Essays in optimization of commodity procurement, processing and trade operations

by

Sripad Krishnaji Devalkar

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2011

Doctoral Committee:

Professor Ravi M. Anupindi, Co-Chair
Assistant Professor Amitabh Sinha, Co-Chair
Professor Izak Duenyas
Professor Haitao Li

To all my teachers.

Acknowledgements

This dissertation would not have been possible without the help and support of many people. I am very fortunate to have had Ravi Anupindi and Amitabh Sinha as my advisors. Their insights, guidance, patience and willingness to let me make mistakes on the path to discovering better answers were invaluable to my growth as a researcher and enhancing the quality of this dissertation. They were generous with their time and were always available, even if it was only to listen to my half-baked and hare-brained ideas. Thank you, Ravi and Amitabh for your faith and confidence in me and for defying the norms of a stereotypical doctoral advisor that one reads about in PhD comics. I would also like to thank the other members of my doctoral committee, Izak Duenyas and Haitao Li. Their questions and feedback have helped in improving this dissertation and given me new ideas to extend this research.

The support of faculty members within the Operations and Management Science (OMS) department has played an important role in my growth as a scholar. I would especially like to thank Owen Wu for tolerating my unannounced visits to his office and letting me pick his brains. His help and advice on modeling and estimating commodity price processes have enhanced the numerical studies in this dissertation. I am also grateful to Roman Kapuscinski and Hyun-Soo Ahn, the chairs of the OMS doctoral program during the course of my PhD, for their support.

My journey through the doctoral program would have been far less enjoyable, had it not been for my fellow students in the OMS PhD program. The interactions, within the classroom and outside, are things I will always treasure. Amongst this group, there is one person that deserves special mention: Suleyman Demirel. The many discussions that I have had with him and his help, especially in my first two years of the program, have been instrumental in making me a better mathematician and researcher. Nothing that I say here will be enough to convey the admiration and gratitude I have for you Suleyman, so I will just say thank you for everything.

Outside the OMS community, Pranav Garg and Vivek Tandon have been close companions in this journey. They were both around when I needed them the most and I cannot thank them enough for it. Pranav's levelheadedness and calm nature were helpful in overcoming immediate disappointments and looking at things in a larger context. Vivek and I joined the PhD program at the same time. Although we joined different departments, in Vivek I found a fellow traveler and kindred soul with whom I could share the trials and tribulations of discovering an academic identity.

My lack of knowledge about administrative details and the fact that I seldom had to navigate the bureaucracy at an institution as large as the University of Michigan is thanks to the wonderful support provided by the doctoral studies office at the Ross School. Thank you, Brian Jones, Roberta Perry, Kelsey Zill, Chris Gale, Martie Boron and Linda Veltri.

Outside the Ross community, a few people were instrumental in my pursuing a PhD. Kannan Sethuraman planted the germ of an idea to pursue a PhD in Operations Management in my mind. His wholehearted support and guidance played a big part in my joining the program at Michigan. I am also grateful to Shankar Narasimhan for his advice and support since the days of my undergraduate studies. Thanks are also due to Miriam, for constantly goading me and never letting me give up on the idea of pursuing a PhD.

This dissertation would not have been possible without the love and support of my family. Although invisible, they have played a big part in making this dissertation see the light of the day. Sudarsana's presence has been a great source of strength and support through this entire process. While one may question her own sanity, given that she threw in her lot with that of a PhD student just starting the program, there is no question that she helped me preserve mine through this long, and sometimes arduous, journey. I cannot thank my parents and sister enough for their unstinted support and allowing me to pursue my interests, however baffling and inexplicable it may have appeared to them.

Dilip Veeraraghavan has had a great influence in my growth as an individual. It was upon Dilip's suggestion that I spent the final semester of my undergraduate days as a tutor for school kids and discovered the joy of teaching and interacting with curious young minds. A part of me always wonders as to what I would have made of my life, had I not spent those long hours in his office. While he is no longer around, I am sure he would have been one of the happiest people to see me achieve this milestone. He is the model of a friend, philosopher and guide that I shall always strive to be.

Table of Contents

List of Tables

Table

List of Figures

Figure

Abstract

Managing commodity price uncertainty is an integral part of many firms' business process. Firms adopt a variety of operational strategies to manage this uncertainty, subject to operational constraints such as finite procurement and processing capacities. The availability of financial derivative instruments provide firms with additional options to manage the risk from commodity operations. This dissertation explores different aspects of managing the price uncertainty for a commodity processing firm in a series of four related essays.

The first three essays consider the integrated procurement, processing and trade decisions for a firm operating a single location with procurement and processing capacity constraints under risk-neutral and risk-averse objective functions. These essays focus on deriving the optimal policy structure and developing computationally tractable heuristics where required. The first essay considers a risk-neutral firm maximizing expected profits from operations over a multi-period horizon and derives the optimal operational policy for the firm. The second essay deals with the issue of time-consistent decision making in risk-averse settings while the third essay looks at the value of operational hedging, such as excess procurement or processing capacity.

The fourth essay extends the single location problem to a network setting and considers a 'star' network configuration. While solving the network problem optimally is hard, this essay proposes heuristics based on insights from the optimal policy structure for the single node problem to address the computational complexities. In addition, this essay also proposes a myopic heuristic to manage the commodity procurement and processing decisions in a network. Numerical studies indicate that

these heuristics provide a significant improvement in expected profits, compared to heuristics used in practice.

Chapter 1

Introduction

Many firms use commodities as inputs, while for other firms commodities are an output of their production process. In some cases, firms deal with commodities as both inputs to and outputs of their production process. In general, the commodity prices are quite volatile and fluctuate over time, reflecting the dynamics of the underlying demand and supply for the commodities (Figure 1.1 provides an example of the price uncertainty that Soybean processors face). As a result, firms dealing with such commodities use a variety of strategies to manage the commodity price uncertainty.

Consider the example of the ITC Group, one of India's largest private sector companies, whose operations provide the original motivation for this research. 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

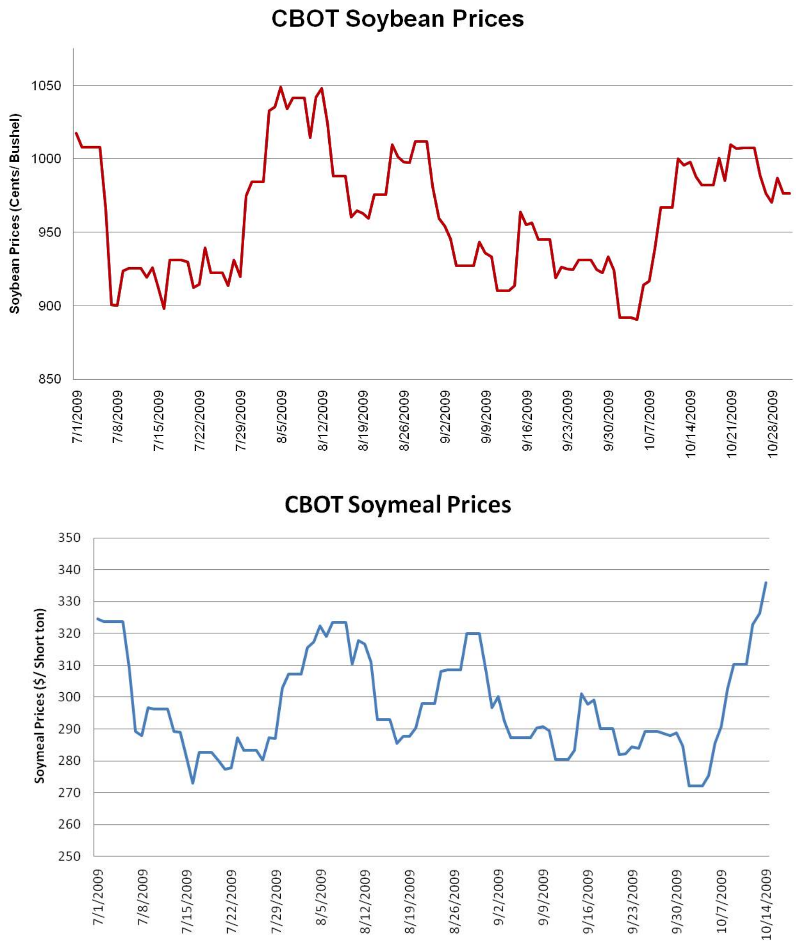

Figure 1.1: Soybean and Soymeal Prices

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 1.2.

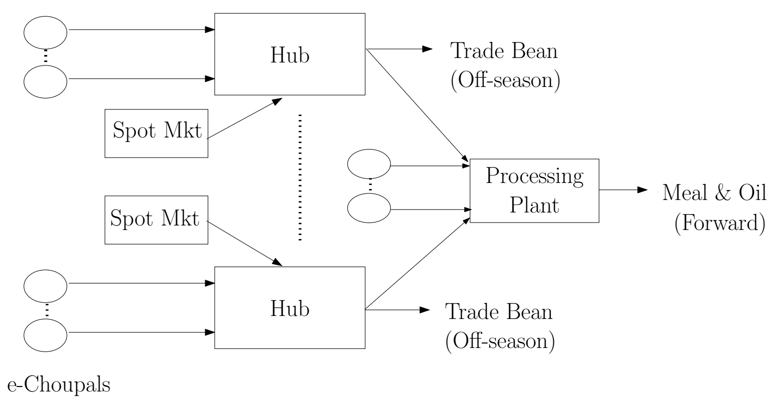

Figure 1.2: ITC e-Choupal Network.

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 soymeal, both of which are traded through various channels. Managing this network requires decisions regarding procurement and trading of commodities to maximize profits and mitigate the losses from adverse commodity price movements. 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 soymeal. Trade options include trading in open markets and with other processors. While operational decisions help manage some of the risk from the underlying commodity price uncertainty, the availability of derivative instruments on various commodities provide additional options to manage the risk effectively. Thus, ITC in addition to physical procurement, processing and trade decisions, needs to make financial hedging decisions to manage the risk in its commodity operations more effectively.

While ITC's operations provide the basic context for research, the problems considered in this dissertation are quite generic and applicable for firms in the commodities processing business. Profits for such firms are affected by both the input and output commodity prices in international exchanges and local spot markets. Additionally, inputs for processing such as agricultural commodities, metals etc. are typically procured from many different locations. In this context, the procurement, transshipment, processing and trade (of commodities) decisions for the firm are interlinked and affect the overall profits of the firm. Also, the integration of operational and financial trading decisions are essential for effective risk management and avoiding the costs of financial distress.

This dissertation consists of four essays that explore different aspects of managing price uncertainty for a commodity processing firm. This introduction briefly describes the problem considered in each essay. Section 1.1 introduces the first essay: the integrated procurement, processing and trade decisions for a commodity processing firm facing operational capacity constraints and interested in maximizing total expected profits. Section 1.2 introduces the second essay which considers risk aversion in a multi-period context and develops optimal operational and financial hedging decisions for a commodity processor. Section 1.4 introduces the problem for a firm operating a network of procurement and processing locations.

1.1. Essay 1: "Integrated Optimization of Procurement Processing and Trade of Commodities"

This essay considers a risk-neutral commodity processing firm, operating a single facility with procurement and processing capacity constraints. The firm procures an input commodity and converts the input into a processed product ('output') using the processing capacity. The firm earns revenues by selling the output using forward contracts and also by trading the input with other processors at the end of the horizon.

We model the multi-period problem for the expected profit maximizing firm using a stochastic dynamic program and characterize the optimal procurement, processing and trade policy structure. We show that the procurement and processing decisions are governed by two price and inventory dependent thresholds, while the output commodity sales commitment is akin to the exercise of a compound exchange option.

Using commodity market data for the soybean complex—soybean, soybean meal and soybean oil—we conduct numerical experiments to compare the performance of the optimal policy with that of heuristics used in practice and option valuation literature.

1.2. Essay 2: "Dynamic Risk Management of Commodity Operations: Model and Analysis"

While Essay 1 considers a risk-neutral firm, many firms in the commodities business exhibit risk aversion and use a variety of operational and financial strategies to manage risk from commodity price uncertainty. In this essay, we consider a risk-averse commodity processing firm concerned about managing the risk over a multi-period planning horizon. The firm procures an input commodity and processes it to produce an output commodity. The output commodity is sold using forward contracts, while the input itself is traded at the end of the horizon. The firm also trades financial derivative instruments to manage the commodity price risk.

In a multi-period setting, efficient risk management requires controlling risk over the entire horizon, and not just in the total payoffs at the end of the planning horizon. We propose a time-consistent dynamic risk measure, DCVaR, based on the conditional value at risk (CVaR) to model the firm's risk aversion over the planning horizon. We obtain the optimal operational—procurement, processing and trade—and financial hedging policies by formulating the risk management problem as a stochastic dynamic program. We show that the optimal operational policies are governed by price dependent inventory thresholds which, conditional on optimal financial hedging decisions, can be calculated without knowing the details of the financial hedging decisions themselves.

We develop tractable heuristics to overcome the computational complexity in determining the optimal policy parameters and provide numerical studies to illustrate the performance of the heuristics. Using numerical experiments, we also show that a time-consistent risk measure (such as the DCVaR proposed here) provides a better mean-risk tradeoff for total profits as well as better risk control over the entire horizon, compared to optimizing static risk measures such as CVaR on terminal wealth.

1.3. Essay 3: "Commodity Operations in Partially Complete Markets"

Essay 2 considers the problem of a risk-averse commodity processing firm and shows the benefit of using time-consistent risk measures to model risk aversion in a multi-period setting. While it considers trading financial instruments to manage the commodity price risk, it does not provide specific details of the structure of the financial trading decisions themselves. Further, the analysis in Essay 2 does not explore in depth the benefits of operation hedging; e.g., the benefit of having excess procurement or processing capacity to manage price uncertainty.

In this essay, we analyze the structure of the optimal financial trading policy and explore the benefits from operational hedging, in addition to financial hedging, for a commodity processing firm. We use the partially complete markets framework to model the underlying uncertainty in commodity prices and distinguish between financial market and firm specific (or private) factors. Extending the time-consistent risk measure DCVaR to this framework, we characterize the optimal financial trading policy explicitly as a portfolio which replicates the CVaR of cashflows measured over states of private uncertainty for each realization of the market uncertainty. Contingent on the optimal financial trading policy, we show that the optimal commitment policy for selling the output is identical to the risk-neutral commitment policy and that the procurement and processing decisions for the input are governed by price and horizon dependent inventory thresholds.

Under the mild restriction that the worst case expected salvage value for the input is no more than the benefit from processing and selling the output, we show that excess processing capacity (relative to procurement capacity) does not provide any additional value. On the other hand, excess procurement capacity serves as an operational hedge to manage the input commodity price uncertainty. We also characterize the value of this operational hedge analytically.

1.4. Essay 4: "Commodity Operations in a Network Environment: Model, Analysis and Heuristics"

Many commodities, e.g., corn, crude oil, are produced in different geographical areas and transported to multiple locations. Firms using these commodities as inputs generally procure them from multiple locations for a variety of reasons, including price differentials across locations and capacity constraints. Further, firms usually have processing capacities at fixed locations, requiring transshipment of the processed product to various locations for delivery. Profits for such firms are affected by the network characteristics such as transshipment costs, capacities at various locations and transportation costs. In the fourth essay, we consider the integrated commodity operations for a firm managing a network. We explore how the results for the single node problem can be extended to a network setting and study the impact of the network characteristics on the optimal policies.

We consider a risk-neutral firm operating a star network with processing capacity at a central location and procurement over multiple locations. Our analysis of the star network, which has a simple structure, shows that characterizing and computing the optimal policies is hard because of high dimensionality of the state space. However, for a special case when the transshipment costs are symmetric and the input salvage values across different locations are sufficiently close, we find that the network problem reduces to a single node problem with piecewise linear and convex cost of procurement. Based on this similarity, we propose a heuristic, termed the 'Equivalent Single Node' (ESN) heuristic, for solving the star network problem by approximating it as an equivalent single node problem. We also propose a myopic heuristic for solving the network problem; this is termed the 'Network Full Commitment' (NFC) heuristic and is based on a heuristic used in practice. We use the technique of information relaxation and dual penalties for stochastic dynamic programs to compute an upper bound on the optimal expected profits. Using commodity market data for the soybean complex, we evaluate the performance of these heuristics through numerical experiments.

Chapter 2

Integrated Optimization of Procurement, Processing and Trade of Commodities

2.1. Introduction

Profits for commodity processing firms are affected by changes in both input and output commodity prices. Typically, such firms have little influence or control over the prices of these commodities which are driven by global supply and demand shocks and determined by trading activities on global exchanges and spot markets. Under these conditions, it is important for the processing firms to coordinate their procurement, processing and trade decisions in order to maximize the total value from their operations. Further, operational constraints such as limited procurement and/or processing capacities impose additional complications, making the various decisions interdependent and the optimization of such operations a non-trivial exercise. While different aspects of the problem—procurement, processing and trade—have been studied earlier, the integrated problem itself, even for operations at a single node, has not received much attention in the literature. In practice, firms consider the interdependencies between procurement, processing and trade decisions (see Plato, 2001, for instance), but do so in a myopic fashion and ignore the dynamic nature of decisions.

In this essay, we consider a firm that procures an input commodity with the marginal cost of procurement equal to the spot price of the commodity. The firm earns revenues by processing the input commodity and committing to sell the processed outputs 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. We model the firm's multi-period optimization problem as a stochastic dynamic program, with the procurement and processing decisions in each period subject to capacity constraints. We provide a precise, analytical description of the optimal policy structure. We investigate the benefits of using a forward looking, optimization based policy relative to myopic spread-option-based policies that are used in practice. We do this by conducting numerical studies using commodity markets data for the soybean complex—soybean, soybean meal and soybean oil. To summarize our results:

1. We show that the optimal value function is separable in the input and output commodity inventories, and piecewise linear and concave in the inventory levels. We derive recursive expressions to quantify the marginal value of the input and output commodity inventories.

2. We find that it is optimal for the firm to postpone the output trade against a forward contract with given maturity to the last possible period; i.e., period just before the maturity of the forward contract and the optimal output commitment policy is similar to the exercise of a compound exchange option.

3. We characterize the optimal procurement and processing policy and find that the optimal decisions are governed by procure up to and process down to inventory thresholds, with these thresholds dependent on the realized prices and remaining horizon length.

4. Using commodity markets data for the soybean complex, we find that a myopic heuristic used in practice performs almost as well as the optimization based dynamic programming policy under normal operating conditions. However, the dynamic programming policy provides significant benefits under conditions of tight processing capacities and high price volatilities.

5. The complexity in computing the dynamic programming policy increases rapidly

as the number of output products increases. We approximate multiple output commodities as a single composite output to address this computational complexity, and find that this approximation is near-optimal.

