell}} - S_n x_n - p m_n - h_I [e_n + x_n - m_n] - h_O [Q_n + m_n - q_n^{\hat{\ell}}] + \beta \mathbb{E}_{\mathcal{I}_n} [V_{n+1}(e_{n+1}, Q_{n+1}, \mathcal{I}_{n+1})] \Bigg\} \quad (2)$$

for $n < N$ and

$$V_N(e_N, Q_N, \mathcal{I}_N) = \begin{cases} S_N e_N & e_N \ge 0 \\ -\infty & \text{otherwise} \end{cases} \quad (3)$$

where the state transition equations are given by

$$e_{n+1} = e_n + x_n - m_n \quad (4)$$

$$Q_{n+1} = Q_n + m_n - q_n^{\hat{\ell}} \quad (5)$$

The constraints on x_n and m_n in equation (2) are capacity and input availability constraints. The constraint on the commitment quantity is the no 'over-commitment' condition, which is without loss of optimality and ensures $(e_n, Q_n, \mathcal{I}_n)$ is sufficient to describe the state of the system.

Consider the commitment decision $q_n^{\hat{\ell}}$. The firm earns a revenue of $\beta^{N_{\hat{\ell}} - n} F_n^{\hat{\ell}} - h_O \sum_{t=0}^{N_{\hat{\ell}} - n - 1} \beta^t$ on each unit committed for sale. The firm can earn the same expected revenue (discounted to period n dollars) by postponing the commitment to period $N_{\hat{\ell}} - 1$, the last opportunity to commit against contract $\hat{\ell}$. By postponing the decision to period $N_{\hat{\ell}} - 1$, the firm retains the option to not commit the unit of output to contract $\hat{\ell}$ if some other contract ℓ' provides a higher revenue. Extending this argument, we have the following result.

LEMMA 1. *It is optimal to commit to sell output using contract ℓ, for $\ell = 1, 2, \ldots, L$, only in period $N_\ell - 1$ and only if $F^\ell_{N_\ell - 1}$ is at least as much as the expected benefit from committing against the remaining $L - l$ contracts. Let $\Phi^{(\ell)}_{N_\ell - 1}$ denote the marginal value of output inventory in period $N_\ell - 1$, when $L - l + 1$ forward contracts are still available. Then, $\Phi^{(\ell)}_{N_\ell - 1}$ is given by*

$$\Phi^{(\ell)}_{N_\ell - 1} = \begin{cases} \beta F^L_{N_L - 1} - h_O & \text{if } \ell = L \\ \max\left\{\beta F^\ell_{N_\ell - 1} - h_O,\ \beta^{N_{\ell+1} - N_\ell} \mathbb{E}_{N_\ell - 1}\left[\Phi^{(\ell+1)}_{N_{\ell+1} - 1}\right] - h_O \sum_{t=0}^{N_{\ell+1} - N_\ell - 1} \beta^t\right\} & \text{if } \ell < L \end{cases} \tag{6}$$

and the optimal commitment quantity against contract ℓ is given by

$$q^{\ell *}_{N_\ell - 1} = \begin{cases} 0 & \text{if } \beta F^\ell_{N_\ell - 1} - h_O < \Phi^{(\ell)}_{N_\ell - 1} \\ Q_n + m_n & \text{otherwise} \end{cases}$$

To gain further intuition about $\Phi^{(\ell)}_{N_\ell - 1}$, consider the case when there are only two forward contracts available for selling the output, with maturities N_1 and N_2 respectively. In period $N_1 - 1$, the firm will commit the available output inventory for sale against contract 1, if doing so provides higher revenue and not otherwise. Thus, the value of a unit of output inventory in period $N_1 - 1$ is equal to $\max\left\{\beta F^1_{N_1 - 1} - h_O,\ \beta^{N_2 - N_1 + 1} F^2_{N_1 - 1} - h_O \sum_{t=0}^{N_2 - N_1} \beta^t\right\}$. This value is simply the payoff from an exchange option on the two discounted forward prices, after adjusting for holding costs. Equation (6) generalizes this to the case when there are L contracts available to commit the output against. At the maturity of contract ℓ, the value of a unit of output is equal to the maximum of the revenue from contract ℓ and the maximum expected benefit from committing against one of the remaining $L - \ell$ contracts at a later date.

Notice that the optimal commitment policy is an 'all or nothing' policy; i.e., if it is optimal to commit against contract ℓ in period $N_\ell - 1$, then it is optimal to commit all the available output inventory, $Q_{N_\ell - 1} + m_{N_\ell - 1}$. Using an induction argument and the result in Lemma 1, we can prove the following result.

LEMMA 2. *The value function $V_n(e_n, Q_n, \mathcal{I}_n)$ is separable in e_n and Q_n, and is linear in Q_n. The marginal value of a unit of output inventory in period n is given by Δ_n, where*

$$\Delta_n = \begin{cases} \left[\beta^{(N_\ell - 1) - n} \mathbb{E}_{\mathcal{I}_n}\left[\Phi^{(\ell)}_{N_\ell - 1}\right] - h_O \sum_{t=0}^{(N_\ell - 1) - n - 1} \beta^t\right] & \text{if } N_{\ell - 1} \le n < N_\ell \text{ for } \ell = 1, 2, \ldots, L \\ 0 & \text{if } n \ge N_L \end{cases} \tag{7}$$

The fact that there are no capacity constraints on the output sale commitments ensures the value function is linear in Q_n. We can write

$$V_n(e_n, Q_n, \mathcal{I}_n) = \Delta_n Q_n + U_n(e_n, \mathcal{I}_n) \text{ for } n < N \text{ and} \tag{8}$$

$$V_N(e_N, Q_N, \mathcal{I}_N) = U_N(e_N, \mathcal{I}_N) \tag{9}$$

where $U_n(e_n, \mathcal{I}_n)$ is given by

$$U_n(e_n, \mathcal{I}_n) = \max_{\substack{0 \le x_n \le K, \\ 0 \le m_n \le e_n + x_n, \\ 0 \le m_n \le C}} \Big\{ [\Delta_n - p] m_n - S_n x_n - h_I [e_n + x_n - m_n] + \beta \mathbb{E}_{\mathcal{I}_n}[U_{n+1}(e_{n+1}, \mathcal{I}_{n+1})] \Big\} \quad \text{for } n < N \tag{10}$$

and

$$U_N(e_N, \mathcal{I}_N) = S_N e_N \tag{11}$$

Notice that in any period $n < N_\ell - 1$, the marginal value of a unit of output inventory is equal to the expected discounted payoff from the optimal commitment decision in period $N_\ell - 1$, after adjusting for holding costs. The payoff from optimal commitment in period $N_\ell - 1$ is nothing but the payoff of a compound exchange option on the remaining $L - \ell + 1$ forward contracts; i.e., an option to exchange revenue from the immediately maturing forward contract ℓ for a compound exchange option on the remaining $L - \ell$ forward contracts, after adjusting for holding costs. Thus, each unit of output inventory can be considered a compound exchange option, with the remaining forward contracts as the underlying assets (cf., Carr (1988)).

We next turn to determining the marginal value of input inventory. If the firm had infinite processing capacity, we can use very similar arguments and show that it would be optimal for the firm to process, if at all, only in periods $N_\ell - 1$. Further, the processing and commitment quantities would be equal to each other and equal $e_{N_\ell - 1} + x_{N_\ell - 1}$, the total available input inventory. In such a situation, the marginal value of input inventory would be equal to the value of a compound exchange option, where the underlying assets of the option include remaining forward contracts net of the processing cost p and the input trade price at the end of horizon, after adjusting for holding costs.

However, the firm does not have infinite processing capacity and cannot afford to limit processing only to periods $N_\ell - 1$. Thus, the true marginal value of input inventory would be less than the value of such a compound exchange option. Further, the value of a unit of input inventory would depend on the total input inventory available. For instance, when the input inventory at the beginning of period n is more than the remaining processing capacity till maturity of the last forward contract L, i.e., $e_n > [N_l - n]C$, the marginal value of input is equal to $\beta^{N-n}\mathbb{E}_{\mathcal{I}_n}[S_N] - h_I \sum_{t=0}^{N-n-1} \beta^t$, the discounted expected salvage value net of total discounted holding costs, irrespective of the value from processing $\Delta_n - p$. We now derive expressions for the marginal value of input inventory that facilitates evaluation of the decision to process in period n.

To this end, let D be the largest value such that the processing capacity $C = aD$ and the procurement capacity $K = bD$, where a and b are positive integers; i.e., D is the greatest common divisor of C and K.[1] Theorem 1 below states that $U_n(e_n, \mathcal{I}_n)$ is piecewise linear, with breaks at integral multiples of D and provides an expression for Θ_n^k, the marginal value of input inventory at the beginning of period n, when $e_n \in [(k-1)D, kD)$, where k is a positive integer. (While Θ_n^k clearly depends on the realization of $\mathcal{I}_n$ for all n and k, for notational convenience, we do not show this dependence explicitly.)

THEOREM 1. *The value function $U_n(e_n, \mathcal{I}_n)$ is continuous, concave and piecewise linear in e_n with changes in slope at integral multiples of D, for each realization of $\mathcal{I}_n$. Let Θ_n^k denote the marginal value of input inventory (i.e., slope of U_n) at the beginning of period n, when $e_n \in [(k-1)D, kD)$ where k is an integer.*

For all n, let $\Theta_n^k \triangleq \infty$ for $k \le 0$. In the last period, $\Theta_N^k = S_N$ for all $k \ge 1$. For any period $n < N$ and positive integer k, the marginal value of inventory Θ_n^k is given by

$$\Theta_n^k = \max\left\{\Omega_n^{(k+b)}, \min\left\{S_n, \Omega_n^{(k)}\right\}\right\} \tag{12}$$

where $\Omega_n^{(j)}$ is the marginal value of $e_n + x_n$, the input inventory after procurement in period n, when $e_n + x_n \in [(j-1)D, jD)$ and is given by

$$\Omega_n^{(j)} = \max\left\{\beta\mathbb{E}_{\mathcal{I}_n}[\Theta_{n+1}^j] - h_I, \min\left\{\Delta_n - p, \beta\mathbb{E}_{\mathcal{I}_n}[\Theta_{n+1}^{j-a}] - h_I\right\}\right\} \tag{13}$$

PROOF: Clearly, $U_N = S_N e_N$ is concave and piecewise linear in e_N for all $e_N \geq 0$. Further, $\Theta_N^k = S_N$ for all positive integers k. Consider the period $N-1$ problem. We have

$$U_{N-1}(e_{N-1}, \mathcal{I}_{N-1}) = \max_{\substack{0 \leq x_{N-1} \leq K, \\ 0 \leq m_{N-1} \leq \min\{C, e_{N-1}+x_{N-1}\}}} \Big\{ [\Delta_{N-1} - p] m_{N-1} - h_I [e_{N-1} + x_{N-1} - m_{N-1}] + \beta \mathbb{E}_{\mathcal{I}_{N-1}}[S_N] \times (e_{N-1} + x_{N-1} - m_{N-1}) \Big\}$$

U_{N-1} is the solution of a linear program and e_{N-1} appears in the right hand side of the constraints. Thus, U_{N-1} is piecewise linear and concave in e_{N-1}. Further, the change in slope of U_{N-1} occurs at integral multiples of D, since the processing and procurement capacities are integral multiples of D.