The rest of the chapter is organized as follows. In the next section, we review literature relevant to this research and position our work. In section 2.3, we solve the integrated procurement, processing and trade decisions for a risk-neutral firm and obtain expressions for the marginal value of inventory. Section 2.3.1 presents the analysis for the case when a single output commodity is produced upon processing the input, while section 2.3.2 generalizes the result to a situation where multiple products are produced upon processing. Section 2.4 provides numerical illustrations using commodity market data for the soybean complex and describes the computation of the optimal policy when all commodity prices are driven by single factor mean reverting processes. The computational policy described here can be extended to the case of multi-factor commodity process models using heuristics which builds on the works of Brown et al. (2010) and Lai et al. (2010a). Details of the heuristic and computational results are given in section 2.5. Section 2.6 concludes with directions for future research.

2.2. Literature Review

The problem studied in this paper is related to the warehouse management problem originally studied by Bellman (1956) and Dreyfus (1957). The warehouse management problem deals with determining the optimal trading policy for a commodity with constraints on the total inventory 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 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. Secomandi (2010b) 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. In contrast to the above papers, we consider multiple commodities and, in addition to the procurement and trading decisions, incorporate a processing decision that irreversibly transforms input to outputs. Moreover, unlike a single commodity procurement and trading operation where the procure up to threshold is always less than or equal to the process down to threshold, the procure up to level can be higher than the process down to level in our model.

The methodology used in the current paper relies on characterizing the value function as a piecewise linear function, with changes in slope at integral multiples of the greatest common divisor of the procurement and processing capacities. While a similar approach has been used by Secomandi (2010b) and Nascimento and Powell (2009), they do so under the assumption of discrete price evolutions. Nascimento and Powell (2009) use the discrete price evolution assumption to prove the convergence of their approximate dynamic program (ADP), while Secomandi (2010b) uses it for computational purposes using lattices. While we use price lattices for computational studies, the characterization of the value function itself, i.e., the piecewise linear property and the marginal values of inventories, is not dependent on the assumption that prices are discretely distributed. In contrast to Secomandi (2010b), where the procure up to threshold is always less than or equal to the sell down to threshold, the procure up to levels can be higher than the process down to threshold in our context, representing arbitrage opportunities across the different commodities; i.e., the value from selling the output is higher than the cost of the input plus the processing cost. In comparison to Nascimento and Powell (2009), who characterize the marginal value of inventory of a single commodity, we model processing decisions and characterize the marginal value of inventory of both input and output commodities.

The decision making framework considered in this paper is related to the valuation of real options and exotic commodity options. The concept of spread options is 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. 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).

The existing literature has mainly focused on valuation of spread options with a given maturity; i.e., options of the European type with a single exercise date. Secomandi (2010a) uses location spread options on natural gas prices to value pipeline capacity. While the pipeline capacity places an upper limit on the total amount of natural gas that can be shipped, the unit spread option value is the same for each unit of the pipeline capacity and is not affected by the total capacity available. 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. The decision to commit processing capacity available on different future dates is modeled as the exercise of a simple spread option on the gross processing margin on that 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. Deng et al. (2001) use spark spread options on the spread between electricity and generating fuel prices to value electricity generation assets. In these papers, no inventory is carried over time and the exercise of spread options maturing on different dates is evaluated independent of each other. In our current paper, unlike the aforementioned papers, decisions across periods are linked through the storage of input inventory and operational capacity constraints, making the processing decision considered here different from the exercise of a simple spread option. In contrast to Secomandi (2010a), we also find that the marginal value of input inventory is affected by the capacity constraints.

Tseng and Barz (2002) and Tseng and Lin (2007) extend Deng et al. (2001) to include operational constraints such as minimum up/downtime, startup/shutdown times, ramp constraints etc. in the electricity generation unit commitment decisions. The main focus of both these papers is to provide a computational framework for valuing the generation assets. We focus on deriving structural results that are useful for decision making and, in the process, derive analytical expressions for the marginal value of input and output commodity inventories. Similar to Tseng and Barz (2002) and Tseng and Lin (2007), our computational study also uses a lattice framework to represent the joint evolution of the multiple commodity prices.

The capacitated procurement of the input commodity over a horizon has similarities to the exercise of swing options (Jaillet et al., 2004; Keppo, 2004). A swing option provides the option holder the flexibility to procure more or less than a baseline amount, at a fixed price and is subject to volume constraints. While we do consider capacitated procurement in the current paper, there is no baseline quantity or price around which the procurement quantity can vary. Further, unlike the swing options pricing literature which typically considers only a single commodity, the procurement decisions in our problem are driven not only by the price of the commodity being procured, but also the price of the output that is produced upon processing.

The single node problem considered here has similarities to the firm level production and inventory control problem studied in Wu and Chen (2010) for a storable input-output commodity pair. While Wu and Chen (2010) consider the optimal procurement and sales policy for the individual firm, their main focus is analysis of 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. Martınez-de Albéniz and Simón (2010) consider a related problem of commodity traders who take advantage of price spreads across locations, and model the impact of the trading decisions on price evolution at the different locations. 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 to these papers, 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 analysis carried out in this paper on the value of a forward looking dynamic programming policy relative to myopic policies is similar to the analysis in Lai et al. (2010b), who consider the real option to store liquified natural gas (LNG) in a LNG value chain. Lai et al. (2010b) develop a model which integrates LNG shipping, natural gas price evolution and inventory control and sales, and find that using a dynamic programming policy is important when the throughput of the LNG shipping process is low compared to the storage capacity. Although in a different context involving multiple commodities and processing decisions, our findings mirror theirs in that the value of a dynamic programming policy is high relative to myopic policies when the processing capacity is tight relative to the procurement capacity.

2.3. Model Formulation and Analysis

2.3.1 Single Output Commodity

Model formulation. 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, the firm can procure the input commodity from the spot market at the current spot price S_n. The firm processes the input and sells all the output using forward contracts.The procurement season for the input commodity may span multiple output forward maturities. The delivery date for forward contract ℓ is given by N_ℓ, with $\ell \in \{1, 2, \ldots, L\}$. 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 $\ell < L$ and $N_L \leq N$. Let F_n^ℓ denote the period n forward price on contract ℓ, for $n < N_\ell \leq N$. In addition to selling the output commodity, 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.

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. The marginal 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 $h_O \geq h_I$. 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.

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 (see Hull (1997) or Bjork (2004) for discussion on risk-neutral measures). 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 forward prices are equal to the futures prices and further, the futures prices are a martingale process (see Hull (1997), Section 3.9 or Bjork (2004), Section 7.6 for details). The output forward prices for each contract thus satisfy

$$\mathbb{E}_n[F_{n+1}^\ell] \quad = \quad F_n^\ell \text{ for } n < N_\ell, \, \forall \, \ell \tag{2.1}$$

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

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 input to be processed into output: m_n and c) the quantity of the output commodity to be committed for sale against forward contract ℓ: q_n^ℓ for all ℓ such that $N_\ell > n$. 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 $\ell^*(n)$ 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 $q_n^{\ell^*(n)} > Q_n + m_n$. This is possible because the output needs to be delivered only in period $N_{\ell^*(n)}$ and the firm can process in some future period(s) t between n and $N_{\ell^*(n)}$ to meet the shortfall $q_n^{\ell^*(n)} - (Q_n + m_n)$, which would require that we keep track of the shortfall against each forward contract. However, in light of the martingale property (equation (2.1)), 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 $\ell^*(n)$ where $\ell^*(n) = \underset{\ell \in \mathcal{L}(n)}{\arg\max} \left\{ \beta^{N_\ell - n} F_n^\ell - h_O \sum_{t=0}^{N_\ell - n-1} \beta^t \right\}$ where $\mathcal{L}(n) = \{\ell \leq L \text{ s.t. } N_\ell > n\}$. The term inside the maximization is the discounted forward price minus the total discounted holding costs incurred from the current period till delivery at the maturity of the forward contract. We can formulate the firm's problem as a stochastic dynamic program (SDP) in the following manner.

$$
\begin{aligned}
V_n(e_n, Q_n, \mathcal{I}_n) \;=\; \max_{\substack{0 \le x_n \le K, \\ 0 \le m_n \le \min\{C, e_n + x_n\}, \\ 0 \le q_n^{\ell^*(n)} \le Q_n + m_n}} & \Bigg\{ \left[\beta^{N_{\ell^*(n)}-n} F_n^{\ell^*(n)} - h_O \sum_{t=0}^{N_{\hat{\ell}}-n-1} \beta^t \right] q_n^{\ell^*(n)} \\
& - S_n x_n - p m_n - h_I [e_n + x_n - m_n] \\
& - h_O [Q_n + m_n - q_n^{\ell^*(n)}] \\
& + \beta \mathbb{E}_n [V_{n+1}(e_{n+1}, Q_{n+1}, \mathcal{I}_{n+1})] \Bigg\} \qquad (2.2)
\end{aligned}
$$

for $n < N$ and

$$
V_N(e_N, Q_N, \mathcal{I}_N) = \begin{cases} S_N e_N & Q_N \ge 0 \\ -\infty & \text{otherwise} \end{cases} \qquad (2.3)
$$

where the state transition equations are given by $e_{n+1} = e_n + x_n - m_n$ and $Q_{n+1} = Q_n + m_n - q_n^{\ell^*(n)}$.

The constraints on x_n and m_n in equation (2.2) 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.

Marginal value of output inventory. Consider the commitment decision in period n. By committing against *any* specific contract ℓ with $N_\ell > n$, the firm earns a revenue of $\beta^{N_\ell - n} F_n^\ell - h_O \sum_{t=0}^{N_\ell - n-1} \beta^t$ 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 $N_\ell - 1$, the last opportunity to commit against contract ℓ. By postponing the decision to period $N_\ell - 1$, the firm retains the option not to commit the unit of output to contract ℓ if some other contract ℓ' provides a higher revenue. Extending this argument, we have the following result.

Lemma 2.1. *It is optimal to commit to sell output using contract ℓ, if at all, only in period $N_\ell - 1$ for $\ell = 1, 2, \ldots, L$.*

To determine if it is optimal to commit against a specific contract ℓ, consider the case of $L = 2$, with maturities N_1 and N_2 respectively. In period N_2, it will be optimal for the firm to commit all the available output inventory against contract 2, as this is the last opportunity to commit the output inventory for sale against *any* forward contract and all uncommitted output inventory beyond period N_2 will earn zero revenue. Therefore, in any period n such that $N_1 \leq n < N_2$, the marginal value of output inventory is equal to $\beta^{N_2-n}\mathbb{E}_n\left[F^2_{N_2-1}\right] - h_O \sum_{t=0}^{N_2-1-n} \beta^t$. In period $N_1 - 1$, it will be optimal for the firm to commit against contract 1, if and only if $\beta F^1_{N_1-1} - h_O > \beta^{N_2-N_1+1}\mathbb{E}_{N_1-1}\left[F^2_{N_2-1}\right] - h_O \sum_{t=0}^{N_2-N_1} \beta^t$. Further, the optimal commitment decision is 'all or nothing'; i.e., if it is optimal to commit against contract 1, then it is optimal to commit all the available output inventory, $Q_{N_1-1} + m_{N_1-1}$. Extending this analysis to a more general case of $L > 2$, we can prove the following result about the marginal value of output inventory and the optimal commitment policy.

Lemma 2.2. *The marginal value of a unit of output inventory in period n, denoted by Δ_n, is given by*

$$\Delta_n = \begin{cases} 0 & \text{if } n \geq N_L \\ \beta \max\left\{F^\ell_n, \mathbb{E}_n\left[\Delta_{n+1}\right]\right\} - h_O & \text{if } n = N_\ell - 1 \text{ for } \ell = 1, \ldots, L \\ \beta\mathbb{E}_n\left[\Delta_{n+1}\right] - h_O & \text{otherwise} \end{cases} \tag{2.4}$$

and the optimal quantity to commit against contract ℓ is given by

$$q^*_{N_\ell-1} = \begin{cases} 0 & \textit{if } F^\ell_{N_\ell-1} \leq \mathbb{E}_{N_\ell-1}\left[\Delta_{N_\ell}\right] \\ Q_{N_\ell-1} + m_{N_\ell-1} & \textit{otherwise} \end{cases} \tag{2.5}$$

Substituting the optimal commitment quantity in the objective function of equation (2.2), and using an induction argument, we can show that the value function is linear in Q_n and moreover, separable in Q_n and e_n. We can write

$$V_n(e_n, Q_n, \mathcal{I}_n) = \Delta_n Q_n + U_n(e_n, \mathcal{I}_n) \text{ for } n < N \tag{2.6}$$

$$V_N(e_N, Q_N, \mathcal{I}_N) = U_N(e_N, \mathcal{I}_N) \tag{2.7}$$

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

$$\begin{aligned} U_n(e_n, \mathcal{I}_n) = \max_{\substack{0 \leq x_n \leq K, \\ 0 \leq m_n \leq \min\{e_n + x_n, C\}}} \Big\{ & [\Delta_n - p] m_n - S_n x_n \\ & - h_I [e_n + x_n - m_n] \\ & + \beta \mathbb{E}_n [U_{n+1}(e_{n+1}, \mathcal{I}_{n+1})] \Big\} \end{aligned} \tag{2.8}$$

for $n < N$ and

$$U_N(e_N, \mathcal{I}_N) = S_N e_N \tag{2.9}$$

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 (cf., Carr, 1988); 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.

Marginal value of input inventory. We next turn to determining the marginal value of input inventory. As the firm has limited processing capacity, the marginal value-to-go of input inventory depends on the total input inventory available. For instance, when the ending input inventory e_{n+1} is greater than the remaining processing capacity $(N-(n+1))\times C$, the marginal value-to-go is equal to the discounted expected salvage value minus the total input holding costs, irrespective of the value from processing, $\Delta_n - p$. The processing decision is therefore dependent on the input inventory levels; i.e., the decision depends on whether $\Delta_n - p$ is higher or lower than the marginal value-to-go of unprocessed input at the given input inventory levels. We now derive expressions for the marginal value of input inventory, with the aim of using them to determine the optimal procurement and processing decisions 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 2.1 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. (For notational convenience, we do not show the dependence of Θ_n^k on $\mathcal{I}_n$, explicitly.)

Theorem 2.1. *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 $\mathcal{I}_n$.*

For all n, let $\Theta_n^k \triangleq \infty$ for $k \in \mathbb{Z}^- \cup \{0\}$. For any period $n \leq N$ and positive

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

integer k, we have

$$\Theta_n^k = \begin{cases} S_N & \text{if } n = N \\ \max\left\{\Omega_n^{(k+b)}, \min\left\{S_n, \Omega_n^{(k)}\right\}\right\} & \text{if } n < N \end{cases} \tag{2.10}$$

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*

$$\Omega_n^{(j)} \quad = \quad \max\left\{\beta\mathbb{E}_n[\Theta_{n+1}^j] - h_I, \min\left\{\Delta_n - p, \beta\mathbb{E}_n[\Theta_{n+1}^{j-a}] - h_I\right\}\right\} \tag{2.11}$$

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. 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 for $e_t \in [(k-1)D, kD)$. Also, U_t is continuous in e_t and $\Theta_t^k \geq \Theta_t^{k+1}$ for all integers $k \geq 1$.

When $e_t \in [(k-1)D, kD)$ for $k \geq (N-t)a+1$, we have $e_t \geq (N-t)aD = (N-t)C$; i.e., there is not enough processing capacity available over the remaining horizon to process all the available input inventory. Thus, the marginal unit of input inventory can only be salvaged and the marginal value of input for all $e_t \geq (N-t)C$ is equal to the expected salvage value net of input holding costs; i.e., $\Theta_t^k = \Theta_t^{(N-t)a+1} = \beta^{N-t-1}\mathbb{E}_n[S_N] - h_I \sum_{m=0}^{N-t-1} \beta^m$ for all $k \geq (N-t)a+1$.

We have

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

where

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

Let $y_n = e_n + x_n$ denote 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 underlying L_n as

$$\big((\Delta_n - p) - (\beta \mathbb{E}_n[\Theta^j_{n+1}] - h_I) \big) \times m_n + \big(\beta \mathbb{E}_n[\Theta^j_{n+1}] - h_I \big) \times y_n + \beta \mathbb{E}_n[\lambda^j_{n+1}] \quad (2.12)$$

for $y_n - m_n \in [(j-1)D, jD)$ and 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, $\big((\Delta_n - p) - (\beta \mathbb{E}_n[\Theta^j_{n+1}] - h_I) \big)$, decreases since $\Theta^j_{n+1} \ge \Theta^{(j+1)}_{n+1}$. 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 processing quantity m^*_n

as

$$m_n^* = \begin{cases} C & \text{if } \beta\mathbb{E}_n[\Theta_{n+1}^{s-a}] - h_I \leq \Delta_n - p \\ y_n - \hat{r}_n D & \text{if } \beta\mathbb{E}_n[\Theta_{n+1}^{s}] - h_I \leq \Delta_n - p < \beta\mathbb{E}_n[\Theta_{n+1}^{s-a}] - h_I \\ 0 & \text{if } \Delta_n - p < \beta\mathbb{E}_n[\Theta_{n+1}^{s}] - h_I \end{cases} \tag{2.13}$$

where $\hat{r}_n = \max\left\{r \in \mathbb{Z}^+ \cup \{0\} \text{ s.t. } \beta\mathbb{E}_n[\Theta_{n+1}^{r}] - h_I > \Delta_n - p\right\}$. Upon substituting m_n^* corresponding to each of the three cases in the objective function (2.12), we have for $y_n \in [(s-1)D, sD)$

$$L_n(y_n, \mathcal{I}_n) = \begin{cases} (\beta\mathbb{E}_n[\Theta_{n+1}^{s-a}] - h_I)y_n + \Upsilon_n^{s,1} & \text{if } \beta\mathbb{E}_n[\Theta_{n+1}^{s-a}] - h_I \leq \Delta_n - p \\ (\Delta_n - p)y_n + \Upsilon_n^{s,2} & \text{if } \beta\mathbb{E}_n[\Theta_{n+1}^{s}] - h_I \leq \Delta_n - p \\ & \text{and } \Delta_n - p < \beta\mathbb{E}_n[\Theta_{n+1}^{s-a}] - h_I \\ (\beta\mathbb{E}_n[\Theta_{n+1}^{s}] - h_I)y_n + \Upsilon_n^{s,3} & \text{if } \Delta_n - p < \beta\mathbb{E}_n[\Theta_{n+1}^{s}] - h_I \end{cases}$$

where $\Upsilon_n^{s,\cdot}$ are constants independent of y_n for $y_n \in [(s-1)D, sD)$. Combining all three cases above, we can write

$$L_n(y_n, \mathcal{I}_n) = \max\left\{\beta\mathbb{E}_n[\Theta_{n+1}^{s}] - h_I, \min\left\{\Delta_n - p, \beta\mathbb{E}_n[\Theta_{n+1}^{s-a}] - h_I\right\}\right\}y_n + \Upsilon_n^{s}$$

for $y_n \in [(s-1)D, sD)$, where Υ_n^s denotes the relevant 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 $\Omega_n^{(s)}$, where $\Omega_n^{(s)}$ is given by equation (2.11). Thus, $\Omega_n^{(s)}$ denotes the marginal value of a unit of input inventory *after* procurement but *before* processing. We now have

$$U_n(e_n, \mathcal{I}_n) = \max_{e_n \le y_n \le e_n + K} \{L_n(y_n, \mathcal{I}_n) - S_n(y_n - e_n)\} \tag{2.14}$$

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

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

$$y_n^* = \begin{cases} e_n + K & \text{if } \Omega_n^{(k+b)} \ge S_n \\ \hat{s}_n D & \text{if } \Omega_n^{(k)} \ge S_n > \Omega_n^{(k+b)} \\ e_n & \text{if } S_n > \Omega_n^{(k)} \end{cases} \tag{2.15}$$

where $\hat{s}_n = \max\left\{s \in \mathbb{Z}^+ \cup \{0\} \text{ s.t. } \Omega_n^{(s)} > S_n\right\}$. Substituting y_n^* in the objective function of (2.14), we get

$$U_n(e_n, \mathcal{I}_n) = \max\left\{\Omega_n^{(k+b)}, \ \min\left\{S_n, \Omega_n^{(k)}\right\}\right\} e_n + \Psi_n^k \text{ for } e_n \in [(k-1)D, kD)$$

where Ψ_n^k is a constant independent of e_n for $e_n \in [(k-1)D, kD)$.

In the above expression, notice that the slope of e_n is constant for $e_n \in [(k-1)D, kD)$ for all positive integers k. Further, by the induction hypothesis, we have $\Theta_n^k \ge \Theta_n^{k+1}$, where $\Theta_n^k = \max\left\{\Omega_n^{(k+b)}, \ \min\left\{S_n, \Omega_n^{(k)}\right\}\right\}$. Thus, U_n is piecewise linear with non-increasing slopes which change only at integral multiples of D. Finally, by equation (2.11), we have $\Omega_n^{(s)} = \beta \mathbb{E}_n\left[\Theta_{n+1}^{(N-[n+1])a+1}\right] - h_I$ for all $s \ge (N-n)a + 1$, which leads to $\Theta_n^k = \Omega_n^{(k)} = \beta \mathbb{E}_n\left[\Theta_{n+1}^{(N-[n+1])a+1}\right] - h_I$ for all $k \ge (N-n)a+1$, completing the proof. $\square$

Optimal policy structure. Theorem 2.1 shows that the optimal procurement and processing policy is governed by two price and horizon dependent inventory thresholds, $\hat{s}_n D$ and $\hat{r}_n D$. In order to compare these thresholds, it is useful to restate the optimal processing policy, obtained by substituting the optimal procure up to level given by equation (2.15) into equation (2.13), as follows.

$$m_n^* = \begin{cases} C & \text{if } \Omega_n^{(k)} < \Delta_n - p \\ \min\{(y_n^* - \hat{r}_n D)^+, C\} & \text{if } \Omega_n^{(k)} \geq \Delta_n - p \geq \Omega_n^{(k+b)} \\ 0 & \text{if } \Omega_n^{(k+b)} > \Delta_n - p \end{cases} \tag{2.16}$$

where $\hat{r}_n = \max\left\{r \in \mathbb{Z}^+ \cup \{0\} \text{ s.t. } \Omega_n^{(r)} > \Delta_n - p\right\}$.

Consider the situation when the 'processing margin' from procuring and processing is negative; i.e., $\Delta_n - p - S_n \leq 0$. Any procurement in the current period is beneficial only if the expected marginal value-to-go of the procured unit is greater than S_n. Similarly, it is optimal to process whenever the benefit from processing, $\Delta_n - p$, is greater than the expected marginal value-to-go. By concavity of the value function, $\hat{s}_n D \leq \hat{r}_n D$ and the starting inventory level can be divided into three regions: a) $e_n \in [0, \hat{s}_n D)$ where it is optimal to only procure input, b) $e_n \in [\hat{s}_n D, \hat{r}_n D]$ where it is optimal to neither procure nor process any input and c) $e_n \in (\hat{r}_n D, \infty)$ where it is optimal to only process the input. The optimal procurement and processing quantities are given by $x_n^* = y_n^* - e_n = \min\{K, (\hat{s}_n D - e_n)^+\}$ and $m_n^* = \min\{C, (e_n - \hat{r}_n D)^+\}$. It is important to notice that even though the processing margin is negative, it is still optimal to process when the input inventory is sufficiently high. On the other hand when the processing margin is positive, i.e., $\Delta_n - p - S_n > 0$, there is benefit from procuring and processing the input immediately. Thus, for some starting input inventory levels, it may be optimal to both procure and process the input. This fact makes it difficult to divide the starting input inventory level into mutually exclusive regions where only one of the actions, procurement or processing, is optimal. At least one of the two activities is at capacity for all starting inventory levels and both, procurement and processing of the input, are optimal for some inventory levels.

We illustrate the features of the optimal policy using an example. To make the intuition clear, and keep the exposition simple, we consider a 3–period problem with deterministic prices, no holding costs and $\beta = 1$.

Example. Consider the situation where $K = C$ and the output commodity prices are such that $\Delta_1 - p = \Delta_2 - p \equiv \Delta - p$. The input spot prices in periods 1 and 2 and the salvage value in period 3 are such that $S_3 < S_1 < \Delta - p < S_2$.

Now, consider the procurement decision in period 1. Because $S_1 < \Delta_1 - p$, it is optimal for the firm to procure input to meet period 1's processing requirements. Because $S_1 < \Delta_2 - p < S_2$, it is optimal to procure for period 2's processing requirements in period 1 itself. Finally, because $S_1 > S_3$, it is not optimal to procure for salvaging at the end of the horizon. The total quantity that can be processed over periods 1 and 2 is equal to $2C$, and therefore the optimal procurement quantity in period 1 is given by $x_1^* = \min\{K, (2C - e_1)^+\}$.

In this example, we have $a = b = 1$. Using equations (2.11) and (2.10), we can calculate $\Omega_1^{(1)} = \Omega_1^{(2)} = \Delta - p$ and $\Omega_1^{(k)} = S_3$ for $k \geq 3$. We see that $\hat{s}_1 D = 2D$ for period 1. The optimal procurement quantity in period 1 is therefore given by $x_1^* = y_1^* - e_1 = \min\{K, (2C - e_1)^+\}$, corresponding to the first two cases in equation (2.15).

The benefit from processing is identical in periods 1 and 2 and greater than the salvage value. Thus, it is optimal for the firm to process all the available input inventory up to processing capacity. The optimal processing quantity in period 1 is thus given by $m_1^* = \min\{C, e_1 + x_1^*\}$. We see that $\hat{r}_1 D = 0$, and the optimal processing quantity corresponds to the second case in equation (2.16).

Figure 2.1 illustrates the optimal procurement and processing quantities in period 1 along with $\Omega_1^{(k)}$ values, for different starting inventory levels.

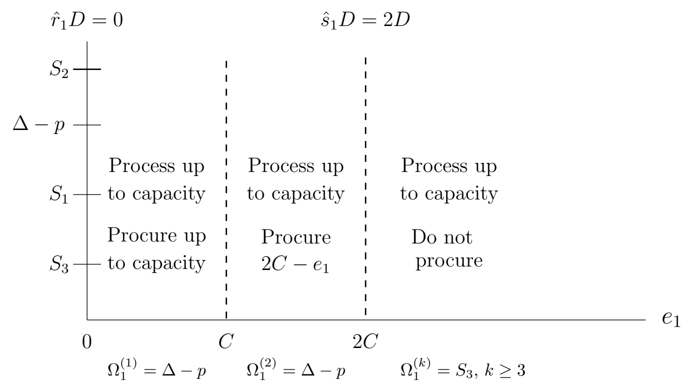

Figure 2.1: Illustration of optimal policy

2.3.2 Multiple Output Commodities

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 obtained in the previous section can be extended to the case when multiple output commodities are produced upon processing the input. To keep the exposition simple, we illustrate the case when two products are produced upon processing the input; the extension to more products is straightforward.

Let one unit of input when processed yield α_M units of product M and α_O units of product O, with α_M and α_O non-negative and $0 < \alpha_M + \alpha_O \leq 1$ (one could think of M and O to denote meal and oil in the soybean processing context). 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. 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, as there are no capacity constraints on the commitment decision itself. Thus, similar to the single output case, the optimal commitment policy for *each* output commodity is given by Lemma 2.1. Also, the marginal value of inventory for output j, denoted by $\Delta_n^{(j)}$, is given by equation (2.4). The expected benefit from processing in period n is therefore equal to $\sum_{j=M,O} \alpha_j \Delta_n^{(j)} - p$. The marginal value of input inventory, optimal procurement and processing policy are given by equations (2.10), (2.15) and (2.16), with $\Delta_n = \sum_{j=M,O} \alpha_j \Delta_n^{(j)}$.

2.4. Numerical Study

In this section, we illustrate our analytical results using numerical studies. We consider the soybean procurement and processing decisions as the context and use commodity market data for the soy complex for our numerical studies.

While the analytical results derived in Section 2.3 did not depend on the specific dynamics of the various commodity prices, computing the marginal values and optimal policy parameters does depend on the specific price processes. Single-factor mean-reverting price processes have often been used to model the spot price processes for various commodities, including agricultural commodities (cf. Geman (2005), Chapter 3). These models capture an essential feature of commodity spot prices, which is that commodity prices tend to revert to a mean level. An attractive feature of the single-factor mean-reverting price processes is their analytical tractability. While other multi-factor price processes are also used to model commodity prices (see discussion at the end of this section), in this section we model the various commodity prices as single-factor mean-reverting processes and demonstrate the computation of the optimal policy using binomial lattices to model the joint evolution of the commodity prices. We compare the performance of the optimal policy (described in Sections 2.3.1 and 2.3.2) with that of heuristics used in practice and the option valuation literature. Specifically, we consider two heuristics: a) modeling multiple outputs produced upon processing as a single, composite product to determine the input procurement and processing policies and b) a myopic, full commitment policy which uses the net margin from processing and committing all the output immediately to determine the procurement and processing decisions.

2.4.1 Implementation

Modeling the commodity price processes. We use a single factor, mean-reverting price process as in Jaillet et al. (2004) to describe the evolution of the spot prices of the various commodities under the risk-neutral measure. Specifically, $S_i(t)$, the spot price of commodity i at time t is modeled as $\ln S_i(t) = \chi_i(t) + \mu(t)$, where $\chi_i(t)$ is the logarithm of the deseasonalized price and $\mu(t)$ is a deterministic factor which captures the seasonality in spot prices. The deseasonalized price $\chi_i(t)$ follows a mean-reverting process given by $d\chi_i(t) = \kappa_i(\xi_i - \chi_i(t))dt + \sigma_i dW_i(t)$ where κ_i is the mean-reversion coefficient, ξ_i is the long run log price level, σ_i is the volatility and $dW_i(t)$ is the increment of a standard Brownian motion.

Data and estimation of the price process parameters. The parameters of the spot price process under the risk-neutral measure can be estimated by calibrating them to the observed futures prices for the various commodities, as described in Jaillet et al. (2004). Specifically, the futures price at time t, for delivery at $T \geq t$ is given by $F_i(t,T) = \mathbb{E}^{\mathbb{Q}}\left[S_i(T)|\mathcal{I}(t)\right]$ where $\mathbb{Q}$ denotes the risk-neutral probability measure and we have

$$\ln F_i(t,T) = \mu(T) + \left(1 - e^{-\kappa_i(T-t)}\right)\xi_i + e^{-\kappa_i(T-t)}\chi_i(t) + \frac{\sigma_i^2}{4\kappa_i}\left[1 - e^{-2\kappa_i(T-t)}\right]$$

The futures price information on futures contracts traded on the Chicago Board of Trade (CBOT) for different maturities on each trading day of the month of June 2010 was used to calibrate the parameters for soybean, soybean meal and soybean oil spot price processes. Futures contracts with the nearest 9 maturities for soybean, nearest 13 maturities for soybean meal and nearest 12 maturities for soybean oil were used for the calibration. While contracts with further maturities are traded for each commodity, these contracts were not included in the calibration as they had very little

Table 2.1: Price Process Parameters

	Soybean	Soybean Meal	Soybean Oil
Mean-Reversion Coeff κ_i	0.229	0.656	0.399
Longrun Log level ξ_i	6.738	5.500	3.734
Volatility σ_i	0.244	0.270	0.233
Seasonality Factor $e^{\mu(t)}$			
Jan	0.992	0.981	1.000
Feb	0.992	0.981	1.000
Mar	0.998	0.987	1.000
Apr	0.998	0.987	1.000
May	1.000	0.989	1.000
Jun	1.000	0.989	1.000
Jul	1.017	1.066	1.000
Aug	1.010	1.026	1.000
Sep	0.991	0.995	1.000
Oct	0.991	0.976	1.000
Nov	0.989	0.976	1.000
Dec	0.989	0.980	1.000

trading volume on each of the trading days used in the sample. For each trading day, the parameters were estimated by minimizing the sum of the absolute deviations between the actual and estimated futures prices for various maturities.

The minimization was carried out by approximating the seasonality factors to be constant between different maturity months in a year and imposing a normalization constraint such that $\sum_{t=1}^{t=12} \mu(t) = 0$. In addition, we also impose the constraint that the estimated 30-day and 60-day volatilities match the implied volatility information for each commodity. The implied volatility is the volatility implied by the market price of the option based on an option pricing model, and this data was obtained from the Bloomberg service. The average of the estimated parameters obtained over each trading day are given in Table 2.1 and used to model the price processes. The standard errors for the key parameters and root mean squared errors (RMSE) between the observed and estimated prices for each commodity are given in Table 2.2.

The various commodities are related through input-output processes, and the un-

Table 2.2: Estimation Errors

	Soybean	Soybean Meal	Soybean Oil
Standard Errors			
Mean-Reversion Coeff κ_i	0.0310	0.0252	0.0049
Longrun Log level ξ_i	0.0028	0.0035	0.0032
Volatility σ_i	0.0006	0.0005	0.0002
RMSE (% terms)			
	1.40%	2.15%	1.58%

Table 2.3: Correlation Between Weekly Returns

	Soybean	Soybean Meal	Soybean Oil
Soybean	1	0.846	0.654
Soybean Meal		1	0.411
Soyean Oil			1

derlying uncertainty in their price processes are likely to be correlated. We estimated the correlation between the Brownian motion increments for the three commodities using historical weekly returns on the nearest maturing futures contracts over the time period 1/1/2000 – 12/31/2009. The estimated correlations are given in Table 2.3.

Computation of the optimal policy. For computing the optimal policy, we use the re-combining binomial tree procedure described in Peterson and Stapleton (2002), which can handle mean reversion in prices, to discretize the dynamics of the price processes and approximate the joint evolution of the spot price of the various commodities. Each period in the discrete binomial tree corresponds to a week and we discretize the price process with δ steps between each period. In our computational studies, we set $\delta = 5$. We have $[(n-1)\delta + 1]^J$ nodes in the tree for period n, where J is the total number of commodities whose joint price evolution is approximated. At each node in the tree, we can compute the forward price F_n^ℓ for $l = 1, 2, \dots, L$ for each output commodity using the discrete transition probabilities at that node. Finally, using equations (2.4) and (2.10), we can compute Δ_n for each output and Θ_n^k for $k = 1, \dots, (N-n)a + 1$, and thereby the procurement and processing policy at each node in the tree.

We evaluate the performance of this policy using Monte Carlo simulation. We generate sample paths of prices for each period $n = 1, 2, \ldots, N$ by sampling from the continuous time price process for the commodities. We round the realized input and output spot prices to the closest node in the binomial tree and obtain the procurement, processing and commitment quantity corresponding to the node and inventory level. Expected profits from the policy are computed as the average profit over 10,000 sample paths. All the numerical studies were conducted using `MATLAB (version 7.9.0 R2009b)` software on a `Dell Optiplex 755` with `E8400 3.00 GHz Intel Core2 Duo CPU and 2 GB RAM`, running `Windows Vista Ent.`.

While we refer to the policy computed above as the optimal policy, it is optimal only if the various commodity prices evolve as the binomial lattices. In reality, the binomial lattice is an approximation of the true, continuous time and space price processes and strictly speaking, the policy is only an optimization based policy (our numerical experiments indicate that the gap between the value of the policy computed using lattices and the value estimated using Monte Carlo simulation is small enough that making this distinction is not important).

Other operational parameters. For all the numerical studies, we set the variable cost of processing p to equal 72 cents / bushel, which corresponds to about 35% of the gross margin from processing one bushel of soybean, based on the long run average prices of the three commodities. This value of the processing cost is close to the average processing costs estimated for the US soybean processing industry (Soyatech, 2008). The procurement and processing capacities are set to 5 and 3 units respectively. These capacities can be considered to be in multiples of bushels, e.g., million bushels. For the base case, we set processing capacity to 60% of total procurement capacity, which is roughly the percentage of soybeans produced in the United States that were estimated to have been crushed 2008 and 2009 (Ash, 2011).

Table 2.4: Optimal Expected Profits for Different Horizon Lengths

Horizon Length (# of Forwards)	Forward Maturities	Expected Profits (Std. Error)	Avg. CPU Time
5	(5)	1557.65	8.68 sec.
(1)		(0.33%)	
10	(5,9)	3095.66	206.74 sec.
(2)		(0.44%)	
20	(5,9,18)	6954.21	4284.39 sec.
(3)		(0.69%)	

We leave the exact units for the capacities unspecified as only the relative values of the procurement and processing capacities matter for computing the policies and multiplying both the capacities by a common factor will scale the expected profits also by the same factor. We assume the physical holding costs for the various commodities are negligible and normalize them to zero.

2.4.2 Numerical Results

We conduct numerical studies to compute the expected profits for the firm from its procurement and processing operations over the procurement season ranging from August to December. We initialize the prices for all the commodities to their long run average values at the beginning of the planning horizon and evaluate the performance of the policy for different horizon lengths. Table 2.4 gives the optimal expected profits for different horizon lengths, when the firm uses all forward contracts available over the horizon for each output commodity.