Suppose U_t is piecewise linear and concave, with change in slope at integral multiples of D for all $t = n+1, n+2, \ldots, N$. That is, for each $t \geq n+1$, we have

$$U_t(e_t, \mathcal{I}_t) = \Theta_t^k e_t + \lambda_t^k \text{ for } e_t \in [(k-1)D, kD)$$

where λ_t^k is a constant independent of e_t, U_t is continuous in e_t and $\Theta_t^k \geq \Theta_t^{k+1}$ for all integers $k \geq 1$.

Now, consider the SDP equation (10)

$$U_n(e_n, \mathcal{I}_n) = \max_{\substack{0 \leq x_n \leq K, \\ 0 \leq m_n \leq \min\{C, e_n + x_n\}}} \Big\{ [\Delta_n - p] m_n - S_n x_n - h_I [e_n + x_n - m_n] + \beta \mathbb{E}_{\mathcal{I}_n}[U_{n+1}(e_{n+1}, \mathcal{I}_{n+1})] \Big\}$$

By the induction assumption on U_{n+1}, $U_n(e_n, \mathcal{I}_n)$ is the solution of a linear program with e_n in the right hand side of the constraint. Thus, $U_n(e_n, \mathcal{I}_n)$ is concave and piecewise linear in e_n.

We can re-write the above maximization problem as

$$\begin{aligned} U_n(e_n, \mathcal{I}_n) &= \max_{0 \leq x_n \leq K} \Bigg\{ \max_{0 \leq m_n \leq \min\{C, e_n + x_n\}} \Big\{ [\Delta_n - p] m_n - h_I [e_n + x_n - m_n] \\ &\qquad + \beta \mathbb{E}_{\mathcal{I}_n}[U_{n+1}(e_n + x_n - m_n, \mathcal{I}_{n+1})] \Big\} - S_n x_n \Bigg\} \\ &= \max_{0 \leq x_n \leq K} \{ L_n(e_n + x_n, \mathcal{I}_n) - S_n x_n \} \quad \text{for } n < N \end{aligned}$$

where

$$L_n(y_n, \mathcal{I}_n) = \max_{0 \leq m_n \leq \min\{C, y_n\}} \Big\{ [\Delta_n - p] m_n - h_I [y_n - m_n] + \beta \mathbb{E}_{\mathcal{I}_n}[U_{n+1}(y_n - m_n, \mathcal{I}_{n+1})] \Big\}$$

By the induction assumption, $U_{n+1}(e_{n+1}, \mathcal{I}_{n+1})$ is concave and piecewise linear in e_{n+1}. Thus, L_n is the solution of a linear programming problem and hence piecewise linear and concave in y_n, where y_n is the input inventory after procurement, but before processing. For y_n and m_n such that $y_n - m_n \in [(j-1)D, jD)$ for some positive integer j, we can write the objective function in the maximization above as

$$\begin{aligned}&[\Delta_n - p]m_n + \left[\beta \mathbb{E}_{\mathcal{I}_n}[\Theta^j_{n+1}] - h_I\right][y_n - m_n] + \beta \mathbb{E}_{\mathcal{I}_n}[\lambda^j_{n+1}] \\ =\ & \left[[\Delta_n - p] - [\beta \mathbb{E}_{\mathcal{I}_n}[\Theta^j_{n+1}] - h_I]\right] m_n + \left[\beta \mathbb{E}_{\mathcal{I}_n}[\Theta^j_{n+1}] - h_I\right] y_n + \beta \mathbb{E}_{\mathcal{I}_n}[\lambda^j_{n+1}] \end{aligned} \tag{14}$$

where λ^j_{n+1} is a constant independent of y_n and m_n.

For a given y_n, as m_n increases, j such that $y_n - m_n \in [(j-1)D, jD)$ decreases. Therefore, as m_n increases, the coefficient of m_n, given by $[\Delta_n - p - [\beta \mathbb{E}_{\mathcal{I}_n}[\Theta^j_{n+1}] - h_I]]$, decreases since $\Theta^j_{n+1} \geq \Theta^{(j+1)}_{n+1}$. Thus, the optimal value of m_n is the maximum possible value for which the coefficient remains non-negative or zero, which ever is higher. For $y_n \in [(s-1)D, sD)$ where s is a positive integer and recalling that the processing capacity $C = aD$, we can determine the optimal value of m_n as

$$m^*_n = \begin{cases} C & \text{if } \beta \mathbb{E}_{\mathcal{I}_n}[\Theta^{s-a}_{n+1}] - h_I \leq \Delta_n - p \\ y_n - \hat{r}D & \text{if } \beta \mathbb{E}_{\mathcal{I}_n}[\Theta^{s}_{n+1}] - h_I \leq \Delta_n - p < \beta \mathbb{E}_{\mathcal{I}_n}[\Theta^{s-a}_{n+1}] - h_I \\ 0 & \text{if } \Delta_n - p < \beta \mathbb{E}_{\mathcal{I}_n}[\Theta^{s}_{n+1}] - h_I \end{cases} \tag{15}$$

where $\hat{r}D = \arg\max_r \left\{\beta \mathbb{E}_{\mathcal{I}_n}[\Theta^r_{n+1}] - h_I > \Delta_n - p\right\}$. Upon substituting m^*_n corresponding to each of the three cases in the objective function (14), we have for $y_n \in [(s-1)D, sD)$

$$L_n(y_n, \mathcal{I}_n) = \begin{cases} (\beta \mathbb{E}_{\mathcal{I}_n}[\Theta^{s-a}_{n+1}] - h_I)y_n + (\Delta_n - p - \beta[\mathbb{E}_{\mathcal{I}_n}[\Theta^{s-a}_{n+1}] - h_I])C + \beta \mathbb{E}_{\mathcal{I}_n}[\lambda^{s-a}_{n+1}] \\ (\Delta_n - p)y_n - (\Delta_n - p - [\beta[\mathbb{E}_{\mathcal{I}_n}[\Theta^{\hat{r}}_{n+1}] - h_I])\hat{r}D + \mathbb{E}_{\mathcal{I}_n}[\lambda^{\hat{r}}_{n+1}] \\ (\beta \mathbb{E}_{\mathcal{I}_n}[\Theta^{s}_{n+1}] - h_I)y_n + \mathbb{E}_{\mathcal{I}_n}[\lambda^{s}_{n+1}] \end{cases}$$

Thereby,

$$L_n(y_n, \mathcal{I}_n) = \max\left\{\beta \mathbb{E}_{\mathcal{I}_n}[\Theta^s_{n+1}] - h_I, \min\{\Delta_n - p, \beta \mathbb{E}_{\mathcal{I}_n}[\Theta^{s-a}_{n+1}] - h_I\}\right\} y_n + \Upsilon^s_n \tag{16}$$

where Υ^s_n denotes constant terms not dependent on y_n.

Notice that the slope of $L_n(\cdot, \cdot)$ with respect to y_n when $y_n \in [(s-1)D, sD)$ is equal to $\Omega^{(s)}_n$, where

$\Omega_n^{(s)}$ is given by equation (13). Thus, $\Omega_n^{(s)}$ denotes the marginal value of a unit of input inventory *after* procurement but *before* processing. We now have

$$U_n(e_n, \mathcal{I}_n) = \max_{e_n \le y_n \le e_n + K} \{L_n(y_n, \mathcal{I}_n) - S_n(y_n - e_n)\}$$

For $y_n \in [(s-1)D, sD)$, substituting $L_n(y_n, \mathcal{I}_n)$ from equation (16), the objective function in the maximization above can be written as $\left[\Omega_n^{(s)} - S_n\right] y_n + \Upsilon_n^s + S_n e_n$.

By the induction assumption, we have $\Theta_{n+1}^j \ge \Theta_{n+1}^{(j+1)}$ for all j and as a result $\Omega_n^{(s)}$ is non-increasing in s. Thus, the slope of y_n decreases as y_n increases. For $e_n \in [(k-1)D, kD)$ where k is a positive integer and recalling that the procurement capacity $K = bD$, we can determine the optimal value of y_n as

$$y_n^* = \begin{cases} e_n + K & \text{if } \Omega_n^{(k+b)} \ge S_n \\ \hat{s}D & \text{if } \Omega_n^{(k)} \ge S_n > \Omega_n^{(k+b)} \\ e_n & \text{if } S_n > \Omega_n^{(k)} \end{cases} \tag{17}$$

where $\hat{s} = \arg\max_s \left\{\Omega_n^{(s)} > S_n\right\}$. Substituting y_n^*, we get

$$U_n(e_n, \mathcal{I}_n) = \max\left\{\Omega_n^{(k+b)},\ \min\left\{S_n, \Omega_n^{(k)}\right\}\right\} e_n + \Psi_n^k$$

where Ψ_n^k is a constant independent of e_n.

Thus, $U_n(e_n, \mathcal{I}_n)$ is piecewise linear in e_n with breaks at integral multiples of D. Further, because $\Omega_n^{(k+1)} \le \Omega_n^{(k)}$, we have $\Theta_n^{k+1} \le \Theta_n^k$ for all non-negative integers k. □

The optimal processing quantity m_n^* given by equation (15) is based on comparing the value of $\Delta_n - p$ relative to the marginal value-to-go, $\beta \mathbb{E}_{\mathcal{I}_n}[\Theta_{n+1}^s] - h_I$, of the input inventory evaluated at y_n, the input inventory level *after* procurement. It is useful to state the optimal policy in terms of parameters that can be evaluated based on the state variables at the beginning of the period, instead. Substituting the optimal procure up to level for the input given by equation (17), we can re-state the optimal procurement and processing quantities for a given realization of $\mathcal{I}_n$ and a starting input inventory level e_n as follows.

PROPOSITION 1. *For all $n < N$, let $\Omega_n^{(k)}$ be as defined in equation (13). For a starting input inventory level e_n such that $e_n \in [(k-1)D, kD)$ where k is a positive integer,*

1. *The optimal procurement quantity is given by*

$$x_n^* = \begin{cases} K & \text{if } \Omega_n^{(k+b)} > S_n \\ \hat{s}D - e_n & \text{if } \Omega_n^{(k)} \geq S_n \geq \Omega_n^{(k+b)} \\ 0 & \text{if } \Omega_n^{(k)} < S_n \end{cases} \tag{18}$$

where $\hat{s} = \arg\max_{s \in \mathbb{Z}} \{\Omega_n^{(s)} > S_n\}$.

2. *The optimal quantity to process is given by*

$$m_n^* = \begin{cases} C & \text{if } \Omega_n^{(k)} < \Delta_n - p \\ \min\{(e_n + x_n^* - \hat{r}D)^+, C\} & \text{if } \Omega_n^{(k)} \geq \Delta_n - p \geq \Omega_n^{(k+b)} \\ 0 & \text{if } \Omega_n^{(k+b)} > \Delta_n - p \end{cases} \tag{19}$$

where $\hat{r} = \arg\max_{r \in \mathbb{Z}} \{\Omega_n^{(r)} > \Delta_n - p\}$.

The above result implies that in each period there exist two inventory and state dependent thresholds $\Omega_n^{(k)}$ and $\Omega_n^{(k+b)}$ with $\Omega_n^{(k)} \geq \Omega_n^{(k+b)}$ such that it is optimal to procure up to capacity (respectively, procure nothing) if the marginal cost of procurement S_n is less than the lower threshold (respectively, greater than the higher threshold). Similarly, it is optimal to process up to capacity (respectively, process nothing) if the benefit from processing $\Delta_n - p$ is greater than the higher threshold (respectively, less than the lower threshold). For values of S_n and $\Delta_n - p$ in between the two thresholds, the procurement and processing quantities are interdependent.