The optimal policy above is obtained by modeling the joint evolution of the input and individual output commodity prices and using the results in Section 2.3.2. As the number of output commodities produced upon processing increases, the number of nodes in the binomial tree used to represent the joint evolution of the various prices increases exponentially. Thus, the computational complexity increases quickly as the number of output commodities increases, requiring one to consider tractable approximations.

Composite Output Approximation. A potential approach to compute heuristic policies when multiple output commodities are produced is to model all the outputs together as a single 'composite' product and model the price process for this composite product. This is similar to the approach used by Borovkova et al. (2007) in the valuation of basket options, i.e., options on a linear combination of different assets, where the entire basket of commodities is modeled as a single 'composite' product. We model a hypothetical, composite output whose price in any period is equal to the total value of soybean meal and soybean oil produced upon processing one bushel of soybeans, where the value is calculated based on the current prices of the two products. As only futures instruments are publicly traded for the different output commodities, we consider futures instruments for the composite output as well, where the futures price for a particular maturity is a combination of the futures prices of the individual output commodities , to estimate the price process parameters for the composite output. The parameters of the price process for the composite output can be estimated as described in Section 2.4.1 by considering the hypothetical futures instruments for the composite output. The joint evolution of the input and composite output can be modeled using a binomial tree and a heuristic policy computed using the results for the single output case in Section 2.3.1. This heuristic yields a feasible policy for the original model with separate output commodities and the total expected profits from following such a heuristic policy and the gap with respect of the optimal expected profits for different horizon lengths are shown in Table 2.5.[2]

As seen from Table 2.5, approximating the multiple outputs as a single composite output comes with very little loss in optimality. The composite output approximation is also computationally far less burdensome than the optimal policy, as seen from the CPU times.

[2]Unless indicated, the gaps shown in all tables are significant with $p < 0.05$.

Table 2.5: Expected Profits Using Composite Output (CO) Approximation

Horizon Length (# of Forwards)	Forward Maturities	Expected Profits (Std. Error)	Gap (as % of Optimal)	Avg. CPU Time
5	(5)	1557.71	0.0%†	1.17 sec.
(1)		(0.33%)		
10	(5,9)	3061.71	1.10%	11.89 sec.
(2)		(0.46%)		
20	(5,9,18)	6917.71	0.52%	112.48 sec.
(3)		(0.69%)		

†$p - value > 0.1$

The composite output approximation, in addition to approximating the joint evolution of two output commodity prices as a single composite price, also leads to a lower flexibility in the commitment decision for the two output products. This is because when the composite output is committed for sale against a forward contract, both the underlying output products are committed for sale against their respective forward contracts, maturing in the same period. This is not necessarily the case under the optimal policy, where the commitment decisions for the individual outputs are independent of each other. The results in Table 2.5 imply that the loss in value by ignoring this flexibility in commitment is negligible. Further, the loss in information because of approximating the outputs by a single product has negligible impact on the total expected profits.

Full commitment policy. We evaluate the benefit of following an optimal policy by comparing the optimal expected profits with the expected profits from following a myopic policy, which only considers the value from processing and committing the output immediately in the same period. Under this myopic policy, termed the full commitment policy, the firm procures up to the minimum of procurement and processing capacities if there exists a positive margin from processing and committing the output immediately and nothing otherwise. Notice that full commitment policy ignores the 'option' value from postponing commitment of the output, as also

Table 2.6: Expected Profits From Full Commitment (FC) Policy

Horizon Length (# of Forwards)	Forward Maturities	Expected Profits (Std. Error)	Gap (as % of Optimal)
5	(5)	1563.56	-0.38%†
(1)		(0.11%)	
10	(5,9)	3106.67	-0.36%†
(2)		(0.16%)	
20	(5,9,18)	6829.88	1.79%
(3)		(0.23%)	

†$p - value > 0.1$

the value from holding and trading the input inventory at the end of the horizon. The expected profits and gap with respect to optimal profits from following the full commitment policy for different horizon lengths are shown in Table 2.6.[3]

The results in Table 2.6 suggest that the benefits of integrated decision making are negligible, compared to a myopic policy. However, these results are for the base set of parameters and do not necessarily imply the same behavior under all circumstances. To investigate this issue, we consider sensitivity of the different policies to two key parameters; processing capacity and price volatilities.[4]

Impact of processing capacity. When processing capacity is limited compared to the procurement capacity, we expected the value of integrated decision making to be higher. This is because when the input spot prices are low, the optimal policy is likely to procure input for current period processing as well as for the future. The myopic policy however does not do so. Further, including the option to trade input inventory at the end of the horizon is more valuable when processing capacity is limited. The results in Table 2.7, which shows the expected profits under the three policies as the

[3]The negative gaps are because the optimal policy is computed assuming the various commodities prices evolve in discrete space and time, while the performance of the policies are evaluated using a Monte Carlo simulation which samples from the continuous time and space price processes. As indicated, the negative values for the gaps are statistically insignificant. The same explanation holds for negative gaps seen in Tables 2.7 and 2.8 also.

[4]We also ran sensitivity analysis by varying the correlation between the different price processes. For various values of the correlation factors, we observed gaps that ranged from 0.85% to 1.32%.

Table 2.7: Impact of Processing Capacity
($N = 20$, $L = 3$, $N_\ell = \{5, 9, 18\}$)

Processing Capacity (C) (as % of K)	Expected Profits FC	Expected Profits CO	Expected Profits Optimal	Gap (% of Optimal) FC	Gap (% of Optimal) CO
20%	2274.03	2495.24	2511.97	9.4%	0.67%
40%	4545.13	4673.40	4710.69	3.5%	0.79%
60%	6813.65	6852.33	6907.15	1.35%	0.79%
80%	9064.14	9030.05	9103.55	0.43%†	0.81%
100%	11381.80	11013.50	11295.94	-0.76%†	2.50%

$^\dagger p - value > 0.1$

processing capacity is varied from 20% to 100% of the procurement capacity, support this intuition.

Compared to a myopic policy, the benefits from using the optimal policy can be as high as 9.4% for highly constrained processing firms and reduces as the processing capacity increases (the negative value in the last row is statistically insignificant). On the other hand, we notice that the gap between the optimal profits and the profits using the composite output approximation increases with the processing capacity, because more of the input procured is processed and sold as output. Thus, the loss in flexibility to commit the different outputs to contracts maturing at different dates has a higher impact. However, the maximum gap, at 2.5%, is still low.

Impact of price volatilities. All three policies–full commitment, composite output and optimal–have option like features. While the full commitment policy is equivalent to the exercise of an European spread option between the output and input prices, the composite output and optimal policies model the output commitment decision as a compound exchange option, in addition to modeling the procurement and processing decisions based on spread options. Given this, the expected profits under each policy increases with commodity price volatilities, as seen in Table 2.8.

We also notice that the gap between the optimal policy and the full commitment

Table 2.8: Impact of Price Volatility
($N = 20,\ L = 3,\ N_\ell = \{5, 9, 18\}$)

Volatility	Expected Profits			Gap (% of Optimal)	
$(\sigma_B, \sigma_M, \sigma_O)$	FC	CO	Optimal	FC	CO
0.25	6787.03	6827.70	6849.08	0.90%	0.31%
0.30	6899.90	6997.81	7029.71	1.85%	0.45%
0.35	7041.81	7105.27	7253.73	2.92%	2.05%
0.40	7216.42	7541.87	7546.57	4.37%	0.06%
0.45	7429.53	7873.10	7893.37	5.88%	0.26%
0.50	7675.51	8305.80	8297.09	7.49%	-0.10%†

$^\dagger p - value > 0.1$

policy increases as the volatility increases. This is because the full commitment policy models the marginal value of output based on the realized output prices and does not account for the exchange option inherent in the output commitment decision. Further, the full commitment policy does not procure additional input in periods with low input spot price for processing needs in future periods. The value of this opportunistic input procurement also increases as price volatility increases.

In summary, approximating multiple outputs using a hypothetical, single composite output product is near-optimal and captures almost the entire value from following an optimal policy. The value of integrated decision making can be significant for firms with tight processing capacities and facing high commodity price volatility.

While single factor, mean-reverting processes are good approximations for modeling commodity price processes, they also have drawbacks. For instance, under a single factor, mean-reverting process, the volatility of futures prices decreases exponentially to zero, as time to maturity increases. To overcome these drawbacks and explain other empirical features of commodity prices, various multi-factor models have also been proposed to model different commodity prices (see Schwartz and Smith, 2000; Geman and Nguyen, 2005, for instance). While these models provide a better description of the price processes, they come with added computational complexity. Computing the optimal policy is difficult because modeling the joint evolution of these multi-factor price processes becomes computationally inefficient (the number of nodes in the lattice increases exponentially with the number of factors). As a result, one has to resort to tractable heuristics. We discuss one particular heuristic to address these computational challenges. Details of the heuristic, along with the computation of an upper bound against which the heuristic can be compared and numerical illustrations can be found in the online appendix.

2.5. Heuristic and Upper Bound for Multi-Factor Models

While the analytical results in Section 2.3 are not dependent on the specific price process, the numerical experiments implemented in Section 2.4 assumed all the commodity prices follow single factor, mean-reverting processes. In this section, we describe a heuristic to compute policy parameters when the output commodity price dynamics follow multi-factor processes. We also describe a computationally tractable upper bound on the optimal expected profits and perform numerical studies to evaluate the performance of the heuristic.

2.5.1 Heuristics for Computation of Marginal Values

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

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

$$\hat{\mathcal{I}}_n = (S_n, F_n^\ell, F_1^{\ell+1}, F_1^{\ell+2}, \ldots, F_1^L) \text{ for } n \text{ such that } N_{\ell-1} \leq n < N_\ell \tag{2.17}$$

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

Next, we approximate the marginal value of output inventory given in equation (2.4) by conditioning the expectations on $\hat{\mathcal{I}}_n$ as follows.

$$\hat{\Delta}_n = \begin{cases} 0 & \text{if } n \geq N_L \\ \beta \max\left\{F_n^\ell, \mathbb{E}_{\hat{\mathcal{I}}_n}\left[\hat{\Delta}_{n+1}\right]\right\} - h_O & \text{if } n = N_\ell - 1 \text{ for } \ell = 1, \ldots, L \\ \beta \mathbb{E}_{\hat{\mathcal{I}}_n}\left[\hat{\Delta}_{n+1}\right] - h_O & \text{otherwise} \end{cases} \tag{2.18}$$

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

$$\hat{\Theta}_n^k = \max\{\hat{\Omega}_n^{(k+b)}, \min\{S_n, \hat{\Omega}_n^{(k)}\}\} \tag{2.19}$$

where

$$\hat{\Omega}_n^{(k)} = \max\left\{\beta \mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Theta}_{n+1}^k] - h_I, \min\left\{\hat{\Delta}_n - p, \beta \mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Theta}_{n+1}^{k-a}] - h_I\right\}\right\} \tag{2.20}$$

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

The heuristic procurement, processing and commitment quantities $(\hat{x}_n, \hat{m}_n, \hat{q}_n)$ are then given by equations (2.15), (2.16) and (2.5) respectively, with the approximate marginal values replacing the true marginal values. Notice that this heuristic requires only modeling the joint evolution of single factor price processes in any given period. Thus, the binomial discretization approaches mentioned in section 2.4.1 can be used to compute the approximate marginal values efficiently. Further, the heuristic is exact in the case when the input and output commodity prices truly follow single factor price processes. We now describe a computationally tractable upper bound on the optimal expected profits, which can be used to evaluate the performance of the heuristic.

2.5.2 Upper Bound on Optimal Expected Profits

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

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

Γ_N, let $H_n^{UB}(e_n, Q_n; \Gamma_N)$ be defined as

$$
\begin{aligned}
H_N^{UB}(e_N, Q_N; \Gamma_N) &= S_N e_N && (2.21)\\
H_n^{UB}(e_n, Q_n; \Gamma_N) &= \max_{q_n, x_n, m_n \in \mathcal{B}_n} \Bigg\{ \Bigg[\beta^{N_\ell - n} F_n^\ell - h_O \sum_{t=0}^{n_\ell - n - 1} \beta^t \Bigg] q_n - p \times m_n \\
&\qquad -S_n \times x_n - h_I e_{n+1} - z_n(e_n, q_n, x_n, m_n, \Gamma_N) \\
&\qquad + \beta H_{n+1}^{UB}(e_{n+1}, Q_{n+1}; \Gamma_N) \Bigg\} \\
&\qquad \text{for } n < N && (2.22)
\end{aligned}
$$

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

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

Notice that H_n^{UB} is the same as V_n given by equations (2.2)–(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

$$
V_1^{UB}(e_1, Q_1, \mathcal{I}_1) = E_{\mathcal{I}_1}[H_1^{UB}(e_1, Q_1; \Gamma_N)] \tag{2.23}
$$

where the expectation is taken over all Γ_N.

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

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

where the marginal values, $\hat{\Delta}_{n+1}$ and $\hat{\Theta}^k_{n+1}$, are given by equations (2.18) and (2.19) and $\hat{\lambda}^k_{n+1}$ are constants such that $\hat{V}_{n+1}$ is continuous in e_{n+1} and $\hat{\lambda}^1_n = 0$ for all n.

We then have

Proposition 2.1. *$V_1^{UB}(e_1, Q_1, \mathcal{I}_1)$ as defined in equation (2.23), with dual penalties given by*

$$\begin{aligned} z_n(e_n, q_n, x_n, m_n, \Gamma_N) &= \beta \Big[\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1}) \\ &\qquad - \mathbb{E}_{\hat{\mathcal{I}}_n}\big[\hat{V}_{n+1}(e_{n+1}, Q_{n+1}, \hat{\mathcal{I}}_{n+1})\big]\Big] \end{aligned} \quad (2.24)$$

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

Proof: The dual penalty in equation (2.24) is a feasible penalty and hence, by Proposition 3.1 in Brown et al. (2010), $V_1^{UB}(e_1, Q_1, \mathcal{I}_1) \geq V_1(e_1, Q_1, \mathcal{I}_1)$. □

Notice that the DP given by (2.22) 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 as a mixed-integer linear program is described in Appendix 2.7.

2.5.3 Numerical Study

We study the performance of the heuristic described in section 2.5.1 by comparing the expected profits using the heuristic with expected profits under the full commitment policy, and the upper bound on optimal expected profits. We consider a single composite output for the purposes of illustrating the heuristic (as seen from previous results, the composite output approximation itself does not lead to significant loss of optimality). The performance of the heuristic is quantified in Section 2.5.3.2. In the next section, we describe the implementation of the heuristics. Our implementation of the heuristic follows Lai et al. (2010a) closely.

2.5.3.1 Implementation

We model the risk-neutral dynamics of the output forward price with maturity at time T_ℓ, $F(t, T_\ell)$, by a driftless geometric Brownian motion and with constant volatility $\sigma_\ell > 0$ as $\frac{dF(t,T_\ell)}{F(t,T_\ell)} = \sigma_\ell dW_\ell(t)$ where $dW_\ell(t)$ is the increment of a standard Brownian motion. The Brownian motion increments corresponding to forward prices with maturities T_ℓ and T_k have a constant correlation coefficient $\rho_{\ell k} \in (-1, 1)$. The Brownian motion increment corresponding to forward price with maturity T_ℓ has a constant correlation coefficient $\rho_{\ell s} \in (-1, 1)$ with the Brownian motion increment corresponding to the input spot price. We continue to model the input spot price as a mean-reverting process, as described in section 2.4.1.

The volatilities of the forward prices and correlation coefficients were estimated using historical data for futures contracts traded on CBOT. These parameters for the four immediately maturing contracts were estimated using the closing futures price on each trading day in the months of June and July, for the years 2001 to 2010. We use only June and July trading dates because futures contracts with maturities every month are not available during other calendar months of the year (soybean meal and oil futures contracts traded on CBOT have maturities in Jan., Mar., May, Jul., Aug., Sept., Oct., and Dec.). The volatilities of the four output forward contracts and the correlation matrix are given in Table 2.9.

The dynamics of the input spot and output forward prices are discretized using the same procedure described in section 2.4.1. We construct recombining binomial trees to

Table 2.9: Multi-factor Price Parameters

(a) Output Forward Price Volatility

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

(b) Correlation Matrix

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

represent the joint evolution of $(S(t), F(t, T_\ell))$ for each $\ell \in \{1, 2, \ldots, L\}$, conditional on F_0^k for $k > \ell$. We also generate binomial trees to represent the evolution of $(F(t, T_\ell), F(t, T_{\ell,\ell+1}))$ for each $\ell \in \{1, 2, \ldots, L-1\}$.

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

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

to approximate the transition probabilities at the expiration of forward contract ℓ.

We compute the heuristic marginal values $\hat{\Delta}_n$ and $\hat{\Theta}_n^k$ for each period n at each node in the binomial tree by using the input spot price value S_n at the node, forward price F_n^ℓ for n such that $n_{\ell-1} \leq n < n_\ell$, and the probability mass functions G_n^ℓ for $n_{\ell-1} \leq n < n_\ell - 1$ and $\hat{G}_{n_\ell - 1}^\ell$ for $n = n_\ell - 1$. The heuristic policy parameters $(\hat{x}_n, \hat{m}_n, \hat{q}_n)$ are computed based on the values of $\hat{\Delta}_{n+1}$ and $\hat{\Theta}_{n+1}^k$ stored at each node. We evaluate the policies using Monte Carlo simulation and compute the expected profits from using the heuristic by averaging the performance on the generated sample paths. We also compute the upper bound for each sample path by solving the mixed-integer program described in Section 2.5.2.

The operational parameters are the same as in section 2.4.1; $p = 72$, the physical holding costs are zero, $\beta = 1$ and the procurement and processing capacities are set to 5 and 3 units.

2.5.3.2 Performance of Heuristic

We consider the procurement, processing and trade operations for the firm over the period June to October. We initialize the input spot price to its long run average value, while the output forward prices are set to the average closing price over June 2010. Table 2.10 gives the optimal expected profits for different horizon lengths when using the full commitment and heuristic policies, along with the gap with respect to the upper bound.

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

As seen in section 2.4.2, the processing capacity and price volatilities had a significant impact on the performance of the full commitment policy. We investigate the impact of varying these parameters on the performance of the heuristic.

Table 2.10: Performance of Heuristic for Different Horizon Lengths

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

Table 2.11: Impact of Processing Capacity
($N = 10$, $L = 2$, $N_\ell = \{5, 10\}$)

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

Impact of processing capacity. As seen from section 2.4.2, the performance of the full commitment policy deteriorates as the processing capacity becomes tight. Table 2.11 shows the expected profits under the full commitment and heuristic policies for different processing capacities. These numerical results are for a horizon length of 10 periods, with two forward contracts available for the output commodity.

For the same procurement capacity, the fraction of the total profits contributed by the output sales are lower for tighter processing capacities. As a result, it is important to extract the full value of the option to postpone commitment. As the results in Table 2.11 show, the heuristic captures this option value and performs well even for tight processing capacities. In comparison, the full commitment policy has a gap as high as 24.26% with respect to the upper bound for very tight capacities.

Table 2.12: Impact of Price Volatility
($N = 10$, $L = 2$, $N_\ell = \{5, 10\}$)

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

Impact of price volatilities. As price volatilities increase, we expect the value of the options inherent in the various decisions to increase. As the results in Table 2.12 show, the value of the upper bound increases more than the expected profits under both the policies as the price volatilities increase. The heuristic is able to capture the option value inherent in the various decisions better than the full commitment policy, as seen from the results in Table 2.12.

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

2.6. Conclusion

In this chapter, we have considered the integrated procurement, processing and trade decisions for a firm dealing in commodities and subject to procurement and processing constraints. We solved the problem optimally and showed that the procurement and processing decisions in any period are governed by price dependent inventory thresholds. We developed recursive expressions to computes these thresholds and illustrate our analytical results using commodity markets data for the soybean complex. Through numerical studies, we find that approximating multiple outputs produced upon processing (e.g., soybean meal and oil) by a single, composite output product is near-optimal. Integrated decision making provides significant benefits compared to a myopic policy which considers only the benefit from processing and trading the output immediately. The value of integrated decision making is especially high under conditions of tight processing capacities and high price volatilities. We also propose a computationally tractable heuristic for computing procurement and processing decisions when commodity prices are driven by multi-factor processes.

The results in the current chapter can easily be extended to incorporate convex, piecewise linear procurement costs and/or concave piecewise linear salvage values for the input inventory. Further, input trade opportunities throughout the horizon can also be easily incorporated into the analysis.

This work lays the foundation for further research in commodity processing and trading operations. Typically, commodity processors operate networks, with procurement and processing activities spread across multiple locations. For instance, the ITC *e-Choupal* network has multiple procurement hubs, along with a few central processing locations. While commodity production and distribution networks have been studied earlier (cf., Markland (1975); Markland and Newett (1976)), these papers assume deterministic commodity prices and no operational constraints. In Chapter 5, we extend the model in the current chapter to a network setting, incorporating stochastic commodity prices and operational constraints. We find that the results developed in the current chapter are useful in developing computationally tractable heuristics for the network problem.

2.7. Appendix: Upper Bound Calculation

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

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

where the dual penalty is given by

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

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

$$\begin{aligned} z_n^k &= \left[\Theta_{n+1}^k(\hat{\mathcal{I}}_{n+1}) - \mathbb{E}_{\hat{\mathcal{I}}_n}[\Theta_{n+1}^k(\hat{\mathcal{I}}_{n+1})]\right] kD + [\hat{\lambda}_{n+1}^k - \mathbb{E}_{\hat{\mathcal{I}}_n}[\lambda_{n+1}^k]] \\ &\qquad \text{for } k=1,2,\ldots,\kappa(n), \text{ for } n=1,2,\ldots,N-1 \\ z_n^0 &= 0 \text{ for } n=1,2,\ldots,N-1 \end{aligned}$$

We can then write the upper bound maximization problem as follows

$$\max \sum_{n=1}^{N-1}\left(\left[\beta^{N_\ell-n}F_n^\ell - h_O\sum_{t=0}^{n_\ell-n-1}\beta^t\right]q_n - pm_n - S_nx_n - h_Ie_{n+1} - \right. \\ \left. \beta\left[\hat{\Delta}_{n+1} - \mathbb{E}_{\hat{\mathcal{I}}_n}[\hat{\Delta}_{n+1}]\right]Q_{n+1} - \beta\sum_{k=1}^{\kappa(n)}[z_n^{k-1}\varphi_n^{(k,l)} + z_n^k\varphi_n^{(k,r)}]\right) + S_Ne_N$$

subject to

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

The above problem can then be solved using a standard mixed-integer programming solver.

Chapter 3

Dynamic Risk Management of Commodity Operations: Model and Analysis

3.1. Introduction

Chapter 2 shows the importance of coordinating procurement, processing and trade decisions in the face of commodity price uncertainty, to maximize expected profits. However, many firms are interested in not just maximizing expected profits but also managing the risk in operations, because an adverse change in commodity prices can greatly impact the firm's profitability and viability. As a result, firms use a variety of operational and financial strategies, including procuring from multiple suppliers and processing to transform the commodities. The option to transform the commodity is especially pertinent to commodity processors, as the output and input commodity price changes are correlated and provide natural hedges. For instance, Yanglin Soybean Inc., a soybean processing firm in China faced increasing soybean prices, but the effect of the increased input prices was offset to some extent by increased soyoil prices, the output produced by the firm (PR Newswire, 2008). With the growth of commodity exchanges and availability of financial derivatives on commodities, firms can now supplement their operational hedging strategies with financial instruments. However, misplaced hedges on commodities using these financial instruments can also have disastrous effects. For instance, VeraSun Energy, a bio-fuel company and one of the biggest producers of ethanol, could not take advantage of falling corn prices because of hedges it had entered into at a time when the corn prices were high. A combination of the high input prices and falling ethanol prices, along with other factors led to the company filing for bankruptcy (Mandaro, 2008).

The examples cited above underscore the need for efficient risk management strategies to mitigate the loss in profitability and financial distress caused by adverse commodity price movements. In fact, managers cite the costs of financial distress and high costs of external financing as some of the key drivers for risk management and hedging (Bickel, 2006). And yet there are no agreed upon practical models for risk management (by a firm) in a dynamic context of operational decision making. One approach to model risk aversion treats firms as an entity or individual to whom a utility function can be attributed. However, a firm's risk preferences are never defined in terms of a corporate utility function and thus, in a normative sense, there is no 'right' corporate utility function to use. An alternate approach to model the risk management problem for firms is to use mean-risk objective functions that represent a tradeoff between the expected value and the cost associated with uncertainty or variability of cash flows. Perhaps the most well known instance of such objective functions is the mean-variance objective function, where risk is measured by the variance in profits. Another risk measure to quantify risk is the value at risk (VaR), which is popular in the financial industry. However, measures such as variance and VaR have drawbacks (e.g., variance penalizes both under and over performance, while VaR is not sub-additive). Coherent risk measures, such as conditional value at risk (CVaR), have been proposed that overcome these drawbacks.

In a multi-period setting, efficient risk management requires controlling risk in intermediate periods in addition to controlling the risk in total payoffs at the end of the planning horizon. For instance, grain elevators use futures contracts to reduce the risk of falling prices at the time of delivery. However, the increased volatility and near doubling of corn and soybean prices in 2008 led to significant margin calls on these futures positions and put severe pressure on the working capital requirements for many grain elevators (Serres, 2008). Mean-risk objective functions on terminal wealth that ensure the risk in final payoffs are minimized, but do not account for adverse outcomes in intermediate periods in a consistent manner, may not be truly effective when modeling risk averse decisions in a multi-period setting. Such objective functions do not incorporate the timing of uncertainty resolution over the horizon in a consistent manner and can lead to situations where significantly more risky positions are taken in intermediate periods. This aspect of dynamic risk management has not received much attention in the OM literature and we specifically focus on the issue of time consistency in multi-period risk averse decision making in this essay. Specifically, we model the risk averse firm's objective by extending the single period coherent risk measure CVaR to incorporate uncertainty resolution in a consistent manner and control risk over the entire horizon, and not just in the terminal wealth.

We model a commodity processing firm that procures an input commodity from the spot market and processes the commodity to produce an output commodity. At the beginning of each period, the firm decides how much input commodity to procure from the spot market and how much of the total available input to process, subject to capacity constraints on procurement and processing. The firm uses forward contracts to sell the output commodity and in each period, given the current forward prices, the firm decides the quantity to commit for sale. All forward sale commitments are delivered at the maturity of the forward contract. In addition to operational decisions, the firm also makes financial trading decisions using derivative instruments on the input commodity to manage the commodity risk.

We model the firm's multi-period risk management problem using a proposed time-consistent risk measure based on the conditional value at risk (CVaR). Broadly speaking, CVaR measures the conditional expectation of the worst case profits, where worst case profits are profits in the left tail of the distribution. However, using CVaR of net present value of total profits does not lead to consistent decisions over time;

i.e., optimal policies in the current period may not appear optimal when evaluated in earlier periods (see Section 3.3 for an example). To overcome this problem and ensure that the firm's risk preferences are represented in a consistent manner over the planning horizon, we extend CVaR to dynamic settings using a recursive definition (see Section 3.4 for details of our objective function). The new objective function is time-consistent in the sense that an optimal policy in the current period will also be optimal when evaluated in earlier periods. Using the time-consistent objective function, we characterize the structure of the optimal operational and financial hedging policy in the presence of capacity constraints and inventory. To summarize our contributions and results,

1. We propose a model to capture risk aversion in a multi-period setting that is time-consistent.

    (a) Specifically, we propose a time-consistent variation of the conditional value at risk measure to model risk aversion in a multi-period setting. While the CVaR measure provides a specific illustration, our approach can be used more generally to extend other coherent risk measures in a time-consistent manner.

    (b) Using numerical studies, we find that a time-consistent risk measure dominates static risk measures defined on the total profits in a mean-CVaR sense. That is, the mean-CVaR tradeoff achieved using policies for optimizing a time-consistent objective function outperforms the mean-CVaR tradeoff achieved using policies for optimizing a static risk measure on the total profits over the horizon.

2. For the proposed time-consistent risk measure, we obtain the optimal operational and financial risk management policy. We show that

    (a) The optimal procurement and processing decisions are characterized by price dependent ‘procure up to’ and ‘process down to’ inventory thresholds. These thresholds are however hard to compute because of high dimensionality of the state space. We develop efficient heuristics to approximate the target inventory levels. Using numerical results, we demonstrate near-optimality of the heuristics.

(b) The optimal financial hedging decisions are a function of the ending input inventory levels. Thus, integration of financial and operational decisions is crucial for effective risk management.

(c) Conditional on the optimal financial trading decisions, the optimal operational decisions can be obtained as the solution to a dynamic programming equation that does not involve the financial decisions.

The remainder of the chapter is organized as follows. In section 3.2, we review the relevant OM literature. Section 3.3 provides an overview of different approaches to model risk averse decision making and describes the issue of time consistency in multi-period settings. Our time-consistent objective function based on CVaR is presented in section 3.4, while section 3.5 describes the model and analysis of the multi-period integrated risk management problem for a commodity processor. Section 3.6 illustrates the analytical results using numerical experiments and we conclude in section 3.7.

3.2. Previous Work

This essay spans two areas of literature. The first concerns integrated operational and financial risk management. The second pertains to models for risk aversion; we postpone its discussion to the next section, including identifying issues that arise in a dynamic context. A majority of the existing literature related to frameworks of integrated operational and financial risk management (Kleindorfer, 2008) deals with integrating long term contracts with short term spot market procurement to manage price and demand uncertainty in single period contexts. Kleindorfer and Wu (2003) review the literature on integrating B2B commodity exchanges in procurement contracts. Wu and Kleindorfer (2005) and Martínez-de Albéniz and Simchi-Levi (2006) consider the optimal portfolio of contracts for a buyer who can sign option/capacity reservation contracts with multiple suppliers and also has access to a spot market. In contrast to these papers which consider single period models, we study dynamic decision making in a multi-period setting. Further, risk management which is an important aspect of decision making in the current work, is not dealt with in the above literature.

The research on multi-period integrated operational and financial hedging is fairly limited and deals primarily with non-storable commodities. Since decisions across periods are not linked by inventory, the problem can be decomposed into single period problems thus simplifying the analysis. For instance, Kleindorfer and Li (2005) consider the multi-period risk management problem for a electricity generator using a mean-risk objective, where the risk measure used is the value at risk (VaR). Broadly speaking, VaR is the level of losses that is exceeded only with a small probability, e.g., 1% (or alternately, the value below which profits fall only with a small probability). By decomposing the problem, Kleindorfer and Li show the equivalence of the mean-VaR and mean-variance frontiers. Further, they solve the problem in an open-loop manner, where all decisions are made at the beginning of the horizon. In a dynamic decision making context, Zhu and Kapuscinski (2006) model operational and financial hedging for a risk-averse multinational newsvendor exposed to exchange rate risk, with the newsvendor maximizing an additive exponential utility function on consumption streams. Notably, Zhu and Kapuscinski (2006) do not allow inventory to be carried across periods, with only the financial decisions affecting income across periods.

The papers by Kouvelis et al. (2009) and Geman and Ohana (2008) are the closest to ours in terms of modeling the integrated risk management problem for a storable commodity in a multi-period context. Kouvelis et al. model the dynamic procurement and financial trading decisions for a firm managing a storable commodity used to satisfy uncertain demand in each period. A critical difference between our work and Kouvelis et al. (2009) is in terms of how the firm's objective function is modeled. They use a mean-variance objective function, where the objective at the start of each period is the expected profits over the remaining periods minus a multiple of the sum of variances of profits from the current and future periods. This objective function is not necessarily time-consistent and indeed, the authors mention that their choice of the risk averse objective function can lead to inconsistent decisions over time and as a result they cannot use dynamic programming techniques to solve the multi-period problem (see Kouvelis et al., 2009, pg. 12–13). In contrast, we use a time-consistent risk measure to model the risk averse firm's objective function. Fundamentally, time consistency ensures that the objective function represents the firm's underlying preferences in a consistent manner over time and leads to consistent decision making, thus enabling the use of dynamic programming techniques to solve the multi-period risk management problem. While our modeling of the objective function is similar to Geman and Ohana (2008), they only model trading decisions for a single commodity and their model does not include decisions to irreversibly transform some of the commodities through processing. They do not consider capacity constraints that play a central role in our model and analysis. Finally, in contrast to our work, Geman and Ohana (2008) neither provide any results for the optimal policy structure nor do they study the interaction of operational and financial decisions.

3.3. Modeling Risk Averse Decision Making

Use of utility functions to model risk averse preferences with expected utility as a criteria for decision making is a common approach in economics. The OM literature has often followed a similar approach; see for instance Eeckhoudt et al. (1995); Agrawal and Seshadri (2000); Bouakiz and Sobel (1992). This approach treats firms as an entity or individual to whom a utility function can be attributed. Smith (2004) notes that some of the suggested approaches to model a corporate utility function include using the individual unit manager / decision maker's preferences or have the corporate leaders define a utility function for the firm as a whole, which the employees then adopt. Finance theorists argue that a firm's management should maximize the benefits of the shareholders of the firm, thereby suggesting that the firm should use the shareholders' utility function. However, it is unlikely that even when acting to benefit the shareholders there will be unanimous agreement between all shareholders and maximizing shareholders' benefit can at best place limits on the risk preferences that inform the firm's decision, but not yield a utility function for the firm itself (Smith, 2004).

Corporate finance also argues that firms should only care about systematic risk, because shareholders can diversify away any unsystematic risk (Brealey and Myers, 2003). However, managers do exhibit risk averse behavior and cite risk management as a priority (Walls and Dyer, 1996; Bickel, 2006). The need for risk management in corporations result from costs of financial distress, differences between cost of external and internal financing and principal-agent problems between shareholders and management (Bickel, 2006). An approach that is more closely aligned with the risk management motives of avoiding or minimizing the costs of financial distress and external financing is the use of risk measures or mean-risk models to capture the tradeoff between expected value and risk and has its roots in the mean-variance framework for optimal portfolio selection (Markowitz, 1952).

Since Markowitz (1952) several refinements have been proposed to capture the mean-risk tradeoff more efficiently and overcome the drawbacks associated with the mean-variance approach (it penalizes over- and under-performance equally and excludes stochastically dominant portfolios, for instance). In a seminal paper, Artzner et al. (1999) proposed the theory of coherent risk measures and introduced mean-risk functionals that satisfy properties desirable from the perspective of intuitive decision making. Specifically, coherent risk measures satisfy a) monotonicity, b) sub-additivity, c) positive homogeneity and d) translation invariance. Some of the popular risk measures do not necessarily satisfy all the properties required of coherent risk measures, e.g., VaR violates sub-additivity and variance in a mean-variance criterion violates monotonicity.

The monotonicity and sub-additivity properties of coherent risk measures ensure that using coherent risk measures as a risk-averse objective function is consistent with the use of increasing, concave utility functions under the expected utility approach. In addition, coherent risk measures are also translation invariant and positive homogenous, properties that are not necessarily true for utility functions. Translation invariance implies that initial endowments do not alter risk preferences and hence coherent risk measures provide a good way to quantify the risk inherent in a specific project or the risk from specific decisions. Positive homogeneity ensures that adding same risks together do not reduce the total risk and neither do changing the units of measurement. For these reasons, coherent risk measures suggest themselves as desirable criteria for modeling risk averse decision making at the firm level. A popular coherent risk measure that has been used in finance is the conditional value at risk (CVaR), a modification of VaR (e.g., Rockafellar and Uryasev, 2000). CVaR is a downside risk measure that captures payoffs in the worst case scenarios. Generally speaking, CVaR is equal to the conditional expectation of losses above the VaR; e.g., conditional expectation of highest 1% of losses (alternately, lowest 1% of profits). Managers often express risk as only those events associated with negative outcomes (March and Shapira, 1987) and a downside risk measure such as CVaR captures these preferences. Increasingly, CVaR has also been used in the OM literature to model risk averse decision making in single period contexts (e.g., Gotoh and Takano (2007); Chen et al. (2008); Choi et al. (2009) for newsvendor problems, Tomlin and Wang

(2005); Tomlin (2006) in the context of flexibility and supply disruptions, etc.)

Whether using utility functions or risk measures to capture risk aversion, dynamic risk averse decision making also requires incorporating preferences over time. It is known that using expected utility over net present value of wealth leads to a temporal risk problem; i.e., it does not differentiate between alternatives that have the same distribution of terminal wealth, but different times for uncertainty resolution (Smith, 1998). Presumably, a risk-averse decision maker is sensitive to the timing of uncertainty resolution, and not just the distribution of total wealth. One approach used to overcome the temporal risk problem in the utility framework is to model utility over consumption streams rather than incomes and allow borrowing and lending to smooth consumption across time periods. This is the approach followed for instance in Chen et al. (2007) to model dynamic, risk averse decision making. Such an approach requires optimization of consumption, borrowing and lending decisions in addition to the operational decisions and other than for specific utility functions, e.g., additive exponential utility, it is hard to solve the problem and gather insights.

Using static risk measures, i.e., single period risk measures, on the net present value of wealth suffers from a related problem of time inconsistency. That is, the preferences represented by the static measure on the NPV is not consistent over time. To illustrate this, consider the following example adapted from Roorda and Schumacher (2007).