To illustrate the results in Proposition 1, consider the example where $C = K$; i.e., $a = b = 1$ and $D = C = K$. Figure 2 shows the value of $\Omega_n^{(k)}$ as a function of e_n (i.e., for different k), for a given realization of $\mathcal{I}_n$ where $S_n < \Delta_n - p$. By the definition of $\hat{s}$ and $\hat{r}$ given in Proposition 1, we have $\hat{s} = 7$ and $\hat{r} = 3$.

Region A, $k = \{1, 2\}$, $e_n < 2D$: At these levels of starting input inventory, the expected marginal value of unprocessed input inventory even after procurement up to capacity is greater than $\Delta_n - p$. Thus, it would not be optimal for the firm to process any input in this region. Further, the expected value of input inventory is greater than S_n, thus making procurement up to capacity optimal. Notice that in this region we have $\Omega_n^{(k+b)} = \Omega_n^{(k+1)} > \Delta_n - p$. This situation corresponds to the first and last cases respectively in the procurement and processing policy given by equations (18) and (19).

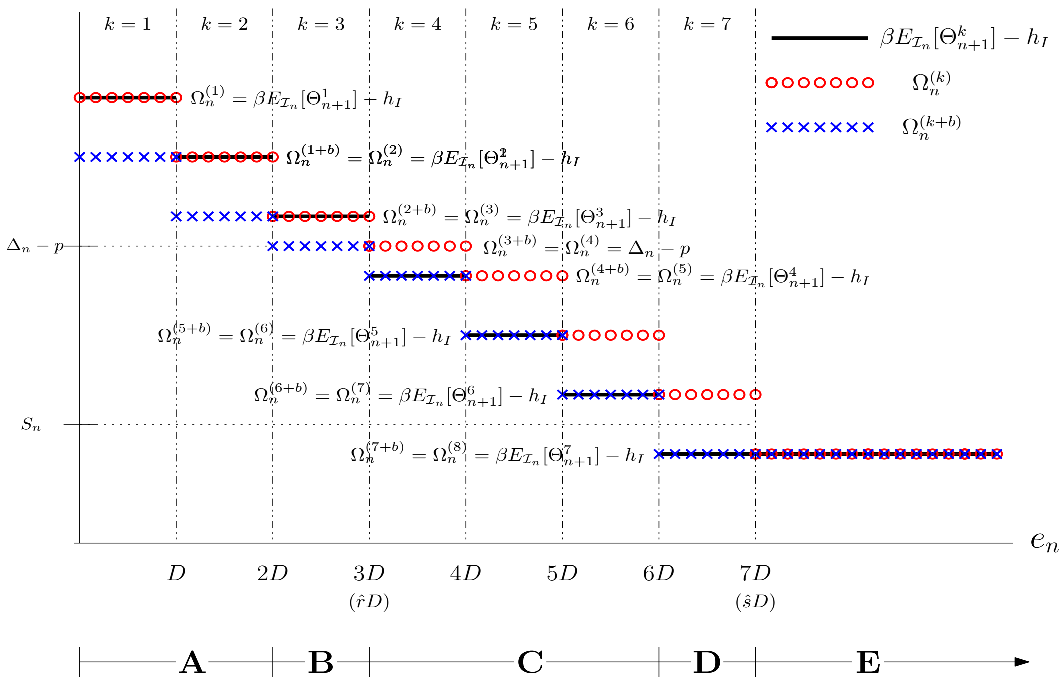

Figure 2 Illustration of optimal policy when $S_n \leq \Delta_n - p$, for $C = K$

Region B, $k = \{3\}$**,** $e_n \in [2D, 3D)$**:** In this region, procuring up to capacity will result in $e_n + x_n = e_n + D \geq \hat{r}D$. Thus, the value from processing $\Delta_n - p$ can be greater than the expected marginal value from keeping the input unprocessed at that inventory level. Thus, the optimal quantity to process is such that $e_n + x_n - m_n = \hat{r}D$. Because $\Delta_n - p > S_n$, there is an instantaneous margin from procurement and processing, and it is optimal for the firm to procure up to capacity and process $e_n + D - \hat{r}D$. Thus, in region **B**, even though $e_n < \hat{r}D$, we find that it is optimal to process a positive quantity, after procuring up to capacity. As in region **A**, we have $\Omega_n^{(k+b)} = \Omega_n^{(k+1)} > S_n$ in this region also, corresponding to the first case in the optimal procurement policy given by equation (18). We also have $\Omega_n^{(k+b)} = \Omega_n^{(4)} = \Delta_n - p < \Omega_n^{(3)} = \Omega_n^{(k)}$, corresponding to the second case in the optimal processing policy given by equation (19).

Region C, $k = \{4, 5, 6\}$**,** $e_n \in [3D, 6D)$**:** The value from processing, $\Delta_n - p$, is greater than the value from keeping the input unprocessed. Thus, it is optimal for the firm to process as long as the final input inventory level is at least $\hat{r}D$, below which the expected marginal value from unprocessed input is greater than $\Delta_n - p$. Also, since $\Delta_n - p > S_n$, it is always optimal to procure additional input to ensure the processing capacity is utilized fully. Thus, in this region it is optimal to process up to capacity. Further, the expected marginal value of input inventory after procurement up to capacity and processing is greater than S_n. Thus, it is also optimal to procure up to capacity in this region. Notice that in this region we have $\Delta_n - p \geq \Omega_n^{(k)} > \Omega_n^{(k+1)} > S_n$, corresponding to the first case in equations (18) and (19) respectively.

Region D, $k = \{7\}$**,** $e_n \in [6D, 7D)$**:** It is optimal to process up to capacity because $\Delta_n - p$ is greater than the expected marginal value of unprocessed input at $e_n - D \in [5D, 6D)$. Thus, the optimal processing quantity is limited only by the processing capacity in this region and any additional inventory procured in the current period will remain unprocessed. Further, it is optimal to procure additional input as long as the expected marginal value of the unprocessed input is greater than S_n; i.e., x_n is such that $e_n + x_n - C = 6D$ or $x_n = 7D - e_n = \hat{s}D - e_n \leq K$. Notice that in this region we have $\Delta_n - p > \Omega_n^{(k)} = \Omega_n^{(7)} > S_n > \Omega_n^{(8)} = \Omega_n^{(k+1)}$, corresponding to the second and first cases respectively in equations (18) and (19).

Region E, $k \geq 8$**,** $e_n \geq 7D$**:** Similar to the earlier case, we can see that processing up to capacity is optimal. Thus any additional input procured at these inventory levels will remain unprocessed. We have $e_n - m_n = e_n - D \geq 6D$. Since the expected marginal value of unprocessed input inventory at these levels is less than S_n, it is optimal not to procure any additional input. In this region, we have $\Delta_n - p > S_n > \Omega_n^{(k)}$, corresponding to the last and first cases respectively in equations (18) and (19).

When $S_n > \Delta_n - p$, as shown in Figure 3, we can similarly divide the state space corresponding to the beginning input inventory into: 1) regions **A** and **B**, where it is optimal to only procure (and procure up to capacity in region **A**), 2) region **C**, where it is optimal to do nothing and 3) regions **D** and **E**, where it is optimal to only process (and process up to capacity in region **E**).

In both the figures, we can see that the regions of positive procurement and processing quantities correspond to regions where $S_n \leq \Omega_n^{(k)}$ and $\Delta_n - p \geq \Omega_n^{(k+b)}$. Further, notice that $\hat{s}D$ and $\hat{r}D$ represent target 'procure up to' and 'process down to' inventory levels. This is similar to the target base stock levels in the single commodity, capacitated warehouse management problem (cf. Secomandi (2009b)). However, unlike the single commodity case, we can have instances where $\hat{s}D >$

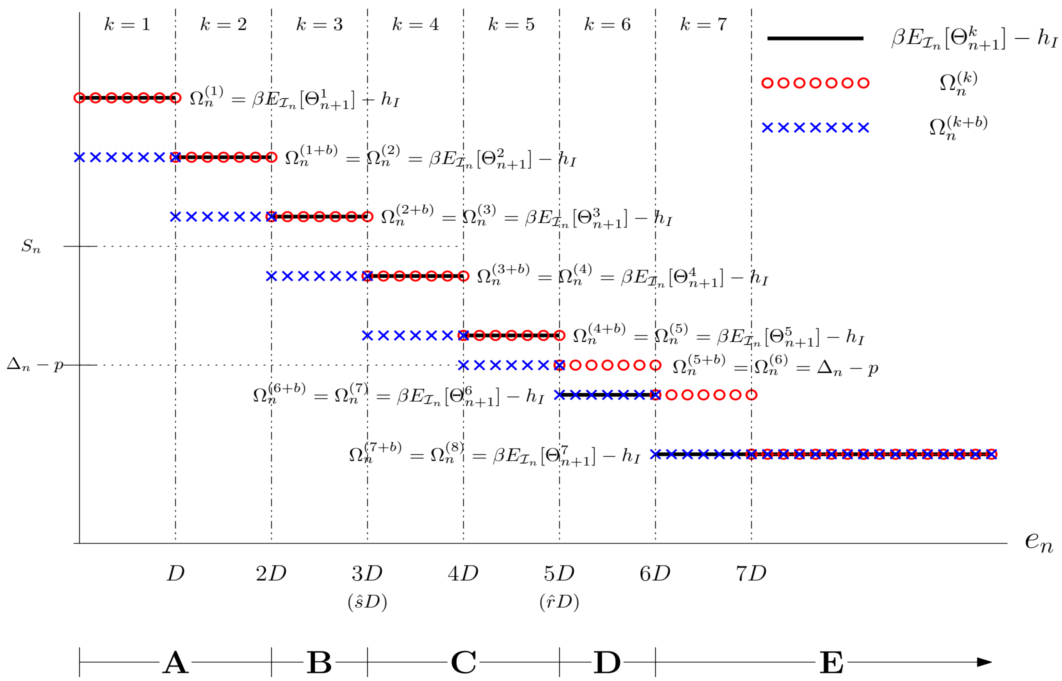

Figure 3 Illustration of optimal policy when $S_n > \Delta_n - p$, for $C = K$

$\hat{r}D$; i.e., there can be inventory levels at which *both* procurement and processing are optimal. This happens whenever there is an immediate margin from procurement and processing; i.e., $\Delta_n - p > S_n$, which is never possible in the single commodity procurement and trade problem.

In the next section, we extend the analysis to consider the more general situation of convex procurement costs encountered in the *e-Choupal* and other commodity processing contexts.

2.2. Convex Cost of Procurement

The analysis thus far assumed that the procurement cost is linear in the quantity procured and the firm pays the spot price per unit. This is generally true when the firm is small and the firm's actions do not affect the market prices. However, even for such firms the cost of procurement may not necessarily be linear. Consider ITC's *e-Choupal* network where at each hub procurement is through the direct channel as well as the spot market. Under such circumstances, the total cost of procurement over both sources would ideally be a piecewise linear convex function because of the 'merit order' of procurement (cf., Bannister and Kaye (1991)); i.e., the firm will procure from the cheaper source first before using the more costly channel.[2] Other instances where a convex cost of procurement may arise is when the firm procures over multiple locations to serve a single processing and trade location. As we discuss later in Section 3, the results developed for the single node convex procurement cost case will be useful when analyzing the integrated problem over a network. With this motivation, we consider the situation when the firm has a convex cost of procurement.