Example 3.1. *An operational investment* **A**, *e.g., inventory investment, pays off at the end of two periods, with payoffs as shown in Figure 3.1. The payoffs are dependent on the state of the world in periods 1 and 2, and the probabilities for the different states of the world are as shown on the branches of the tree.*

Consider a risk-averse decision maker who needs to make this investment and whose risk aversion can be represented by CVaR at the $\eta = \frac{2}{3}$ level on the final payoffs. Broadly speaking, $CVaR$ at the $\frac{2}{3}$ level represents the conditional expectation of the

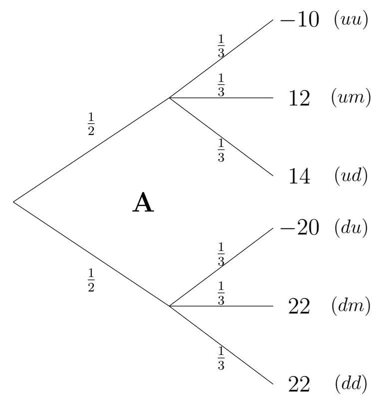

Figure 3.1: Two period investment **A**

payoffs in the worst $\frac{2}{3}$ states of the world. For investment **A**, the payoffs in the worst $\frac{2}{3}$ states of the world as evaluated at $t = 0$ are $(-20, -10, 12, 14)$, each with an unconditional probability of $1/2 \times 1/3 = 1/6$. Thus, we have $CVaR$ at $\eta = \frac{2}{3}$ equal to $\frac{1}{2/3} \times \frac{(-20-10+12+14)}{6} = -1$. We can also evaluate $CVaR$ for **A** at $t = 1$, conditional on the state realized at $t = 1$, in a similar manner. Table 3.1 summarizes the $CVaR$ of **A** at $\eta = \frac{2}{3}$, evaluated at $t = 0$ and $t = 1$. With the default option of doing nothing

Table 3.1: CVaR calculations for **A**

	$CVaR$ at $\eta = \frac{2}{3}$
$t = 0$	-1
$t = 1$, u	1
$t = 1$, d	1

(which has a $CVaR = 0$), the decision maker would undertake the investment only if the $CVaR$ value is greater than 0. The CVaR in period 0 is equal to -1. Thus, **A** will be deemed unacceptable in period 0. Notice however that the CVaR in all possible states of the world in period 1 is equal to 1. Thus, the same investment would have been considered acceptable in *all* states of the world in period 1. Notice that there are no cash flows occurring between periods 0 and 1 and there is no discounting. Thus, using a static risk measure such as CVaR can lead to inconsistent decision making.

3.4. A Time Consistent Objective Function

The problem of time inconsistency has been considered in the area of mathematical finance and different notions of time consistency that a dynamic risk measure should satisfy have been suggested (Wang (1999), Detlefsen and Scandolo (2005), Artzner et al. (2007), Roorda and Schumacher (2007), Roorda et al. (2005) , Geman and Ohana (2008) etc.). The basic idea underlying these notions is that the dynamic risk measure should indicate the underlying preferences in a consistent manner over time. That is, if the risk measure suggests that a particular stream of cashflows is preferred in all possible states of the world in the next period, then the risk measure should also imply the same preference in the current period as well, all other things being the same. Following these ideas, we propose a dynamic, time-consistent risk measure based on CVaR, to represent risk preferences in a multi-period problem.

Formally, $VaR^{\eta}(X)$, the VaR of a stochastic payoff X at a particular probability level η is the value below which profits fall only with a probability of η (Duffie and Pan, 1997). $CVaR^{\eta}(X)$ of a stochastic payoff X is the expectation of X, conditional on $X \leq VaR^{\eta}(X)$. While this definition is exact when the random variable X has a continuous distribution and no atoms at $VaR^{\eta}(X)$, in the general case $CVaR^{\eta}(X)$ is given by (Rockafellar and Uryasev, 2000)

$$CVaR^{\eta}(X) = \max_{\upsilon} \left\{ \upsilon - \frac{1}{\eta} \mathbb{E} \left[(\upsilon - X)^{+} \right] \right\} \tag{3.1}$$

where $\mathbb{E}[\cdot]$ is the expectation (the expression for CVaR above considers X to denote profits or gains, while the definition in Rockafellar and Uryasev (2000) is based on X denoting the losses). The maximizer υ^{*} in the above problem is unique and equal to $VaR^{\eta}(X)$ if the distribution of X has no atoms. Otherwise, the set of optimizers is not necessarily unique, with $VaR^{\eta}(X)$ being one of the optimizers. Of course, $CVaR^{\eta}(X)$ is still uniquely defined.

In the definition of $CVaR^{\eta}(X)$, the probability level η represents the degree of risk aversion, with risk aversion decreasing as η increases. For $\eta = 1$, we have $CVaR^{1}(X) = \mathbb{E}[X]$, representing risk neutrality. The stochastic payoff X can represent a single period payoff or the cumulative payoff over multiple periods. In the latter case, $CVaR^{\eta}(X)$ represents a risk measure on the total cumulative payoff, evaluated based on the current information available.

A dynamic risk measure $\rho = (\rho_1, \ldots, \rho_N)$ is a sequence of risk measures on a cashflow stream $\mathbf{A} = (\mathcal{A}_1(\mathcal{S}_1), \ldots, \mathcal{A}_N(\mathcal{S}_N))$ such that $\rho_n(\mathbf{A}, \mathcal{S}_n)$ measures the risk associated with $\mathbf{A}$, given the state $\mathcal{S}_n$ in period n. We follow Geman and Ohana (2008) to define a time-consistent dynamic risk measure as follows

Definition 3.1. *(Geman and Ohana, 2008) A dynamic risk measure* $\rho = (\rho_1, \rho_2, \ldots, \rho_N)$ *is intrinsically time-consistent if for any two cash flow streams* $\mathbf{A} = (\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_N)$ *and* $\mathbf{B} = (\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_N)$*, for all* $\mathcal{S}_n$

$$\begin{cases} \mathcal{A}_n(\mathcal{S}_n) \geq \mathcal{B}_n(\mathcal{S}_n) \\ \\ \rho_{n+1}(\mathbf{A}, \mathcal{S}_{n+1}) \geq \rho_{n+1}(\mathbf{B}, \mathcal{S}_{n+1}) \quad \forall \mathcal{S}_{n+1} \in \mathcal{H}_{n+1}(\mathcal{S}_n) \end{cases} \Rightarrow \rho_n(\mathbf{A}, \mathcal{S}_n) \geq \rho_n(\mathbf{B}, \mathcal{S}_n)$$

where $\mathcal{S}_n$ *is the state in period* n *and* $\mathcal{H}_{n+1}(\mathcal{S}_n)$ *is the set of all possible states in period* $n+1$ *given that the state in period* n *is* $\mathcal{S}_n$.

The above definition implies that if a particular cash flow stream, $\mathbf{A}$, is preferred over another, $\mathbf{B}$, in all possible states of the world in the next period, and the cash flow in the current period under $\mathbf{A}$ is at least as much as the cash flow under $\mathbf{B}$, then $\mathbf{A}$ should be preferred over $\mathbf{B}$ in the current period also. In Example 3.1, we had $\rho = (\rho_0, \rho_1, \rho_2)$, where $\rho_n(\mathbf{A}, \mathcal{S}_n) = CVaR^{(\frac{2}{3})}\left(\sum_{t=n}^{N} \mathcal{A}_t(\mathcal{S}_t) \Big| \mathcal{S}_n\right)$. If we take $\mathbf{B}$ as the cashflow stream that has zero payoff in all periods, in all states of the world, we saw that $\rho_0(\mathbf{A}, \mathcal{S}_0) < \rho_0(\mathbf{B}, \mathcal{S}_0)$, while $\rho_1(\mathbf{A}, \mathcal{S}_1) > \rho_1(\mathbf{B}, \mathcal{S}_1)$ for all $\mathcal{S}_1 \in \mathcal{H}_1(\mathcal{S}_0)$. Thus, $CVaR$ of total cashflow over the remaining horizon is not necessarily a time-consistent risk measure.

We now define a dynamic version of the conditional value at risk, $DCVaR(\cdot\ ;\eta) = (DCVaR_1(\cdot\ ;\eta_1),\ldots,DCVaR_N(\cdot\ ;\eta_N))$ as follows

$$\begin{aligned} DCVaR_N(\mathbf{A},\mathcal{S}_N\ ;\eta_n) &= \mathcal{A}_N(\mathcal{S}_N) && (3.2)\\ DCVaR_n(\mathbf{A},\mathcal{S}_n\ ;\eta_n) &= CVaR^{\eta_n}\Bigg(\mathcal{A}_n(\mathcal{S}_n) \\ &\quad + DCVaR_{n+1}\Big(\sum_{t=n+1}^{N}\mathcal{A}_t(\mathcal{S}_t),\mathcal{S}_{n+1};\eta_{n+1}\Big)\Big|\ \mathcal{S}_n\Bigg) && (3.3)\\ &\qquad \text{for } n < N \end{aligned}$$

In the above definition, the mean-risk tradeoff evaluated in period n using $CVaR$ is based on the sum of cashflow in the current period $\mathcal{A}_n$ and the $DCVaR_{n+1}$ of the cash flow over the remaining periods. Thus, it is the 'risk adjusted' value of the future period cash flows, rather than just the value of the future cash flows, that is taken into account when evaluating the risk over the remaining horizon. As Proposition 3.1 states, $DCVaR(\cdot\ \eta)$ is a time-consistent risk measure (The proofs for all results are provided in the Appendix).

Proposition 3.1. *According to Definition 3.1, the dynamic risk measure given by equation (3.3) is time-consistent.*

Before continuing, we briefly re-visit Example 3.1 to illustrate the time-consistent measure. We have $\mathcal{A}_0 = 0$ and $\mathcal{A}_1(\mathcal{S}_1) = 0$ for $\mathcal{S}_1 = u, d$. We have,

$$DCVaR_1(\mathbf{A},\mathcal{S}_1; 2/3) = CVaR^{(2/3)}(\mathcal{A}_2(\mathcal{S}_2)|\mathcal{S}_1) = 1$$

for $\mathcal{S}_1 = u, d$. Also,

$$DCVaR_0(\mathbf{A},\mathcal{S}_0; 2/3) = CVaR^{(2/3)}(DCVaR_1(\mathbf{A},\mathcal{S}_1)|\mathcal{S}_0) = 1.$$

Notice that the $DCVaR(\cdot)$ measure accounts for the uncertainty resolved in period 1 in a consistent manner, resulting in $DCVaR_0(\mathbf{A}, \mathcal{S}_0; 2/3) = 1$. Thus, using a time-consistent risk measure such as $DCVaR$ will lead to no inconsistency in decision making since the $DCVaR$ in period 0 is equal to 1, thus indicating that $\mathbf{A}$ is deemed acceptable in period 0 also. For the integrated operational and financial risk management problem considered here, we use $DCVaR$, the time-consistent risk measure defined in equation (3.3) to model the firm's objective.

3.5. Model Description and Analysis

We consider the integrated operational and financial risk management problem for a commodity processor who procures, processes and trades commodities over a finite horizon. The time periods are indexed by $n = 1, \ldots, N$ with $n = 1$ denoting the first decision period. In each period n, the firm procures the input commodity from a spot market, where S_n denotes the spot price in period n. The firm earns revenues by processing the input and selling the output commodity (processed product) using a forward contract, with the forward price in period n given by F_n. We assume the delivery period for the forward contract is period N and period $N-1$ is the last period in which the firm can commit to sell the output commodity using the forward contract. In addition to the output commodity sales, the firm can also earn revenues by trading the input commodity with other processors. For ease of exposition, we assume that all input commodity trading occurs at the end of the horizon, at the trade (salvage) price of S_N. Let $\mathcal{I}_n$ denote the relevant information available to the firm at the beginning of period n regarding the various commodity prices.

On the operational side, the firm has a per-period procurement and processing capacity restriction of K and C units respectively. The firm incurs a variable cost of p to process one unit of input into the output commodity. For simplicity, we assume all physical holding costs for the various commodities are negligible. We first consider the situation when the firm uses only operational decisions, without any trading in the financial markets, to manage the risk. Later, in Section 3.5.3, we analyze the role of financial hedging decisions for risk management.

3.5.1 Operational Hedging

In the current section, we focus on the optimal procurement, processing and physical commodity trade decisions to manage risk in the commodity processing operations. The output sale commitments require physical delivery of the output commodity on the delivery date and are hence included in the operational decisions.

At the beginning of each period n, the firm observes the input spot price, S_n, and the output forward price, F_n, for the period. Based on the input commodity inventory, e_n, and the uncommitted output commodity inventory, Q_n, the firm makes the following decisions: 1) The quantity of input commodity to procure, x_n, 2) the quantity to process, m_n and 3) the quantity of the output commodity to commit to sale against the forward contract, q_n.

The uncommitted output inventory refers to the total output inventory that is in excess (or shortfall) of the total commitments made till the beginning of period n: i.e., $Q_n = \sum_{t=1}^{n-1} m_t - \sum_{t=1}^{n-1} q_t$. It is not necessary that $Q_n \geq 0$, as the firm can commit to sell more output than is available on hand as long as all the output committed for sale against the forward contract is delivered on the delivery date specified in the forward contract, which in our case is period N. Thus, it is not necessary that $Q_n \geq 0$ for all $n < N$. However, all commitments made over the horizon have to be satisfied and hence we require $Q_N \geq 0$. Further, we assume commitments once made cannot be reneged on, i.e., $q_n \geq 0$ for all n.

The procurement and processing decisions in any period are subject to capacity and inventory availability constraints and the feasible set of actions in period n is given by $\mathcal{A}_n(e_n)$ where

$$\mathcal{A}_n(e_n) = \{(x_n, m_n) : 0 \leq x_n \leq K,\ 0 \leq m_n \leq \min\{C, e_n + x_n\}\} \tag{3.4}$$

and the state transitions are given by $e_{n+1} = e_n + x_n - m_n$ and $Q_{n+1} = Q_n + m_n - q_n$.

The profits realized by the firm in period n, for $n \leq N-1$, are given by

$$\Pi_n(x_n, m_n, q_n, \mathcal{I}_n) = \beta^{N-n} F_n q_n - S_n x_n - p m_n \tag{3.5}$$

where β is the discount factor. In the final period, we have

$$\Pi_N(e_N, Q_N, \mathcal{I}_N) = \begin{cases} S_N e_N & \text{if } Q_N \geq 0 \\ -\infty & \text{otherwise} \end{cases} \tag{3.6}$$

The profit function in equation (3.5) above recognizes revenues from output sales at the time of commitment rather than at delivery. Since commodity sale commitments are not reversible and we assume no counter party risk is present, recognizing revenue at the time of commitment rather than at delivery is without loss of generality. As our focus in this research is on managing the commodity price risk, we assume no counter-party risk associated with the buyer of the output commodity. Notice that the profit function for the last period given by equation (3.6) accounts for the fact that all output sale commitments have to be met on the delivery date for the forward contract.

We model the risk-averse firm's objective function by the time-consistent risk measure $DCVaR(\cdot;\eta)$, defined in equation (3.3), on the stream of discounted cashflows. To keep the exposition simple, we assume that the probability levels η are the same across all periods. The firm's risk management problem in period n can then be written as

$$V_n(e_n, Q_n, \mathcal{I}_n) = \max_{\substack{((x_t,m_t)\in\mathcal{A}_t(e_t),\ q_t\geq 0),\\ t=n,\ldots,N-1}} \left\{ DCVaR_n^{\eta}\left(\sum_{t=n}^{N-1} \beta^{t-n}\Pi_t(x_t, m_t, q_t, \mathcal{I}_t) + \beta^{N-n}\Pi_N(e_N, Q_N, \mathcal{I}_N) \right) \right\} \tag{3.7}$$

where we have used the short hand $DCVaR_n^{\eta}(\cdot) \triangleq DCVaR_n(\cdot, \mathcal{S}_n; \eta_n)$ to keep the notation simple.

The fact that $DCVaR_n^{\eta}$ is a time-consistent risk measure ensures that the maximization problem in equation (3.7) can be solved as a stochastic dynamic program (SDP) as stated below.

Theorem 3.1. *The optimal $DCVaR_n^{\eta}$ of profits from period n till end of the horizon, $V_n(e_n, Q_n, \mathcal{I}_n)$ in equation (3.7), is given by the solution of the following stochastic dynamic program*

$$V_n(e_n, Q_n, \mathcal{I}_n) = \max_{(x_n, m_n) \in \mathcal{A}_n(e_n, Q_n),\ q_n \geq 0} \Big\{ \Pi_n(x_n, m_n, q_n, \mathcal{I}_n) + \beta CVaR_n^{\eta}\Big(V_{n+1}(e_{n+1}, Q_{n+1}, \mathcal{I}_{n+1})\Big) \Big\} \quad (3.8)$$

for $n = 1, \ldots, N-1$ and

$$V_N(e_N, Q_N, \mathcal{I}_N) = \begin{cases} S_N e_N & \text{if } Q_N \geq 0 \\ -\infty & Q_N < 0 \end{cases} \quad (3.9)$$

Notice that the objective function in equation (3.8) is very similar to the objective function in SDP formulations for expected value maximization. However, there are crucial differences. In equation (3.8), V_{n+1} is the optimal value of $DCVaR_{n+1}$ of cash flows over periods $n+1, \ldots, N$, conditional on $\mathcal{I}_{n+}$ and the starting inventory levels. Thus, V_{n+1} incorporates the firm's risk aversion over future cash flows and can be thought of as a 'risk adjusted' value of future period cash flows. Evaluated in period n, the 'risk adjusted' value of cash flows over periods $n, \ldots, N$ is itself uncertain. The optimization problem incorporates the firm's risk aversion in period n over the 'risk adjusted' value of future cash flows through the $CVaR_n^{\eta}$ term. As mentioned earlier, varying the value of η varies the level of risk aversion and for $\eta = 1$, the above problem is the same as expected value maximization.

While equation (3.8) enables us to use the tools developed for SDP in the expected value maximization context to analyze and solve the risk-averse problem, the analysis is significantly more complex because $CVaR_n^\eta$ involves only the left tail of the distribution of cash flows. Nevertheless, we can establish a basic concavity property of the value function. Concavity implies that there exists an optimal decision, (x_n^*, m_n^*, q_n^*), for each starting state $(e_n, Q_n, \mathcal{I}_n)$, in each period n.

Lemma 3.1. *For all $n \le N$, the value function $V_n(e_n, Q_n, \mathcal{I}_n)$ is concave in (e_n, Q_n) for each $\mathcal{I}_n$ and increasing in e_n and Q_n.*

In computations, it is often useful to represent the evolution of the prices as a discrete process, e.g., as lattice models. Such discrete price models have been used for instance by Secomandi (2010b) and Jaillet et al. (2004) in the context of risk-neutral commodity trading and option valuation problems. Further, the qualitative insights obtained by considering discrete price processes will hold for the general processes as well. To gain more insights about the dynamic risk management decisions, we restrict attention to the situation where the price processes follow a discrete distribution and the following assumption holds for the remainder of this chapter.

Assumption 3.1. *(Finite price sets). In each period n, the set of possible commodity prices is finite.*

The above assumption implies that in each period n, the random variable $\mathcal{I}_{n+1}$, conditional on $\mathcal{I}_n$, has a discrete probability distribution for all $n < N$. Let $\mathcal{M}_n = \left(\mathcal{I}_n^1, \ldots, \mathcal{I}_n^{M_n}\right)$ denote the set of possible realizations of $\mathcal{I}_n$ and $M_n = |\mathcal{M}_n|$ be the number of possible realizations for $n = 1, 2, \ldots, N-1, N$. We use the shorthand notation $V_n^m(e_n, Q_n) \triangleq V_n(e_n, Q_n, \mathcal{I}_n^m)$ to denote the value function for a specific realization of $\mathcal{I}_n^m \in \mathcal{M}_n$, while retaining $V_n(\cdot, \cdot, \mathcal{I}_n)$ to denote the value function for a general $\mathcal{I}_n$. Also, let $CVaR_n^m(\cdot)$ denote $CVaR_n(\cdot|\mathcal{I}_n^m)$ (we have suppressed η, the level at which the CVaR is calculated, for ease of notation).

When prices belong to a finite set, equation (3.6) can be modified to

$$\Pi_N(e_N, Q_N, \mathcal{I}_N) = \begin{cases} S_N e_N & \text{if } Q_N \geq 0 \\ \Gamma Q_N & \text{if } Q_N < 0 \end{cases}$$

where $\Gamma > 0$ is sufficiently large to ensure that $Q_N \geq 0$ under an optimal policy. For instance, we could choose Γ to be a value larger than the maximum possible output forward price across all periods. As a result of this modification, equation (3.9) becomes

$$V_N(e_N, Q_N, \mathcal{I}_N) = \begin{cases} S_N e_N & \text{if } Q_N \geq 0 \\ \Gamma Q_N & Q_N < 0 \end{cases} \tag{3.10}$$

in the case of discrete prices (there is no change to equation (3.8)). Theorem 3.2 establishes an additional property of the value function under Assumption 3.1.

Theorem 3.2. *Suppose Assumption 3.1 holds. Then, the value function* $V_n^l(e_n, Q_n)$ *is piecewise linear and continuous in* (e_n, Q_n) *for each* $\mathcal{I}_n^l \in \mathcal{M}_n$.

Consider the commitment decision, q_n. Recall that the firm can potentially over-commit; i.e., have $q_n > Q_n + m_n$ in any period $n < N - 1$ as long as it ensures that $Q_N \geq 0$, as implied by equation (3.10). The next lemma proves that under mild restrictions on the output forward price process, it is never optimal to commit more than the current output inventory.

Lemma 3.2. *If the output forward prices are unbiased, i.e.,* $\mathbb{E}_{\mathcal{I}_n}[F_{n+1}] = F_n$, *then the optimal commitment quantity in any period is limited by the available uncommitted output inventory. That is,* $q_n^* \leq [Q_n + m_n^*]^+$.

The firm has to satisfy all output commitments before the end of the planning horizon. Thus, any over-commitment made in period n has to be satisfied by processing (and possibly procuring,) additional input in the future periods. Further, any over-commitment in the current period also means a forgone commitment opportunity in the future. Lemma 3.2 states that the $DCVaR$ of satisfying the over-commitment in terms of the cost of meeting the commitment and/or the forgone future commitment opportunities is higher than the forward price earned by committing the extra unit. Under an optimal policy, we will never have a situation where $Q_n < 0$ for any n.

In general, the optimal procurement, processing and commitment decisions will depend on each other and the starting inventory levels in a non-trivial fashion. For instance, concavity of $CVaR_n^l(V_{n+1})$ in Q_{n+1} implies that there exists a target output inventory level $\underline{Q}_{n+1}^l \geq 0$ such that $q_n^* = (Q_n + m_n - \underline{Q}_{n+1}^l)^+$ for a given realization $\mathcal{I}_n^l$. However, this target output inventory level is not necessarily independent of e_{n+1}. The manner in which $\underline{Q}_{n+1}^l$ changes with e_{n+1} will depend on the joint evolution of the input and output price processes. If the commodity prices satisfy the conditions of Assumption 3.2 below, we can obtain more insights into the optimal decisions.

Assumption 3.2. *For every period $n < N$, (a) the distribution of F_{n+1} conditional on the forward price F_n realized in period n stochastically increases in F_n; (b) the output forward prices are unbiased, i.e., $\mathbb{E}_n[F_{n+1}] = F_n$.*

The two conditions stated in Assumption 3.2 are quite natural. In particular, part (a) of the assumption implies that the expected output forward price in the next period is increasing in the output forward price realized in the current period. Well known models of commodity prices such as the mean-reverting model or geometric Brownian motion models satisfy this condition. The second condition implies that the firm has no speculative motive to hold the output inventory. Under these conditions, Lemma 3.3 shows that it is optimal for the firm to commit all available output inventory in any period.

Lemma 3.3. *If the output forward prices satisfy the conditions in Assumption 3.2, then it is optimal to commit all the available output inventory. That is, $q_n^* = Q_n + m_n^*$*

is an optimal solution.

The above result implies that there always exists an optimal policy where all available output inventory is committed. Under such a policy, we will always have $Q_n = 0$ and further, $q_n = m_n$ for all n. Restricting our attention to only such policies, we do not need to consider the commitment and processing decisions separately. Further, we do not need to carry the output inventory as a state variable since $Q_n = 0$ for all n under such an optimal policy. We can write the SDP formulation in equation (3.8) as

$$V_n^l(e_n) = \max_{(x_n,m_n)\in\mathcal{A}_n(e_n)} \Big\{(\beta^{N-n}F_n^l - p)m_n - S_n^l x_n + \beta CVaR_n^l\Big(V_{n+1}(e_{n+1},\mathcal{I}_{n+1})\Big)\Big\} \quad (3.11)$$

for each $\mathcal{I}_n^l \in \mathcal{M}_n$.

By Theorem 3.2, $V_n^l(e_n)$ is piecewise linear and concave in e_n for each $\mathcal{I}_n^l \in \mathcal{M}_n$, for all n. As a result, the function $H_n^l(e_{n+1}) \triangleq \beta CVaR_n^l(V_{n+1}(e_{n+1},\mathcal{I}_{n+1}))$ is also piecewise linear and concave in e_{n+1} for all $n < N$. We use the piecewise linear nature of the value functions to characterize the optimal procurement and processing decisions next.

Let $0 = b_n^l(0) < b_n^l(1) < \ldots < b_n^l(k) < b_n^l(k+1) < \ldots < b_n^l(\kappa(l)) < \infty$ denote the break points for $H_n^l(e_{n+1})$; these are the points at which there is a change in slope of $H_n^l(e_{n+1})$. As the number of possible price realizations are finite in each period, we can use an induction argument to prove that both the number of break points $\kappa(l)+1$ as well as the magnitude $b_n^l(\kappa(l))$ is finite. We can express the ending input inventory for period n in terms of these break points, enabling us to write the optimization problem in equation (3.11) as a linear program. To see this, for a given e_{n+1}, define

$$\delta e^l_{n+1}(k) = \begin{cases} 0 & \text{if } e_{n+1} \leq b^l_n(k-1) \\ e_{n+1} - b^l_n(k-1) & \text{if } b^l_n(k-1) < e_{n+1} \leq b^l_n(k) \\ g^l_n(k) & \text{if } b^l_n(k) < e_{n+1} \end{cases} \tag{3.12}$$

for $k = 1, 2, \ldots, \kappa(l)$, where $g^l_n(k) = b^l_n(k) - b^l_n(k-1)$ and

$$\delta e^l_{n+1}(\kappa(l)+1) = \begin{cases} 0 & \text{if } e_{n+1} \leq b^l_n(\kappa(l)) \\ e_{n+1} - b^l_n(\kappa(l)) & \text{if } b^l_n(\kappa(l)) < e_{n+1} \end{cases} \tag{3.13}$$

We can write $e_{n+1} = \sum_{k=1}^{\kappa(l)+1} \delta e^l_{n+1}(k)$ and

$$H^l_n(e_{n+1}) = H^l_n(0) + \sum_{k=1}^{\kappa(l)+1} \Upsilon^l_n(k) \delta e^l_{n+1}(k) \tag{3.14}$$

where

$$\Upsilon^l_n(k) = \frac{H^l_n(b^l_n(k)) - H^l_n(b^l_n(k-1))}{g^l_n(k)} \tag{3.15}$$

for $k = 1, 2, \ldots, \kappa(l)$ and $\Upsilon^l_n(\kappa(l)+1)$ is the slope of H^l_n for $e_{n+1} > b^l_n(\kappa(l))$. By concavity of $H^l_n(e_{n+1})$, we have $\Upsilon^l_n(k+1) < \Upsilon^l_n(k)$ for all $k \leq \kappa(l)+1$. As a result, we can write the problem in equation (3.11) as a linear program

$$V_n^l(e_n) = \max\left\{(\beta^{N-n}F_n^l - p)m_n - S_n^l x_n + H_n^l(0) + \sum_{k=1}^{\kappa(l)+1} \Upsilon_n^l(k)\delta e_{n+1}^m(k)\right\} \quad (3.16)$$

subject to

$$0 \le x_n \le K,\ 0 \le m_n \le C$$

$$0 \le \delta e_{n+1}^l(k) \le g_n^l(k) \text{ for } k = 1, \ldots, \kappa(l) \text{ and}$$

$$0 \le \delta e_{n+1}^l(\kappa(l)+1)$$

$$\sum_{k=1}^{\kappa(l)+1} \delta e_{n+1}^l(k) = e_n + x_n - m_n$$

To characterize the optimal solution for the above linear program, we define $\underline{b_n}(l)$ and $\overline{b_n}(l)$ as follows

$$\underline{b_n}(l) = \begin{cases} b_n^l(k) & \text{if } \exists\ k \le \kappa(l) \text{ s.t. } \Upsilon_n^l(k) > \beta^{N-n}F_n^l - p \ge \Upsilon_n^l(k+1) \\ 0 & \text{if } \Upsilon_n^l(1) \le \beta^{N-n}F_n^l - p \\ \infty & \text{if } \Upsilon_n^l(\kappa(l)+1) > \beta^{N-n}F_n^l - p \end{cases} \quad (3.17)$$

$$\overline{b_n}(l) = \begin{cases} b_n^l(k) & \text{if } \exists\ k \le \kappa(l) \text{ s.t. } \Upsilon_n^l(k) > S_n^l \ge \Upsilon_n^l(k+1) \\ 0 & \text{if } \Upsilon_n^l(1) \le S_n^l \\ \infty & \text{if } \Upsilon_n^l(\kappa(l)+1) > S_n^l \end{cases} \quad (3.18)$$

The optimal procurement and processing policy is then given by the following proposition.

Proposition 3.2. *In any period n, for a realization $\mathcal{I}_n^l \in \mathcal{M}_n$ of the prices, there exist two input inventory levels, $\underline{b_n}(l)$ and $\overline{b_n}(l)$ given by equations (3.17) and (3.18) respectively, such that the optimal procurement and processing quantities $(x_n^*(l), m_n^*(l))$*

are given by

$$x_n^*(l) = \begin{cases} \min\{K, (\overline{b_n}(l) - e_n)^+\} & \textit{if } \overline{b_n}(l) \leq \underline{b_n}(l) \\ \min\{K, (\overline{b_n}(l) + C - e_n)^+\} & \textit{if } \overline{b_n}(l) > \underline{b_n}(l) \end{cases} \tag{3.19}$$

$$m_n^*(l) = \begin{cases} \min\{C, (e_n - \underline{b_n}(l))^+\} & \textit{if } \overline{b_n}(l) \leq \underline{b_n}(l) \\ \min\{C, (e_n + K - \underline{b_n}(l))^+\} & \textit{if } \overline{b_n}(l) > \underline{b_n}(l) \end{cases} \tag{3.20}$$

The quantities $\underline{b_n}(l)$ and $\overline{b_n}(l)$ represent input inventory levels above which the marginal risk-adjusted value-to-go is less than the processing margin $\beta^{N-n}F_n - p$ and the spot price S_n respectively. When $S_n \geq \beta^{N-n}F_n - p$, we have $\overline{b_n}(l) \leq \underline{b_n}(l)$ and the starting input inventory e_n can be separated into three regions: a) $0 \leq e_n < \overline{b_n}(l)$ where it is optimal to procure, b) $\overline{b_n}(l) \leq e_n \leq \underline{b_n}(l)$ where it is optimal to do nothing and c) $e_n > \underline{b_n}(l)$ where it is optimal to process. When $\overline{b_n}(l) > \underline{b_n}(l)$, there is value from procuring and processing the input immediately since $S_n < \beta^{N-n}F_n - p$. In this case, for input inventory levels $e_n \geq \underline{b_n}(l)$ it is optimal for at least one of the activities, procurement and processing, to be up to capacity.

Determining the target inventory levels requires knowledge of the break points and slopes of $H_n^l(\cdot) = \beta CVaR_n^l(V_{n+1}(\cdot,\cdot))$. Since $CVaR$ is not a linear operator, it is possible for the number of break points for $H_n^l(\cdot)$ to be more than the total number of break points of V_{n+1}^m. Further, it is not necessary that $H_n^l(\cdot)$ has the same set of break points for all $\mathcal{I}_n^l \in \mathcal{M}_n$. These facts make the computation of V_n^m complicated, in spite of the piecewise linear nature of the function. In the next section, we describe an approximation to the value function by restricting the set of break points of $H_n^l(\cdot)$ to be no more than the set of break points of the approximation to V_{n+1}.

3.5.2 Approximating the Value Function

We define piecewise linear, concave approximations to V_t^m for each $m \leq M_t$, for each $t = 1, \ldots, N$ with the objective of ensuring that the break points for these approximate value functions can be determined in a simple and systematic manner. To this end, let $\hat{V}_t^m$ denote the approximate value function such that $\hat{V}_t^m \leq V_t^m$ for all m, for all t. Starting with

$$\hat{V}_N^l(e_N) = S_N^l e_N \quad \text{for each } \mathcal{I}_N^l \in \mathcal{M}_N \tag{3.21}$$

we define $\hat{V}_n^m$ in a recursive manner.

Let $\mathcal{B}_{n+1}^m$ denote the set of break points of $\hat{V}_{n+1}^m$. Then, $\mathcal{B}_{n+1} = \bigcup_{m=1}^{M_{n+1}} \mathcal{B}_{n+1}^m$ is the set of all break points of $\hat{V}_{n+1}$. Let $\mathcal{K}$ be the number of elements in $\mathcal{B}_{n+1}$ and without loss of generality, assume the elements of $\mathcal{B}_{n+1}$ are in ascending order. Further, let $\mathcal{B}_{n+1}(k)$ denote the k^{th} element of $\mathcal{B}_{n+1}$ and set $\mathcal{B}_{n+1}(0) = 0$. For each $l \leq M_n$, we define $\hat{H}_n^l$, a piecewise linear and concave function with break points in the set $\mathcal{B}_{n+1}$ and slope $\hat{\Upsilon}_n^l(k)$ between $[\mathcal{B}_{n+1}(k-1), \mathcal{B}_{n+1}(k))$ as follows.

We set $H_n^l(\mathcal{B}_{n+1}(k)) = \beta CVaR_n^l\left(\hat{V}_{n+1}(\mathcal{B}_{n+1}(k), \mathcal{I}_{n+1})\right)$ for each k, where $k = 0, \ldots, \mathcal{K}$. The slope between $[\mathcal{B}_{n+1}(k-1), \mathcal{B}_{n+1}(k))$ is given by

$$\hat{\Upsilon}_n^l(k) = \frac{\hat{H}_n^l\left(\mathcal{B}_{n+1}(k)\right) - \hat{H}_n^l\left(\mathcal{B}_{n+1}(k-1)\right)}{\mathcal{B}_{n+1}(k) - \mathcal{B}_{n+1}(k-1)}$$

for $k = 1, \ldots, \mathcal{K}$ and $\hat{\Upsilon}_n^l(\mathcal{K}+1) = \beta^{N-n} DCVaR_n^l(S_N)$. To complete the specification, we set

$$\hat{H}_n^l(e_{n+1}) = \begin{cases} H_n^l(\mathcal{B}_{n+1}(k-1)) & \\ +\hat{\Upsilon}_n^l(k) \times [e_{n+1} - \mathcal{B}_{n+1}(k-1)] & \text{for } e_{n+1} \in [\mathcal{B}_{n+1}(k-1), \\ & \qquad \mathcal{B}_{n+1}(k)) \\ & \text{and } k \leq \mathcal{K} \\ H_n^l(\mathcal{B}_{n+1}(\mathcal{K})) & \\ +\hat{\Upsilon}_n^l(\mathcal{K}+1) \times [e_{n+1} - \mathcal{B}_{n+1}(\mathcal{K})] & \text{for } e_{n+1} \geq \mathcal{B}_{n+1}(\mathcal{K}) \end{cases} \tag{3.22}$$

and define $\hat{V}_n^l$ as

$$\hat{V}_n^l(e_n) = \max_{\substack{0 \le x_n \le K, \\ 0 \le m_n \le C, \\ 0 \le e_{n+1} = e_n + x - m}} \left\{ [\beta^{N-n} F_n^l - p] m - S_n^l x + \hat{H}_n^l(e_{n+1}) \right\} \tag{3.23}$$

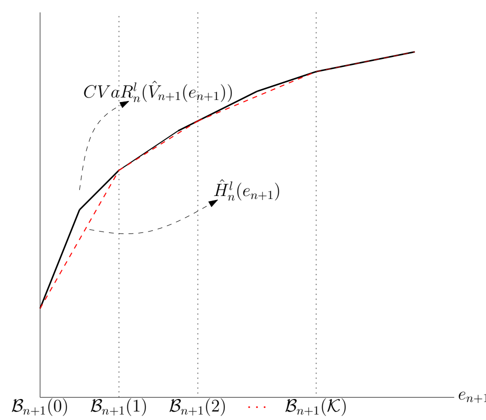

Figure 3.2: $H_n^l(e_{n+1})$ and $\beta CVaR_n^l(\hat{V}_{n+1}(e_{n+1}))$

Notice that equation (3.22) coincides with $\beta CVaR_n^l(\hat{V}_{n+1})$ for $e_{n+1} \in \mathcal{B}_{n+1}$. For e_{n+1} not coinciding with the break points, $H_n^l(e_{n+1})$ is a linear interpolation of the values of $\beta CVaR_n^l(\hat{V}_{n+1})$ at the break points and therefore a lower bound on $\beta CVaR_n^l(\hat{V}_{n+1})$ (see Figure 3.2). As a result, $\hat{V}_n^l$ is a lower bound on V_n^l for each $l \le M_n$. The next theorem states that the set of break points of $\hat{V}_n^l$ are integral multiples of the greatest common divisor of the processing and procurement capacities.