We assume all aspects of the operations remain the same as in Section 2, except for the procurement cost. Let the total cost of procuring x_n units of input when the spot price is S_n be denoted by $\mathscr{C}(S_n, x_n)$. We model $\mathscr{C}(S_n, x_n)$ as a piecewise linear, convex function such that

$$\mathscr{C}(S_n, x_n) = \begin{cases} \gamma^j S_n \times [x_n - K^{j-1}] + \alpha^j & \text{if } K^{j-1} < x_n \leq K^j \\ \gamma^1 S_n \times x_n & \text{if } 0 \leq x_n \leq K^1 \end{cases} \tag{20}$$

where $\gamma^j > \gamma^{j-1}$ and $K^j > K^{j-1}$ for all $j = 1, 2, \ldots, J$ and α^j are such that $\mathscr{C}(S_n, x_n)$ is continuous in x_n. Let $\gamma^0 = 0$ and $K^0 = 0$. Notice that the linear cost of procurement is a special case of this function with $J = 1$ and the values $b^J = b$ and $\gamma^J = 1$. Further, a general convex cost of procurement can be approximated by a piecewise linear function such as this by varying the number of segments in the cost function.

Notice that the optimal commitment policy for selling the output and the marginal value of a unit of output inventory is not affected by the procurement cost. Thus, Lemma 1 holds for this case and the marginal value of output is given by equation (7). Further, the value function $V_n(e_n, Q_n, \mathcal{I}_n)$ is separable in e_n and Q_n as shown.

$$\begin{aligned} V_n(e_n, Q_n, \mathcal{I}_n) &= \Delta_n Q_n + U_n(e_n, \mathcal{I}_n) \quad \text{and} \\ U_n(e_n, \mathcal{I}_n) &= \max_{\substack{0 \leq x_n \leq K^J, \\ 0 \leq m_n \leq \min\{C, e_n + x_n\}}} \Big\{ [\Delta_n - p] m_n - \mathscr{C}(S_n, x_n) \\ &\qquad - h_I [e_n + x_n - m_n] + \beta \mathbb{E}_{\mathcal{I}_n} [U_{n+1}(e_{n+1}, \mathcal{I}_{n+1})] \Big\} \end{aligned}$$

We now focus on computing the marginal value of input inventory when the procurement cost is given by equation (20). To this end, let D be the greatest common divisor (GCD) of $(C, K^1 - K^0, K^2 - K^1, \ldots, K^J - K^{J-1})$. Let $(a, b^1, b^2, \ldots, b^J)$ be positive integers such that $C = aD$ and $K^j = b^j D$ for all $j = 1, 2, \ldots, J$ and $b^0 = 0$. Using arguments similar to those in the proof of Theorem 1, we can prove the next result.

Theorem 2. *The value function $U_n(e_n, \mathcal{I}_n)$ is continuous, concave and piecewise linear in e_n with changes in slope at integral multiples of D, for each realization of $\mathcal{I}_n$ when the procurement cost is given by $\mathscr{C}(S_n, x_n)$, as defined in equation (20). Let Θ_n^k denote the marginal value of input inventory (i.e., slope of U_n) when $e_n \in [(k-1)D, kD)$ where k is an integer.*

For all n, let $\Theta_n^k \triangleq \infty$ for $k \leq 0$. In the last period, $\Theta_N^k = S_N$ for all $k \geq 1$. For any period $n < N$ and $k \geq 1$, the marginal value of input inventory $\Theta_n^k \triangleq \Theta_n^{(k,J)}$ where

$$\Theta_n^{(k,j)} = \begin{cases} \Omega_n^{(k)} & j = 0 \\ \max\left\{\Omega_n^{(k+b^j)}, \min\left\{\gamma^j S_n, \Theta_n^{(k+b^{j-1})}\right\}\right\} & \text{for } j = 1, 2, \ldots, J \end{cases} \tag{21}$$

and $\Omega_n^{(k)}$ is given by equation (13).

Similar to the linear procurement cost case, we can define thresholds based on $\Omega_n^{(k)}$ to characterize the optimal procurement and processing policy when the procurement cost is convex and piecewise linear. However, the procurement policy is more involved and characterized by $J+1$ thresholds. More specifically,

Proposition 2. *For all $n < N$, let $\Omega_n^{(k)}$ be as defined in equation (13). Then, in period n*

1. *The optimal procurement quantity is given by*

$$x_n^* = \begin{cases} K^{j-1} & \text{if } \gamma^{j-1} S_n \leq \Omega_n^{(k+b^{j-1})} < \gamma^j S_n \\ \hat{s}^j D - e_n & \text{if } \Omega_n^{(k+b^{j-1})} \geq \gamma^j S_n \geq \Omega_n^{(k+b^j)} \\ K^j & \text{if } \gamma^{j+1} S_n > \Omega_n^{(k+b^j)} > \gamma^j S_n \end{cases}$$

where $\hat{s}^j = \arg\max_{s \in \mathbb{Z}} \left\{\Omega_n^{(s)} > \gamma^j S_n\right\}$.

2. *The optimal quantity to process is given by*

$$m_n^* = \begin{cases} C & \text{if } \Omega_n^{(k)} < \Delta_n - p \\ \min\{(e_n + x_n^* - \hat{r}D)^+, C\} & \text{if } \Omega_n^{(k)} \geq \Delta_n - p \geq \Omega_n^{(k+b^J)} \\ 0 & \text{if } \Omega_n^{(k+b^J)} > \Delta_n - p \end{cases}$$

where $\hat{r} = \arg\max_{r \in \mathbb{Z}} \{\Omega_n^{(r)} > \Delta_n - p\}$.

The results in Theorem 2 have been derived assuming the γ^j are stationary. However, equation (21) can easily incorporate non-stationary values of γ^j, thus allowing us to model time varying procurement cost functions. More significantly, the γ^j values can also be stochastic, with the realized values of γ^j being used in equation (21). In such a case, the variable $\mathcal{I}_n$ would include

$(\gamma_n^1, \gamma_n^2, \ldots)$ as part of the state variable. Similarly, equation (21) can be modified to easily incorporate non-stationary and stochastic values of b^j; i.e., the procurement capacities in each segment of the piecewise linear cost function need not be the same across periods. Stochastic γ^j and b^j are useful to model multiple sources of procurement, with stochastic marginal cost of procurement at each source. In Section 3, we present a specific instance where these generalizations are useful in developing heuristics for the integrated problem over a star network.

We now discuss computational issues associated with calculating the optimal policy.

2.3. Computation of Optimal Policy

The analytical results derived in Sections 2.1 and 2.2 did not depend on the specific dynamics of the various commodity prices. However, computing the conditional expectations in the marginal value of output and input inventory expressions (equations (6)–(7), (12) and (21)) depends on the specific model used to describe the evolution of $\mathcal{I}_n$. For the purposes of developing the computational procedures, we assume the dynamics of the various commodity prices follow a Markov process; i.e., $\mathcal{I}_n = (S_n, F_n^1, F_n^2, \ldots, F_n^\ell, \ldots, F_n^L)$. For instance, single factor mean-reverting processes and multi-dimensional, driftless geometric Brownian motion processes which are typically used to model commodity spot price and forward curve dynamics would fall under this category.

A standard approach developed in the financial literature for pricing derivatives, especially *American* style options which require evaluation of conditional expectations, involves discretizing the price processes using binomial or trinomial lattices to generate possible states of price realizations with corresponding probabilities of transition in discrete time steps. The objective is to approximate the joint evolution of the continuous time processes over the time period of interest. The option can then be valued on the generated price lattice by using backward stochastic dynamic programming recursion, using the terminal value of the option on the final set of approximated prices (cf., Ho et al. (1995), Nelson and Ramaswamy (1990), Hahn and Dyer (2008) for some examples of discrete-time lattices).

We can use a similar approach to generate a discrete-time lattice for the various prices and calculate the marginal values (and hence the optimal policy) at each node in the tree. Notice that we do not need to discretize the state variable corresponding to the input inventory because of the piecewise linear nature of the value function. These discretization procedures are fairly efficient when modeling bivariate processes. We can readily use these different price discretization procedures when there is a single output forward contract or all output forward price changes are perfectly correlated. In these cases, we can compute the optimal marginal values at all nodes in the price lattice and thus compute the optimal policy efficiently.

While the discretization procedures are theoretically valid for modeling multivariate processes, they become computationally inefficient as the number of processes increases. Thus, in the more general case with multiple forward contracts and imperfectly correlated price changes, we need to resort to tractable approximations to compute the marginal values efficiently. We describe one such computationally tractable approximation next.

2.3.1. Heuristic for computing marginal values. As mentioned, the primary difficulty in using the binomial trees is the fact that modeling more than two processes jointly becomes computationally inefficient. To overcome this, we consider an approximation where only the dynamics of the input spot price and the nearest maturing forward contract are modeled in any period n. More precisely, define

$$\hat{\mathcal{I}}_n = (S_n, F_n^{\ell}, F_1^{\ell+1}, F_1^{\ell+2}, \ldots, F_1^{L}) \text{ for } n \text{ such that } N_{\ell-1} \leq n < N_{\ell} \tag{22}$$

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. (2009a).

Next, we approximate the marginal value of output inventory given in equations (6) and (7) by conditioning the expectations on $\hat{\mathcal{I}}_n$ instead of $\mathcal{I}_n$ as follows.

$$\hat{\Phi}^{(\ell)}_{N_\ell - 1} = \begin{cases} \beta F^L_{N_L-1} - h_O & \text{if } \ell = L \\ \max\left\{\beta F^\ell_{N_\ell-1} - h_O,\ \beta^{N_{\ell+1}-N_\ell} E_{\hat{\mathcal{I}}_{N_\ell-1}}\left[\hat{\Phi}^{(\ell+1)}_{N_{\ell+1}-1}\right] - h_O \sum_{t=0}^{N_{\ell+1}-N_\ell-1} \beta^t\right\} & \text{if } \ell < L \end{cases} \tag{23}$$

and

$$\hat{\Delta}_n = \begin{cases} \left[\beta^{(N_\ell-1)-n}\mathbb{E}_{\hat{\mathcal{I}}_n}\left[\hat{\Phi}^{(\ell)}_{N_\ell-1}\right] - h_O \sum_{t=0}^{(N_\ell-1)-n-1} \beta^t\right] & \text{if } N_{\ell-1} \le n < N_\ell \text{ for } \ell = 1, 2, \ldots, L \\ 0 & \text{if } n \ge N_L \end{cases} \tag{24}$$

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

$$\hat{\Theta}^k_n = \max\{\hat{\Omega}^{(k+b)}_n, \min\{S_n, \hat{\Omega}^{(k)}_n\}\} \tag{25}$$

where

$$\hat{\Omega}^{(k)}_n = \max\left\{\beta\mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Theta}^k_{n+1}] - h_I, \min\left\{\hat{\Delta}_n - p, \beta\mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Theta}^{k-a}_{n+1}] - h_I\right\}\right\} \tag{26}$$

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

The heuristic procurement, processing and commitment quantities $(\hat{x}_n, \hat{m}_n, \hat{q}_n)$ are then given by the results in Proposition 1 and Lemma 1 with the approximate marginal values replacing the true marginal values. We can also define the approximate marginal value of input inventory when the procurement cost is convex and piecewise linear in an analogous manner, using equation (21).

Notice that this heuristic requires only modeling the joint evolution of two price processes in any given period. Further, the heuristic is exact in the case where a single forward contract is available for selling the output commodity. Thus, the binomial discretization approaches mentioned earlier can be used to compute the approximate marginal values efficiently. We discuss the mechanics of implementing this heuristic in Section 2.4.1. We now develop a computationally tractable upper bound on the optimal expected profits, which will be used as a benchmark to evaluate the performance of the heuristic. We quantify the performance of the heuristic using numerical studies in Section 2.4.

2.3.2. 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. (2008). 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. An optimal policy with the information relaxation can potentially take advantage of the additional information available to improve the solution, leading to temporally infeasible policies for 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. (2008) define dual penalties for information relaxations. Akin to the strong duality result for linear programs, the solution to the relaxed problem is equal to the optimal solution of the original problem when an ideal dual penalty is used. Furthermore, for any appropriately defined feasible dual penalties, 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 the upper bound to the single node problem; 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

$$H_N^{UB}(e_N, Q_N; \Gamma_N) = S_N e_N \tag{27}$$

$$\begin{aligned} H_n^{UB}(e_n, Q_n; \Gamma_N) = \max_{q_n, x_n, m_n \in \mathcal{B}_n} \Bigg\{ \left[\beta^{N_\ell - n} F_n^\ell - h_O \sum_{t=0}^{n_\ell - n - 1} \beta^t \right] q_n - p m_n - S_n \times x_n - h_I e_{n+1} \\ - z_n(e_n, q_n, x_n, m_n, \Gamma_N) + \beta H_{n+1}^{UB}(e_{n+1}, Q_{n+1}; \Gamma_N) \Bigg\} \\ \text{for } n = 1, 2, \ldots, N-1 \end{aligned} \tag{28}$$

where e_{n+1} and Q_{n+1} are given by the state transition equations (4)–(5) and $\mathcal{B}_n$ is the constraint set on the decisions in period n. Specifically,

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

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

$$V_1^{UB}(e_1, Q_1, \mathcal{I}_1) = E_{\mathcal{I}_1}[H_1^{UB}(e_1, Q_1; \Gamma_N)] \tag{29}$$

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. Consider an ideal dual penalty in period n defined as

$$z_n^{ideal}(e_n, q_n, x_n, m_n, \Gamma_N) = \beta \left[V_{n+1}(e_{n+1}, Q_{n+1}, \mathcal{I}_{n+1}) - \mathbb{E}_{\mathcal{I}_n}\left[V_{n+1}(e_{n+1}, Q_{n+1}, \mathcal{I}_{n+1}) \right] \right]$$

Substituting this ideal penalty in the optimization in equation (28) will lead to $V_1^{UB} = V_1$, a tight bound. However, using the ideal penalty described above is not practical as it would require computing the exact value function. 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

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

where the marginal values, $\hat{\Delta}_{n+1}$ and $\hat{\Theta}_{n+1}^k$, are given by equations (24) and (25) and $\hat{\lambda}_{n+1}^k$ are constants such that $\hat{V}_{n+1}$ is continuous in e_{n+1} and $\hat{\lambda}_n^1 = 0$ for all n.

We then have

PROPOSITION 3. $V_1^{UB}(e_1, Q_1, \mathcal{I}_1)$ *as defined in equation (29), with dual penalties given by*

$$z_n(e_n, q_n, x_n, m_n, \Gamma_N) = \beta \left[\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1}) - \mathbb{E}_{\hat{\mathcal{I}}_n}\left[\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1}) \right] \right] \tag{30}$$

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

Proof: The dual penalty in equation (30) is a feasible penalty and hence, by Proposition 3.1 in Brown et al. (2008), $V_1^{UB}(e_1, Q_1, \mathcal{I}_1) \geq V_1(e_1, Q_1, \mathcal{I}_1)$. □

Notice that the DP given by (28) 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. The computation of the upper bound problem along each sample path is described in Appendix A.

2.4. Numerical Study

In this section, we describe several numerical studies to support our analysis. We discretize the various price processes and compute the optimal and approximate marginal values, as the case may be, on the resulting discrete time price lattice.

We first demonstrate the computational efficiency of the procedure by computing the optimal policy when a single forward is available for output sales. We compare the expected profits generated by using the optimal policy for different levels of discretization and evaluate the tradeoff between the improvement in expected profits versus additional computational time as the number of discretization steps increases. These results are presented in Section 2.4.2.

We also perform numerical studies for the general case with multiple forward contracts. For the general case, we use the heuristic described in described in Section 2.3.1. We measure the performance of the heuristic by comparing the expected profits using the heuristic with the upper bound on optimal expected profits. We study how the gap between the expected profit and the upper bound changes with various parameters. The performance of the heuristic is quantified in Section 2.4.3. We now describe the implementation of the heuristic.

2.4.1. Implementation The input spot and output forward prices are continuous and evolve continuously in time. In the commodity pricing literature, one factor mean-reverting processes have often been used to model the spot price process for various commodities, including agricultural commodities (cf. Geman (2005), Chapter 3). While multi-factor models have also been used (see for instance Gibson and Schwartz (1990), Schwartz and Smith (2000), Geman and Nguyen (2005)), the single factor mean-reverting dynamics capture many of the essential features of commodity spot prices and are also analytically tractable. We use a single factor mean reverting model as in Schwartz (1997) to describe the evolution of the input spot price over the time interval $[0,T]$. More specifically, the dynamics of the input spot price S_t are modeled as $\ln S(t) = \chi(t) + \mu(t)$ where $\chi(t)$ represents the short-term deviation in prices and $\mu(t)$ the equilibrium price level. The short-term deviation $\chi(t)$ follows a mean-reverting process given by $d\chi(t) = -\kappa\chi(t)dt + \sigma_s dW_s(t)$, where $dW_s(t)$ is the increment of a standard Brownian motion, κ is the mean-reversion coefficient and σ_s the volatility.

Multi-dimensional driftless geometric Brownian motion processes have been used commonly to model the dynamics of commodity forward curves (cf., Geman (2005), Lai et al. (2009a)). We model the risk-neutral dynamics of the output forward price with maturity at time T_ℓ by a driftless geometric Brownian motion, 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 $\rho_{\ell k} \in [-1,1]$. Also, 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. The parameters for the input spot and output forward price processes used in the numerical studies are given in Appendix B.

For computing the heuristic policy, we use a discretization of the dynamics of the input and output spot prices, with $n = 1$ corresponding to time $t = 0$ and $n = N_\ell$ to $t = T_\ell$, for each ℓ. For each $\ell \in \{1, 2, \ldots, L\}$, we construct a 3-dimensional binomial tree as described in Hahn and Dyer (2008) with δ discretization steps between each period n and $n+1$ to represent the joint evolution of $(S(t), F(t,T_\ell))$, conditional on F_0^k for $k > \ell$. From each of these trees, we obtain a probability mass function $G_n^\ell(S_{n+1}, F_{n+1}^\ell | S_n, F_n^\ell)$ for each $n \leq N_\ell - 1$, for each node in the tree at time n. The probability mass function $G_n^\ell(\cdot)$ is used to compute expectations, conditional on $\hat{\mathcal{I}}_n$.

We also generate a 3-dimensional binomial tree to represent the evolution of $(F(t,T_\ell), F(t,T_{\ell+1}))$ for each $\ell \in \{1, 2, \ldots, L-1\}$. From each tree, we obtain a probability mass function

$\hat{H}^{\ell}_{N_\ell-1}(F^{\ell+1}_{N_\ell-1}|F^{\ell}_{N_\ell-1})$ 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 $\hat{H}^{\ell}_n(\cdot)$, along with $G^{\ell}(\cdot)$ is used to compute expectations, conditional on $\hat{\mathcal{I}}_n$, at the expiration of forward contract ℓ. Specifically,

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

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 $\hat{\Theta}^k_n$ 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^{ℓ}_n for n such that $N_{\ell-1}\leq n<N_\ell$, and the probability mass functions G^{ℓ}_n for $N_{\ell-1}\leq n<N_\ell-1$ and $\hat{G}^{\ell}_{N_\ell-1}$ for $n=N_\ell-1$.

We evaluate the policy using Monte Carlo simulation. We generate sample paths of prices for periods $n=1,2,\ldots,N$ by sampling from the true continuous time and space price processes. At period n such that $N_{\ell-1}\leq n<N_\ell$, we round the realized prices of S_n and F^{ℓ}_n to the closest discretized values in the 3-dimensional binomial tree. The heuristic policy parameters $(\hat{x}_n,\hat{m}_n,\hat{q}_n)$ are computed using the values of $\hat{\Delta}_{n+1}$ and $\hat{\Theta}^k_{n+1}$ stored at each node and the probability mass functions. Expected profits from using the heuristic policy are calculated as the average profit over 1000 sample paths.

For each sample path, we determine the dual penalty for each period n by using the approximate marginal values at the node in the 3-dimensional binomial tree which is closest to the realized prices S_n and F^{ℓ}_n. The optimization problem given by equation (28) is then solved for each sample path as a mixed-integer linear program. Given the nature of the problem, it is not guaranteed that an optimal solution to the upper bound problem can be found in reasonable time for each sample path. To ensure the upper bound computations terminate, we impose a limit on the computation times for the upper bound calculation along each sample path. For sample paths in which an optimal solution to the mixed-integer linear program is not found within this time limit, we use the optimal value of the linear programming relaxation of the mixed-integer problem as the upper bound value. The average over all sample paths is computed to obtain the upper bound on the optimal expected profits. Notice that the upper bound value thus computed overstates the true upper bound value obtained from information relaxation. The gap between the upper bound, UB, and the expected profits from the heuristic policy is computed and expressed as a percentage of UB.

For all the numerical studies presented here, the marginal processing cost was set to $p = 5$. Observe that a processing cost of $p = 5$ corresponds to an expected processing margin of 0 that allows us to model situations when the actual realized processing margin may be positive or negative. Also, we assume no discounting, i.e., $\beta = 1$, and the holding costs to be negligible. The procurement capacity was set to $K = 5$ and the processing capacity to $C = 3$. Thus, the processing capacity is 60% of the procurement capacity, reflective of the ITC case. Finally, the procurement cost is linear for all the numerical studies presented here.

2.4.2. Computing the Optimal Policy – Single Forward. Recall that the heuristic proposed in Section 2.3.1 is optimal for the case when only a single forward contract is available for output sales. We investigate the impact of discretization on the optimal expected profits by calculating the expected profits for $\delta \in \{5, 10, 15\}$. Table 1 gives the simulation results when a single forward contract is available for output sales and the number of periods in the horizon is varied from $N = 5$ to $N = 20$, for different values of δ. The values in parentheses are the average CPU time in seconds required to compute the optimal policy and the corresponding increase in the CPU time required as δ increases. We do not have results for $\delta = 15$ and $N = 20$ as the maximum memory available to the program was not enough to solve the problem (we used MATLAB®). The standard errors for the expected profits above range from 1.2% to 2.4% of the expected profits.

As expected, the optimal expected profits increase with the number of discretization steps for each horizon length (The value of -0.19% in the table is statistically insignificant). However, the percentage increase is only marginal, especially when compared to the increase in computational burden as evidenced by the CPU times required to compute the policy. Using a small number of

Table 1 Optimal expected profits and CPU Time (Single forward)

N	Optimal Value (Avg. CPU Time in sec.)			% increase in Value (% increase in CPU Time)	
	$\delta = 5$	$\delta = 10$	$\delta = 15$	$\delta : 5 \rightarrow 10$	$\delta : 10 \rightarrow 15$
5	9.92 (0.55)	10.28 (2.12)	10.26 (7.18)	3.06% (287.33%)	-0.19% (237.70%)
10	47.95 (18.04)	48.78 (52.77)	49.17 (175.77)	1.72% (192.60%)	0.80% (233.07%)
15	114.43 (36.32)	114.56 (294.26)	114.72 (1024.77)	0.11% (710.11%)	0.14% (248.25%)
20	151.54 (140.56)	152.89 (980.68)	- -	0.89% (597.67%)	- -

Table 2 Expected profits using heuristic policy (Multiple forwards)

# of Forwards L (Maturity Dates $\{N_\ell\}$)	Horizon Length N	Heuristic Value	UB Value	Gap (Std. Error)	Avg. CPU Time (sec.)
2 $\{5, 10\}$	10	45.46	49.78	8.68% (4.22%)	57.54
3 $\{5, 10, 15\}$	15	94.73	104.54	9.38% (3.30%)	388.16
4 $\{5, 10, 15, 20\}$	20	141.59	165.55	14.48% (3.01%)	1608.33

discretization steps, especially for longer horizon problems, is therefore computationally efficient with practically no loss in optimality. We fix the number of discretization steps δ to 10 for all the remaining studies.

2.4.3. Performance of Heuristic Policy – Multiple Forwards. In this section, we quantify the performance of the heuristic when multiple forward contracts are available for the output. When more than one forward contract is available for selling the output, the heuristic is no longer optimal. We investigate the performance of the heuristic as the number of forward contracts available for the output commodity varies from $L = 2$ to $L = 4$. The results of the simulation are given in Table 2, with the gap and standard error being expressed as a % of the upper bound value. As seen from the results, the heuristics perform quite well, even when there are multiple forward contracts. Also, recall that the upper bound value reported here is overstated and the true gap between the upper bound and expected profits is potentially less than the gaps reported here. Finally, the time taken to compute the heuristics are not significantly different from the single forward case, as seen by the CPU times.

A key parameter affecting the accuracy of the heuristic is the correlation between the Brownian motion increments of the different forward price processes. For instance, if the Brownian motion increments were all perfectly correlated, the heuristic would actually be optimal. While we do not report the results here, we performed numerical experiments varying the correlation coefficient between the various price processes and found that the results did not vary significantly.

3. The Star Network Problem

In reality, in the *e-Choupal* and other commodity business contexts, especially agricultural commodities, procurement is usually done over multiple locations. We now extend our analysis to consider the integrated problem of procurement, processing and trade over a star network with multiple locations. In addition to providing analytical tractability, a star network configuration also approximates real world commodity processing networks fairly well. In a star network, a procurement source for the input commodity usually serves at most one processing location, while a processing plant may have the input transshipped from multiple locations. This is definitely the case with the *e-Choupal* network, where a set of procurement hubs are associated with a processing plant. Due to the geographic proximity and availability of information, differences in prices across the various procurement hubs are usually not significant enough to justify transshipment of the input between the non-processing locations. While the network problem is more complex than the single node problem, some of the insights from the single node problem extend to the network case. The results for the single node problem with convex cost of procurement are especially useful and instrumental in developing a tractable heuristic to solve the network problem.

We consider a multi-node network of M procurement nodes each with procurement capacity of K^i units per period at location $i \in i = 1, 2, \ldots, M$. Let S_n^i denote the price for the input in the spot market at location i. We consider a star network configuration, with location 1 being the central node with a processing capacity of C units, while all other nodes only have a procurement capacity. The transshipment cost is $t^{(j)}$ per unit between locations 1 and j for $j = 2, 3, \ldots, M$. Since the only source of (direct) revenue at the non-processing locations is through trade of the input commodity, the firm has an incentive to transship input from one non-processing location to another only when there is an arbitrage opportunity on the input commodity between the locations; i.e., if the difference in expected trade prices is more than the transshipment cost between the locations. These arbitrage opportunities are not relevant to the core operations considered in our model and therefore to eliminate such opportunities, we do not allow direct transshipment between the non-processing locations[3]; i.e., $t^{(ij)} = \infty$ for $(i,j) \in \{2,3,\ldots,M\} \times \{2,3,\ldots,M\}$.

Let $\mathbf{e_n} = (e_1, e_2, \ldots, e_M)$ be the vector of input inventories at the M locations. Since there is only a single processing location, the output inventory is still a scalar value Q_n. The firm's decisions include a) the quantity of input to procure at each location: $\mathbf{x_n} = (x_n^1, x_n^2, \ldots, x_n^M)$, b) the quantity of the input commodity to be transshipped between the processing and other procurement locations: $\mathbf{y_n} = (y_n^{(ij)} : i \neq j, i = 1 \text{ or } j = 1)$ where $y_n^{(ij)}$ is the quantity transshipped from location i to location j, c) the quantity of the output commodity to be committed for sale: q_n, and d) the quantity of input to be processed into output in period n: m_n.

Notice that the network structure does not affect the optimal commitment policy for selling the output and the marginal value of a unit of output inventory. Thus, Lemma 1 holds for this case and the marginal value of output is given by equation (7). Further, the value function $V_n(\mathbf{e_n}, Q_n, \mathcal{I}_n)$ is separable in $\mathbf{e_n}$ and Q_n as given by equation (8) and we have

$$U_n(\mathbf{e_n}, \mathcal{I}_n) = \max_{(\mathbf{x_n}, \mathbf{y_n}, m_n) \in \mathcal{B}_n} \Bigg\{ [\Delta_n - p] m_n - \sum_{i=1}^{M} S_n^i x_n^i - \sum_{i=2}^{M} t^{(i)} [y_n^{(1i)} + y_n^{(i1)}] - h_I \left[\sum_{i=1}^{M} (e_n^i + x_n^i) - m_n \right] + \beta \mathbb{E}_{\mathcal{I}_n} [U_{n+1}(\mathbf{e_{n+1}}, \mathcal{I}_{n+1})] \Bigg\} \text{ for } n < N \tag{31}$$

$$U_N(\mathbf{e_N}, \mathcal{I}_N) = \sum_{i=1}^{M} S_N^i e_N^i \tag{32}$$

where the set of feasible actions in period n, $\mathcal{B}_n$ is given by

$$\mathcal{B}_n = \left\{ \begin{array}{ll} 0 \leq x_n^i \leq K^i & \text{for } i = 1, 2, \ldots, M \\ 0 \leq m_n \leq C & \\ m_n + \sum_{i=2}^M y^{(1i)} \leq e_n^1 + x_n^1 + \sum_{j=2}^M y_n^{(j1)} & \\ y^{(i1)} \leq e_n^i + x_n^i & \text{for } i = 2, 3, \ldots, M \\ \mathbf{x_n} \geq 0, \mathbf{y_n} \geq 0, m_n \geq 0 & \end{array} \right\} \tag{33}$$

and the state transition equations are given by

$$e_{n+1}^i = \begin{cases} e_n^i + x_n^i + \displaystyle\sum_{j=2}^M y_n^{(j1)} - \sum_{j=2}^M y_n^{(1j)} - m_n & \text{for } i = 1 \\ e_n^i + x_n^i + y_n^{(1i)} - y_n^{(i1)} & \text{for } i = 2, \ldots, M \end{cases} \tag{34}$$

Notice that (32) is linear in $\mathbf{e_N}$ and thereby, also piecewise linear. Similar to the single node problem, we can use induction arguments to show that $U_n(\mathbf{e_n}, \mathcal{I}_n)$ is piecewise linear and concave in $\mathbf{e_n}$. While it is theoretically possible, it is hard to derive expressions for the marginal value of inventory at location i as it depends not just on e_n^i, but the entire inventory vector $\mathbf{e_n}$. As such, it is hard to obtain further insights into the network problem without additional simplifications. In the next section, we consider one such simplified network and use the insights to develop a heuristic for the network problem.

3.1. Heuristic for the Star Network

The heuristic described in Section 2.3.1 overcame the high dimensionality introduced by the output price processes by only considering the joint evolution of the input spot and nearest maturing output forward price. In case of the network problem, we have to consider multiple input spot prices, since the input price changes across locations are usually imperfectly correlated. We could potentially use the same information approximation techniques developed for the single node problem to account for these additional price processes. However, solving the network problem optimally is further complicated by the fact that the marginal value of input inventory is generally different across the various locations and dependent on the inventory levels at the different locations and not just the aggregate input inventory.

Under some simplifying assumptions, the network problem is tractable and is equivalent to the single node problem with piecewise linear, convex cost of procurement. To see this, consider a situation where all the procurement nodes are close to the central processing location such that the transshipment costs between the nodes are a very small fraction of the commodity prices. However, the input commodity prices realized in the spot markets can still be different across locations. Further, consider the case when the trade price at the end of the horizon for the input is the same, irrespective of which node the input is physically stored at. Thus we have, $t^{(i)} \simeq 0$ and $S_N^i \simeq S_N$ for all i. Thus, we can write the SDP equations (31)–(32) as

$$U_n(\mathbf{e_n}, \mathcal{I}_n) = \max_{(\mathbf{x_n}, \mathbf{y_n}, m_n) \in \mathcal{B}_n} \left\{ [\Delta_n - p] m_n - \sum_{i=1}^M S_n^i x_n^i - h_I \left[\sum_{i=1}^M (e_n^i + x_n^i) - m_n \right] + \beta \mathbb{E}_{\mathcal{I}_n} [U_{n+1}(\mathbf{e_{n+1}}, \mathcal{I}_{n+1})] \right\}$$

$$U_N(\mathbf{e_N}, \mathcal{I}_N) = S_N \sum_{i=1}^M e_N^i$$

with the same state transition equations as before.

Notice that the input inventory across different locations are indistinguishable in their marginal values in this case. Thus, we can replace $\mathbf{e_n}$ by $\hat{e}_n = \sum_i e_n^i$ and drop the transshipment decisions from the optimization problem to write

$$U_n(\hat{e}_n, \mathcal{I}_n) = \max_{(\mathbf{x_n}, m_n) \in \hat{\mathcal{B}}_n} \left\{ [\Delta_n - p] m_n - \sum_{i=1}^M S_n^i x_n^i - h_I \left[(\hat{e}_n + \sum_{i=1}^M x_n^i) - m_n \right] + \beta \mathbb{E}_{\mathcal{I}_n} [U_{n+1}(\hat{e}_{n+1}, \mathcal{I}_{n+1})] \right\} \text{ for } n < N$$

$$U_N(\hat{e}_N, \mathcal{I}_N) = S_N \hat{e}_N$$

where $\hat{\mathcal{B}}_n$ is the set of constraints on the procurement and processing quantities given by

$$\hat{\mathcal{B}}_n = \left\{ \begin{array}{ll} 0 \le x_n^i \le K^i & \text{for } i = 1, 2, \dots, M \\ 0 \le m_n \le C & \\ m_n \le \hat{e}_n + \sum_{i=1}^M x_n^i & \end{array} \right\}$$

Notice that even though the input inventory across various locations are indistinguishable, the marginal cost of procurement, S_n^i, is still different across locations and is retained in the above optimization. The SDP equations above are the same as those for the single node, convex procurement cost case, albeit with stochastic γ^j because the S_n^i are stochastic. We can therefore use the results from Section 2.2 to solve this simplified network problem. The heuristic for the general star network is based on the equivalence between the simplified network and the single node problem and has two stages. The first stage approximates the star network as a single node with a piecewise linear convex cost of procurement. In the second stage, we use the approximations developed for the single node problem to model the joint evolution of the various price processes.

Let $S_n^{(j)}$ be the j^{th} order statistic of $\mathbf{S_n} = (S_n^1, S_n^2, \ldots, S_n^M)$. Let i_j be the index of the location corresponding to the j^{th} order statistic of $\mathbf{S_n}$. Define

$$\gamma_n^j = \mathbb{E}_{\mathcal{I}_1}\left[\frac{S_n^{(j)}}{S_n^1}\middle| S_1^1\right] \tag{35}$$

$$\bar{K}_n^j = \mathbb{E}_{\mathcal{I}_1}\left[\sum_{k=1}^{j} K^{i_k}\middle| S_1^1\right] \tag{36}$$

for $j = 1, 2, \ldots, M$, for all n.

Let D be the greatest common divisor of $(C, \bar{K}^1, \bar{K}^2 - \bar{K}^1, \ldots, \bar{K}^M - \bar{K}^{M-1})$ where $\bar{K}^j$ is the average $\bar{K}_n^j$ over all n. Define $(a, b^1, \ldots, b^M)$ to be positive integers such that $C = aD$ and $\bar{K}^j = b^j D$. We approximate the star network by an equivalent single node with a procurement cost function given by equation (20), where $S_n = S_n^1$ and the γ_n^j and K^j are given as above. For this single node network, we can calculate the approximate marginal value of input inventory $\hat{\Theta}_n^k$, according to equation (21) and using the approximation scheme described in Section 2.3.1.

To compute the heuristic procurement, transshipment and processing quantities for the general network problem, we define the approximate value function as

$$\hat{U}_n(\mathbf{e_n}, \hat{\mathcal{I}}_n) = \hat{\Theta}_n^k \sum_i e_n^i + \hat{\lambda}_n^k \text{ if } (k-1)D \le \textstyle\sum_i e_n < kD \tag{37}$$

where the $\hat{\lambda}_n^k$ are constants such that $\hat{U}_n$ is continuous in $\sum_i e_n^i$ and $\hat{\lambda}_n^1 = 0$ for all n and all $\mathcal{I}_n$. The heuristic policy for the general network is then given as the solution to the following optimization problem

$$\max_{\mathbf{x_n},\, \mathbf{y_n},\, m_n \in \mathcal{B}_n} \Bigg\{ [\hat{\Delta}_n - p] m_n - \sum_{i=1}^{M} S_n^i \times x_n^i - \sum_{i=2}^{M} t^{(i)} [y_n^{(1i)} + y_n^{(i1)}] - h_I \sum_{i=1}^{M} e_{n+1}^i + \beta \mathbb{E}_{\mathcal{I}_n}\left[\hat{U}_{n+1}(\mathbf{e_{n+1}}, \hat{\mathcal{I}}_{n+1})\right] \Bigg\} \; n < N$$

We can use dual penalties based on these heuristics and compute an upper bound on the optimal expected profits for the network case using the procedure described in Section 2.3.2. This upper bound can then be used to evaluate the performance of the heuristics for the network case.

3.2. Numerical Study

We investigate the performance of the heuristic for a two-node and a five-node network respectively. As in the single node case, the output forward prices are modeled as driftless geometric Brownian motions. The input spot price at each location, $S^i(t)$, have the same dynamics as in the single node case. The correlations between the various prices are given in Appendix B.

For the two-node network, the processing capacity was set to $C = 3$ and the procurement capacities at the two locations was set to $K^1 = 3$ and $K^2 = 2$ respectively. For the five-node network, the processing capacity was set to $C = 6$ and the procurement capacity was set to $K^i = 2$ for all i. The transshipment cost in both networks was set to 0.5 per unit.

The performance of the network heuristic is summarized in Table 3. The true cost of procurement over the network in any period is convex and piecewise linear with stochastic coefficients. As the input price changes across locations become more correlated, the variability in the coefficients of the convex, piecewise linear cost function decreases. Further, the salvage value across all locations will also become closer as the correlation increases. Recall that the heuristic described in Section 3.1 approximates the multi-node network as an equivalent single node network with convex and piecewise linear costs, where the coefficients γ_n^j are deterministic and given by equation (35). Thus, as the correlation of input price changes across locations increases, the heuristic approximates the network better and the performance of the network heuristic will be close to the performance of the single node heuristic. In the limit with perfect correlation and zero transshipment cost, the heuristic will be optimal when there is a single forward available for output sale commitments. In fact, when the input price change correlations across all locations is varied from 0.95 to 0.98, we found that the gap decreases to 9.02% (from 21.31%) for the five-node, 5 period problem. The results in Table 4, for a two node network with three forward contracts and horizon length of 15

Table 3 Performance of heuristic for network problem

# of Forwards L (Maturity Dates $\{N_\ell\}$)	Horizon Length N	Two-Node Network			Five-Node Network		
		Heuristic Value	UB Value	Gap (Std. Error)	Heuristic Value	UB Value	Gap (Std. Error)
1 $\{5\}$	5	10.53	11.30	6.76% (7.89%)	17.37	23.22	25.19% (7.74%)
2 $\{5, 10\}$	10	42.67	47.83	10.80% (4.20%)	83.22	95.46	12.82% (4.31%)
3 $\{5, 10, 15\}$	15	85.67	96.30	11.04% (3.54%)	163.53	191.61	14.65% (3.58%)
4 $\{5, 10, 15, 20\}$	20	148.49	167.77	11.49% (2.95%)	271.75	317.32	14.36% (3.18%)

Table 4 Sensitivity of heuristic to transshipment costs

Transshipment Cost	Heuristic Value	UB Value	Gap (Std. Error)
0	91.90	105.99	13.29% (3.74%)
0.5	91.35	100.18	8.82% (3.95%)
1	86.53	98.45	12.11% (3.90%)
2	91.76	98.66	7.00 % (3.99%)
3	89.54	100.31	10.73 % (3.80%)

periods, also indicate that the network heuristic is fairly robust to changes in transshipment costs and performs quite well even for high transshipment costs. Thus, the approximation of the convex procurement cost function appears to have a much bigger impact on the performance than the transshipment costs. Future improvements to the network heuristic could look at better ways to approximate the input procurement cost and end of horizon salvage functions.

The policy computation times for the network heuristic are comparable to the those in the single node case. Evaluating the upper bound in the network case however takes significantly longer. This is not a concern for implementing the heuristic itself, because we need to compute the upper bound only for a benchmark and not implementing the policy itself.

4. Conclusion

In this paper we have considered the integrated procurement, processing and trade decisions for a firm dealing in commodities and subject to procurement and processing capacity constraints. We solved the problem optimally and showed that the procurement and processing decisions in any period are governed by inventory dependent thresholds and develop recursive expressions to compute these thresholds. We also extended our results to incorporate convex, piecewise linear costs of procurement and star networks; i.e., networks with a central processing node and multiple procurement nodes connected to the processing node. We developed an efficient heuristic to solve the single node problem when multiple forward contracts with different maturities are available to sell the output. Additionally, for the network, we developed a heuristic based on approximating the network as a single node, with piecewise linear convex cost of procurement. The numerical studies indicate that the heuristic policies are near optimal, with the gap between the expected profits and an upper bound on the optimal profits between $6\% - 14\%$ for most cases.

Our work lays the foundation for further research in commodity trading networks. The focus of this paper has been on a single input that can be processed into a single output. In reality, multiple output commodities may be produced upon processing the input; e.g., soybean is crushed to produce soybean meal and oil, both of which are commodities that can be traded. The results developed here extend to the case when multiple output commodities are produced upon processing. We illustrate the case for the single node problem when two products are produced upon processing the input, but the extension to more products and the network case is straightforward.

Multiple output products. Let one unit of input when processed yield α_M units of product M and $(1-\alpha_M)$ units of product O, with $0 < \alpha_M < 1$ (one could think of M and O to denote meal and oil in ITC's soybean commodity network). Let ℓ_m and ℓ_o index the forward contracts available for output M and O respectively with maturity at N_{ℓ_m} and N_{ℓ_o}. Let $M_n^{\ell_m}$ and $O_n^{\ell_O}$ be the forward prices on these contracts. Further, let h_M and h_O be the unit holding cost per period for M and O.

After processing, the decision to commit commodity M or O for sale against a forward contract can be made independent of the decision for the other commodity. Thus, similar to the single output case, each unit of M and O is a compound exchange option on the remaining forward contracts for that commodity and Lemma 1 holds for each commodity. Also, we can calculate the marginal value of inventory for each output in a manner similar to the single output case. Define $\Phi^{(\ell_m)}_{N_{\ell_m}-1}$ and $\Delta^{(m)}_n$, the marginal value of inventory for commodity M using the forward contracts and holding cost for M. Define $\Phi^{(\ell_o)}_{N_{\ell_o}-1}$ and $\Delta^{(o)}_n$ similarly for O. The benefit from processing in period n is equal to $\Delta_n - p$ where $\Delta_n = \alpha_M \Delta^{(m)}_n + (1-\alpha_M)\Delta^{(o)}_n$ in this case. The marginal value of input inventory and the optimal procurement and processing decisions are then given by Theorem 1, with Δ_n as defined.

An obvious extension would be to consider a multi-product setting where the firm can choose from different input commodities to procure and / or process the input into different output commodities. The methodology and analysis used in this paper can be useful in analyzing more general commodity processing and trading networks which include other operational constraints such as transshipment capacities and stochastic transshipment costs among others. Extending the research to incorporate these aspects could result in newer insights.

The focus of our paper is to determine integrated optimal policies for a risk-neutral firm. Firms dealing in commodity markets are usually risk-averse, with limited appetite for taking on risk. There is a substantial body of literature in the finance and economics streams concerning risk management in commodity and financial markets, and the use of market instruments for managing risk. The operations management literature concerning risk-aversion is fairly limited and focused mostly on single period models; see for example, Eeckhoudt et al. (1995), Agrawal and Seshadri (2000), Gaur and Seshadri (2005). Chen et al. (2007) is an exception, and considers a multiperiod inventory problem for a risk-averse decision maker. The problem context for this paper provides an opportunity to contribute to the literature on risk-aversion in a multiperiod problem. Thus, an important extension to our present work would be to incorporate a firm's risk aversion in the decision making explicitly and integrate the financial and operational decisions in a multiperiod network setting.

Appendix A: Upper Bound Calculation

The upper bound computation along a sample path Γ_N is given by

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

where the dual penalty is given by

$$\begin{aligned}
z_n(e_n, q_n, x_n, m_n, \Gamma_N) &= \beta \left[\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1}) - \mathbb{E}_{\hat{\mathcal{I}}_n} \left[\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1}) \right] \right] \\
&= \beta \left[\Theta_{n+1}^k(\hat{\mathcal{I}}_{n+1}) - E_{\hat{\mathcal{I}}_n}[\Theta_{n+1}^k(\hat{\mathcal{I}}_{n+1})] \right] e_{n+1} + \beta[\hat{\lambda}_{n+1}^k - \mathbb{E}_{\hat{\mathcal{I}}_n}[\lambda_{n+1}^k]] \\
&\quad + \beta \left[\hat{\Delta}_{n+1} - \mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Delta}_{n+1}] \right] Q_{n+1} \text{ for } e_{n+1} \in [(k-1)D, kD)
\end{aligned}$$

Notice that the penalty function above is piecewise linear in e_{n+1}, with change in slopes at integral multiples of D. Since the procurement and processing capacities are integral multiples of D, we can solve the upper bound computation as a mixed-integer linear program, where the binary integer variables identify the segment that e_{n+1} lies in, for each n.

Specifically, $(N-(n+1))a+1$ is the maximum number of segments with different slopes in the penalty function. Further, $e_{n+1} \in [0, nbD]$ always. Therefore, in period n we need $\min\{nb, (N-(n+1))a+1\}$ binary variables to indicate which segment the ending input inventory lies in, in order to compute the dual penalty value at the corresponding inventory level. Let

$$\begin{aligned}
\kappa(n) &= \min\{nb, (N-(n+1))a+1\} \\
a_n^k &= kD \text{ for } k = 0, 1, \ldots, \kappa(n)-1 \text{ and } a_n^{\kappa(n)} = nbD
\end{aligned}$$

Following Sherali (2001), let $\varphi_n^{(k,l)}$ and $\varphi_n^{(k,r)}$ be continuous variables and y_n^k a binary variable for each $k = 1, 2, \ldots, \kappa(n)$ and $n = 1, 2, \ldots, N-1$. Also define

$$z_n^k = \left[\Theta_{n+1}^k(\hat{\mathcal{I}}_{n+1}) - \mathbb{E}_{\hat{\mathcal{I}}_n}[\Theta_{n+1}^k(\hat{\mathcal{I}}_{n+1})]\right] kD + [\hat{\lambda}_{n+1}^k - \mathbb{E}_{\hat{\mathcal{I}}_n}[\lambda_{n+1}^k]]$$

$$\text{for } k = 1, 2, \ldots, \kappa(n), \text{ for } n = 1, 2, \ldots, N-1 \text{ and}$$

$$z_n^0 = 0 \text{ for } n = 1, 2, \ldots, N-1$$

We can then write the upper bound maximization problem as follows

$$\max \sum_{n=1}^{N-1} \left(\left[\beta^{N_\ell - n} F_n^\ell - h_O \sum_{t=0}^{n_\ell - n - 1} \beta^t \right] q_n - p m_n - S_n x_n - h_I e_{n+1} - \beta \left[\hat{\Delta}_{n+1} - \mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Delta}_{n+1}] \right] Q_{n+1} - \beta \sum_{k=1}^{\kappa(n)} [z_n^{k-1} \varphi_n^{(k,l)} + z_n^k \varphi_n^{(k,r)}] \right) + S_N e_N$$

subject to

$$\begin{aligned}
x_n &\le K && n = 1, 2, \ldots, N-1 \\
m_n &\le C && n = 1, 2, \ldots, N-1 \\
q_n &= 0 && n \neq N_\ell - 1 \text{ for } \ell \in \{1, 2, \ldots, L\} \\
q_n &\le Q_n + m_n && n = N_\ell - 1 \text{ for } \ell \in \{1, 2, \ldots, L\} \\
Q_{n+1} &= Q_n + m_n - q_n && n = 1, 2, \ldots, N-1 \\
e_{n+1} &= e_n + x_n - m_n && \text{for } n = 1, 2, \ldots, N-1 \\
e_{n+1} &= \sum_{k=1}^{\kappa(n)} [a_n^{k-1} \varphi_n^{(k,l)} + a_n^k \varphi_n^{(k,r)}] && \text{for } n = 1, 2, \ldots, N-1 \\
\varphi_n^{(k,l)} + \varphi_n^{(k,r)} &= y_n^k && \text{for } k = 1, 2, \ldots, \kappa(n), \\
& && n = 1, 2, \ldots, N-1 \\
\sum_{1}^{\kappa(n)} y_n^k &= 1 && \text{for } n = 1, 2, \ldots, N-1 \\
y_n^k &\in \{0, 1\} && k = 0, 1, \ldots, \kappa(n), \\
& && n = 1, 2, \ldots, N-1 \\
x_n, m_n, q_n, e_{n+1}, \varphi_n^{(k,l)}, \varphi_n^{(k,r)} &\ge 0 && n = 1, 2, \ldots, N-1
\end{aligned}$$

The above problem can then be solved using a standard mixed-integer programming solver (we used the CPLEX® solver in our implementation).

Appendix B: Price Process Parameters for Numerical Studies

This section describes the parameters used for generating the results for the numerical studies described in Sections 2.4 and 3.2. Specifically,

1. Table 5 gives the parameters underlying the dynamics of the input spot and output forward price processes.

2. Table 6 gives the correlation matrix for the output commodity forward prices. These values have been adapted from the correlation structure for natural gas forward contracts described in Lai et al. (2009b).

3. Table 7 gives the correlation matrix for the output commodity forward prices and the input commodity spot prices at each location.

4. Table 8 gives the correlation matrix for the input spot prices across locations.

Table 5 Input and Output Price Process Parameters

(a) Input Price Process

Parameter	Value
κ	0.332
σ	0.490
μ	3.218

(b) Output Price Process

Maturity N_ℓ	Volatility σ_ℓ
5	0.42
10	0.35
15	0.35
20	0.35

Table 6 Output Forward Price Correlation Matrix

	Maturity			
Maturity	5	10	15	20
5	1	0.958	0.933	0.91
10		1	0.983	0.959
15			1	0.982
20				1

Table 7 Output Forward Price and Input Spot Price Correlation Matrix

	Location				
Maturity	1	2	3	4	5
5	0.91	0.91	0.91	0.91	0.91
10	0.91	0.91	0.91	0.91	0.91
15	0.91	0.91	0.91	0.91	0.91
20	0.91	0.91	0.91	0.91	0.91

Table 8 Input Spot Price Correlation Matrix

	Location				
Location	1	2	3	4	5
1	1	0.983	0.959	0.935	0.919
2		1	0.982	0.962	0.946
3			1	0.99	0.975
4				1	0.991
5					1

Endnotes

1. Technically, a greatest common divisor may not exist if either C or K is not a rational number. We assume that both C and K are rational.

2. We should note that while this is true in general for ITC, there are instances when the firm procures from the direct channel at a higher price, even if the price in the spot market is lower. Because the firm has better control over the quality of the soybean procured in the direct channel, however, the true marginal cost after adjusting for quality is still lower in the direct channel. Thus, the total procurement cost is still convex.

3. This restriction on possible transshipment is also consistent with the actual features of the ITC network, where a processing plant is supported by a set of procurement hubs, but transshipment of soybean between the procurement hubs is very rarely observed.

Acknowledgments

The authors gratefully acknowledge the collaboration and support of The ITC Group. Their generous hospitality and access to key personnel and data greatly enhanced this research. The authors thank the associate editor and the referees whose suggestions have led to an improved version of this paper.

References

Agrawal, V., S. Seshadri. 2000. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing & Service Operations Management* **2**(4) 410–423.

Anupindi, R., S. Sivakumar. 2006. Supply chain re-engineering in agri-business: A case study of ITC's e-choupal. Hau L. Lee, Chung-Yee Lee, eds., *Supply Chain Issues in Emerging Economies*. Elsevier-Springer.

Bannister, CH, RJ Kaye. 1991. A rapid method for optimization of linear systems with storage. *Operations Research* **39**(2) 220–232.

Bellman, R. 1956. On the theory of dynamic programming-a warehousing problem. *Management Science* **2**(3) 272–275.

Bjork, T. 2004. *Arbitrage Theory in Continuous Time*. Oxford University Press, New York.

Brown, David B., J.E. Smith, P. Sun. 2008. Information Relaxations and Duality in Stochastic Dynamic Programs. Working Paper, Fuqua School of Business, Duke University.

Carr, P. 1988. The valuation of sequential exchange opportunities. *The Journal of Finance* **43**(5) 1235–1256.

Charnes, A., J. Dreze, M Miller. 1966. Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica* **34**(2) 307–330.

Chen, X., M. Sim, D. Simchi-Levi, P. Sun. 2007. Risk aversion in inventory management. *Operations Research* **55**(5) 828–842.

Deng, S.J., B. Johnson, A. Sogomonian. 2001. Exotic electricity options and the valuation of electricity generation and transmission assets. *Decision Support Systems* **30**(3) 383–392.

Dreyfus, S.E. 1957. An analytic solution of the warehouse problem. *Management Science* **4**(1) 99–104.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The risk-averse (and prudent) newsboy. *Management Science* **41**(5) 786–794.

Gaur, V., S. Seshadri. 2005. Hedging inventory risk through market instruments. *Manufacturing and Service Operations Management* **7**(2) 103–120.

Geman, H. 2005. *Commodities and Commodity Derivatives: Modeling and Pricing for Agriculturals, Metals and Energy*. John Wiley & Sons, West Sussex.

Geman, H., V. Nguyen. 2005. Soybean inventory and forward curve dynamics. *Management Science* **51**(7) 1076.

Gibson, R., E.S. Schwartz. 1990. Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance* **45**(3) 959–976.

Hahn, W.J., J.S. Dyer. 2008. Discrete time modeling of mean-reverting stochastic processes for real option valuation. *European Journal of Operational Research* **184**(2) 534–548.

Ho, T.S., R.C. Stapleton, M.G. Subrahmanyam. 1995. Multivariate binomial approximations for asset prices with nonstationary variance and covariance characteristics. *Review of Financial Studies* **8**(4) 1125–1152.

Hull, J.C. 1997. *Options, Futures and Other Derivatives*. Prentice Hall, Upper Saddle River, New Jersey.

Lai, G., F. Margot, N. Secomandi. 2009a. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research* **forthcoming**.

Lai, G., F. Margot, N. Secomandi. 2009b. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research* **forthcoming**.

Markland, R.E. 1975. Analyzing Multi-Commodity Distribution Networks Having Milling-in-Transit Features. *Management Science* **21**(12) 1405–1416.

Markland, R.E., R.J. Newett. 1976. Production-distribution planning in a large scale commodity processing network. *Decision Sciences* **7**(4) 579–594.

Nelson, D.B., K. Ramaswamy. 1990. Simple binomial processes as diffusion approximations in financial models. *The Review of Financial Studies* **3**(3) 393–430.

Plato, G. 2001. The soybean processing decision: Exercising a real option on processing margins. Electronic Report from the Economic Research Service, Technical Bulletin Number 1897, United States Department of Agriculture.

Prahalad, C.K. 2005. *Fortune at the Bottom of the Pyramid*. Wharton School Publishing.

Routledge, B.R., D.J. Seppi, C.S. Spatt. 2001. The spark spread: An equilibrium model of cross-commodity price relationships in electricity. Working Paper, Tepper School of Business, Carnegie Mellon University.

Schwartz, E., J.E. Smith. 2000. Short-term variations and long-term dynamics in commodity prices. *Management Science* **46**(7) 893–911.

Schwartz, E.S. 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance* **52**(3) 923–973.

Secomandi, N. 2009a. On the pricing of natural gas pipeline capacity. *Manufacturing & Service Operations Management* **forthcoming**.

Secomandi, N. 2009b. Optimal commodity trading with a capacitated storage asset. *Management Science* **forthcoming**.

Sherali, H.D. 2001. On mixed-integer zero-one representations for separable lower-semicontinuous piecewise-linear functions. *Operations Research Letters* **28**(4) 155–160.

Wu, O.Q., H. Chen. 2009. Optimal control and equilibrium behavior of production-inventory systems. Working Paper, Stephen M. Ross School of Business, University of Michigan.