Theorem 3.3. *Let D be the greatest common divisor of the procurement and processing capacities C and K such that $C = aD$ and $K = bD$ where a and b are positive integers. Then $\mathcal{B}_n^l$, the set of break points of $\hat{V}_n^l$, is the set of integer multiples of D such that $\mathcal{B}_n^l = \{0, D, \ldots, [N-n]a \times D\}$ for each $l \le M_n$, for each $n = 1, \ldots, N$.*

By Theorem 3.3, the break points for each $\hat{V}_n^l$ occur only at integral multiples of D. This greatly simplifies the computation of the approximate value function and we can use known results to obtain recursive expressions to compute the slopes of

$\hat{V}_n^l$ between the break points (see Nascimento and Powell, 2009, proposition 2.1) and thereby obtain the heuristic operational policy. We will test the performance of this heuristic policy numerically in Section 3.6.2.

3.5.3 Role of Financial Instruments

As discussed in the introduction and motivation for this research, the growth of commodity exchanges provides firms with additional options to manage the risk from commodity procurement and processing operations. For instance, commodity futures are used extensively by oilseed processors and grain elevators to manage risk in their operations (see Soyatech, 2008; Plato, 2001, for instance). In this section, we consider a firm that maximizes the $DCVaR$ of cash flows over the horizon where the cash flows are a result of both operational and financial activities. Thus, the firm's problem in each period n is still described as in equation (3.7). However, the profits in each period Π_n now include the proceeds from financial trading activity.

In the problem considered here, all output inventory is traded using forward contracts and no counter party risks are assumed. If well traded financial instruments exist for the output commodity, it is possible to perfectly replicate the revenues from output sale commitments and the firm can completely hedge the risk from the output commodity. For instance, when interest rates and commodity price changes are not correlated, the forward prices will coincide with traded futures prices for the commodity (see Hull, 1997, Chap. 3). On the other hand, the input commodity is procured from local spot markets. Further, any input inventory left at the end of the horizon is salvaged / traded. While the spot market and salvage prices may be correlated with the prices of financial derivatives on the input commodity, they are usually not perfectly correlated. Thus, it is of considerable interest to analyze how the firm can use the financial instruments for the input commodity, along with its operational decisions to manage the commodity risk. For the rest of the section, we shall restrict attention to the set of financial instruments available for the input commodity.

Let $j = 1, 2, \ldots, J$ index the financial instruments available for the input commodity and $\mathbf{H}_n = (H_{1n}, \ldots, H_{Jn})$ indicate the firm's position in these instruments at the beginning of period n. We assume that all of the financial positions of the firm are marked to market in each period; that is, at the end of each period the loss or gain in the value of the firm's financial position is assessed and the value of the financial portfolio reset to the current period value (see Hull, 1997, Chap. 2). Thus, without loss of generality, we assume that the firm liquidates its financial positions each period and decides on the positions for the next period. Let $\Theta_n = (\theta_{1n}, \theta_{2n}, \ldots, \theta_{Jn})$ denote the vector of payoffs for these financial instruments in period n (For example, the payoff in period n on a futures contract would be equal to $K_n - \frac{K_{n-1}}{\beta}$, where K_n is the futures price in period n). Thus, the payoff from the financial trading in period n is equal to $\sum_{j=1}^{J} \theta_{jn} H_{jn} = \Theta_n^{\mathsf{T}} \mathbf{H}_n$. We can write the firm's optimization problem as

$$
\begin{aligned}
V_n^l(e_n, \mathbf{H}_n) &= \max_{(x_n, m_n) \in \mathcal{A}_n(e_n),\ \mathbf{H}_{n+1}} \Bigg\{ (\beta^{N-n} F_n^l - p) m_n - S_n^l x_n + (\Theta_n^l)^{\mathsf{T}} \mathbf{H}_n \\
&\qquad + \beta CVaR_n^l \Big(V_{n+1}(e_{n+1}, \mathbf{H}_{n+1}, \mathcal{I}_{n+1}) \Big) \Bigg\} \\
&\qquad\qquad \text{for } n < N \\
V_N^l(e_N, \mathbf{H}_N) &= S_N^l e_N + (\Theta_N^l)^{\mathsf{T}} \mathbf{H}_N
\end{aligned}
$$

for each $\mathcal{I}_n^l \in \mathcal{M}_n$.

Using an induction argument, we can show that $V_n^l(e_n, \mathbf{H}_n)$ is separable in e_n and $\mathbf{H}_n$ and can be written as $V_n^l(e_n, \mathbf{H}_n) = (\Theta_n^l)^{\mathsf{T}} \mathbf{H}_n + U^l(e_n)$, where

$$
\begin{aligned}
U_n^l(e_n) \triangleq U_n(e_n, \mathcal{I}_n^l) = \max_{(x_n, m_n \in \mathcal{A}_n(e_n)} \Bigg\{ & (\beta^{N-n} F_n^l - p) m_n - S_n^l x_n \\
& + \max_{\mathbf{H}_{n+1}} \Big\{ \beta CVaR_n^l \Big(\Theta_{n+1}^{\mathsf{T}} \mathbf{H}_{n+1} + U_{n+1}(e_{n+1}, \mathcal{I}_{n+1}) \Big) \Big\} \Bigg\}
\end{aligned}
$$

and $U_N^l(e_N) \triangleq U_N(e_N, \mathcal{I}_N^l) = S_N^l e_N$.

The next theorem states for each $\mathcal{I}_n \in \mathcal{M}_n$, conditional on the optimal financial hedging decisions, $U_n^l(e_n)$ can be computed without knowing the optimal hedging decision, $\mathbf{H}_{n+1}^*$, explicitly.

Theorem 3.4. *The optimal procurement and processing decisions, conditional on optimal financial hedging decisions being made, are given by the solution to the following SDP:*

$$U_n^l(e_n) = \max_{(x_n, m_n) \in \mathcal{A}_n(e_n)} \left\{ (\beta^{N-n} F_n^l - p) m_n - S_n^l x_n + \beta \mathcal{C}_n^l(e_{n+1}) \right\} \tag{3.24}$$

where

$$\mathcal{C}_n^l(e) = \min_{\psi} \sum_{m=1}^{M_{n+1}} \psi^m U_{n+1}^m(e) \tag{3.25}$$

subject to

$$\sum_{m=1}^{M_{n+1}} \psi^m = 1, \; 0 \leq \psi^m \leq \frac{p^{(l,m)}}{\eta} \; \textit{for } m = 1, \ldots, M_{n+1} \tag{3.26}$$

$$\sum_{m=1}^{M_{n+1}} \psi^m \theta_{j(n+1)}^m = 0 \; \textit{for } j = 1, \ldots, J \tag{3.27}$$

and $p^{(l,m)} = \mathbb{P}\{\mathcal{I}_{n+1} = \mathcal{I}_{n+1}^m | \mathcal{I}_n = \mathcal{I}_n^l\}$, *the transition probability from* $\mathcal{I}_n^l$ *to* $\mathcal{I}_{n+1}^m$.

The SDP in equation (3.24) involves $\mathcal{C}_n^l$, which is the result of financial trading decisions. However, calculating $\mathcal{C}_n^l$ does not require knowledge of the financial trading decisions, but only knowledge of the set of financial instruments traded. This allows for a sequential determination of the optimal operational and financial decisions in the following sense. For a given realization of $\mathcal{I}_n$ and starting input inventory level e_n, the firm can determine the optimal procurement and processing decisions as a solution to equation (3.24). The firm can then solve equation (3.25) for the resulting ending input inventory and obtain the optimal position in financial instrument j as the dual variable to the j^{th} constraint in (3.27). While information exchange between operational and financial trading divisions is necessary for efficient risk management, it is sufficient that operations only needs to know the set of financial instruments traded, and the financial trading division only needs to know the ending input inventory.

3.6. Numerical Study

In this section, we illustrate the analytical results developed in Section 3.5 through numerical experiments. The main goals of the numerical experiments are to 1) study the performance of the heuristic developed in Section 3.5.2, 2) illustrate the benefit of consistent decision making by comparing the performance of a time-consistent risk measure with that of static risk measure on the terminal wealth and 3) quantify the benefit of integrated operational and financial hedging compared to only operational hedging.

3.6.1 Implementation

We model the input spot price as a single factor, mean-reverting process as in Schwartz (1997), while the output forward price is modeled as a driftless geometric Brownian motion with exponentially decreasing volatility (consistent with a mean-reverting price process for the spot price). Specifically, $S_i(t)$, the spot price of the input at time t is modeled as $\ln S_i(t) = \chi_i(t) + \mu(t)$, where $\chi_i(t)$ is the logarithm of the deseasonalized price and $\mu(t)$ is a deterministic factor which captures the seasonality in spot prices. The deseasonalized price $\chi_i(t)$ follows a mean-reverting process given by $d\chi_i(t) = \kappa_i(\xi_i - \chi_i(t))dt + \sigma_i dW_i(t)$ where κ_i is the mean-reversion coefficient, ξ_i is the long run log price level, σ_i is the volatility and $dW_i(t)$ is the increment of a standard Brownian motion. The output forward price $F(t,T)$ is given by $F(t,T) = e^{\mu_f(T)}\hat{F}(t,T)$ where $\mu_f(t)$ is a deterministic factor to capture the seasonality in output prices and $\hat{F}(t,T)$ is the deseasonalized output forward price with dynamics $\frac{d\hat{F}(t,T)}{\hat{F}(t,T)} = \sigma_f e^{-\kappa_f(T-t)} dW_f(t)$. The Brownian motion increments underlying the input and output prices have a constant correlation, ρ. We discretize the input spot and output forward prices and represent the joint evolution of the prices as a multi-dimensional, recombining binomial tree (cf., Peterson and Stapleton (2002) for details on approximating the joint evolution of correlated price processes). The parameters underlying the dynamics of the price processes were estimated based on the soybean, soymeal and soyoil futures prices information for contracts traded on the Chicago Board of Trade (CBOT). For the purposes of the numerical study, we model a single composite output, instead of two individual outputs - meal and oil. The parameters for the composite output are estimated by combining the prices of the soymeal and oil in the proportion in which they are produced upon processing one unit of the input commodity. Such approximations of multiple commodities as a single, composite product have been used in the context of valuing basket options (see Borovkova et al., 2007, for instance) and are sufficient to illustrate the main goals of the numerical study, namely the benefit of time-consistent decision making. The estimated parameters for the price processes of the input and output commodity are given in Table 3.2.

For all the numerical studies, we set the variable cost of processing p to equal 72 cents / bushel, which corresponds to about 35% of the gross margin from processing one bushel of soybean, based on the long run average prices of the three commodities. This value of the processing cost is close to the average processing costs estimated for the US soybean processing industry (Soyatech, 2008). The procurement and processing capacities K and C, were set to 5 and 3 units respectively. These capacities can be considered to be in multiples of bushels, e.g., million bushels. We assume all holding costs are negligible and normalize them to zero, and no discounting, i.e., $\beta = 1$.

3.6.2 Performance of Heuristic

Table 3.2: Price Process Parameters

	Input	Output
Mean-Reversion Coeff κ	0.229	0.5348
Volatility σ	0.244	0.4360
Longrun Log level ξ	6.738	
Seasonality Factor $e^{\mu(t)}$		
Jan	0.992	0.988
Feb	0.992	0.988
Mar	0.998	0.993
Apr	0.998	0.993
May	1.000	0.995
Jun	1.000	0.995
Jul	1.017	1.037
Aug	1.010	1.013
Sep	0.991	1.000
Oct	0.991	0.987
Nov	0.989	0.987
Dec	0.989	0.984

The heuristic described in Section 3.5.2 is based on approximating the optimal value-to-go function by a piecewise linear, concave function with breaks at integral multiples of D. The approximation helps to make the computation tractable by ensuring that the break points of $\hat{V}_n$ for all n occur only at integral multiples of D. Notice that the statement of Theorem 3.3 will still be true if the optimal value-to-go function was approximated by a piecewise linear, concave function with breaks at integral multiples of $D/2^m$, where m is any positive integer; i.e., following the approximation procedure will ensure that the break points of $\hat{V}_n$ for all n will occur only at integral multiples of $D/2^m$. As m increases, the approximation will capture all the break points in the optimal value function and the heuristic value computed using the approximation will converge to the optimal value as $m \to \infty$.

We measure the incremental improvement in the heuristic value as m increases. That is, let $LB_n^\eta(m)$ be the $DCVaR$ value of total profits at the η level for a n period problem when using an approximation with break points at integral multiples

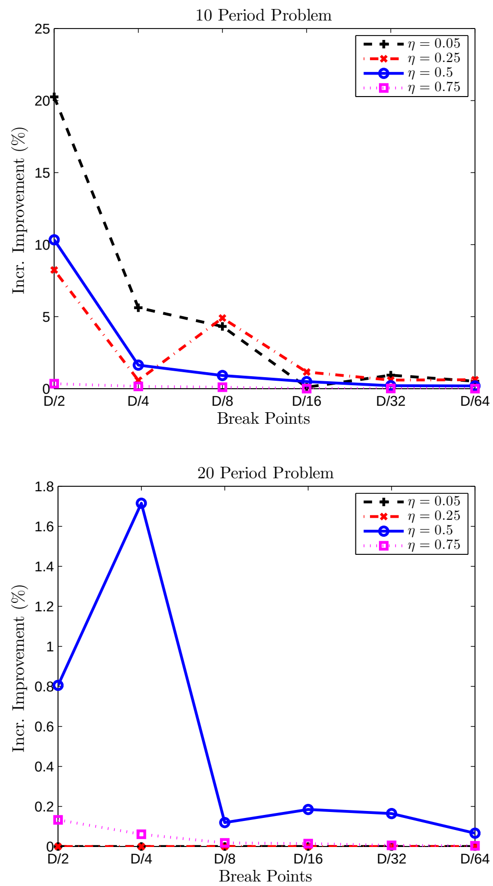

Figure 3.3: Performance of the heuristic: Incremental improvement as number of break points increases

of $D/2^m$ for $m = 0, 1, \ldots$. Then, the % incremental improvement is measured as $\frac{LB_n^\eta(m) - LB_n^\eta(m-1)}{LB_n^\eta(m-1)} \times 100$ for $m = 1, 2, \ldots$. Figure 3.3 illustrates the % incremental improvement for different η levels for horizon lengths of 10 and 20 periods. As seen in Figure 3.3, the incremental improvement for $m \geq 4$ is negligible for all η values and horizon lengths, especially longer horizon lengths. Notice that $D/64 \approx 0.015$ is typical of the step size that may be used when discretizing the inventory values to solve the problem optimally. Thus, we can use the approximation with break points at $D/8$ or $D/16$ with very little loss in optimality and much less computational burden (the computational times for $D/8$ and $D/16$ are roughly 15% to 20% of the computational time using a step size of $D/64$). While we do not report it, we observe similar behavior for other horizon lengths and η values, illustrating that the heuristics are near-optimal for sufficiently small values of m.

3.6.3 Does Time Consistency Matter?

In this section, we compare the performance of a time-consistent risk measure such as $DCVaR$ with that of a static risk measure on the total profits, such as $CVaR^{\eta}$. In order to keep the comparison meaningful, we look at the trade-off between the expected terminal wealth and CVaR of terminal wealth when different objective functions are used for decision making. More precisely, we calculate the distribution of terminal wealth using the policies generated when the decision maker is maximizing $DCVaR^{\eta}$ and $CVaR^{\eta}$ of total profits respectively, for different values of η. One of the motivations for using a time-consistent risk measure was also to manage the risk over the entire horizon, and not just the risk in total payoffs. We consider the distribution of the lowest accumulated profits at any point in the horizon (with negative values representing losses), W_{min}, using the policies generated when $DCVaR^{\eta}$ and $CVaR^{\eta}$ objective functions are used. Since $CVaR^{\eta}$ of total profits is not a time-consistent objective function, we need to consider two situations when the decision maker maximizes $CVaR^{\eta}$ of total profits: 1) a commitment strategy where the decision maker commits to following the optimal policy determined at the beginning of the horizon for all periods and 2) a re-evaluation strategy when the decision maker maximizes $CVaR^{\eta}$ of total profits over the remaining horizon in each period and implements the optimal decision for the current period.

We compare the performance of the different strategies when the initial margin from processing and committing to sell is positive (Figure 3.4), and when it is zero (Figure 3.5). The tradeoff between expected total profits and $CVaR^{\alpha}$ of total profits is generated by computing the distribution of total profits under each strategy for different values of η, and computing the expectation and $CVaR^{\alpha}$, for different values of α. The distribution of minimum wealth over the entire horizon is also computed under each strategy for different values of η. Figures 3.4 and 3.5 show the trade-off between the expected total profits and $CVaR^{\alpha}$ of total profits for $\alpha = 0.10$ and the distribution of minimum wealth over the horizon for policies generated with $\eta = 0.25$.

We see that the time-consistent risk measure dominates the others in the Mean-CVaR sense (Figure 3.4) and distribution of minimum wealth over the horizon (Figure 3.5). While we do not report it here, similar results were obtained when the Mean-CVaR tradeoff of total profits were evaluated for different values of α and the distribution of minimum wealth generated for different values of η were compared. This suggests that the firm can achieve a better risk-return tradeoff in terminal payoffs as well as better risk control over the entire horizon by using time-consistent risk measures.

3.6.4 Benefits of Financial Hedging

We evaluate the benefit of integrating financial trading, in addition to operational decisions, as part of the risk management process. We consider a single financial instrument, a futures contract with maturity at the end of the horizon, and illustrate the benefits of trading the futures contract. Clearly, having the additional option to trade the futures instrument will improve the $DCVaR$ of total profits. Intuitively, using the financial instrument to hedge uncertainty in future period cash flows should also enable the firm to make operational decisions that will maximize the expected value of total profits.

Figure 3.6 shows the tradeoff between the expected total profits and the $DCVaR^{\alpha}$ of total profits for $\alpha = 0.05$ and $\alpha = 0.10$ with and without financial trading. As seen in the figure, integrating financial trading as part of the risk management strategy clearly has benefits for the firm in terms of increased expected profits for the same amount of risk undertaken.

3.7. Conclusions

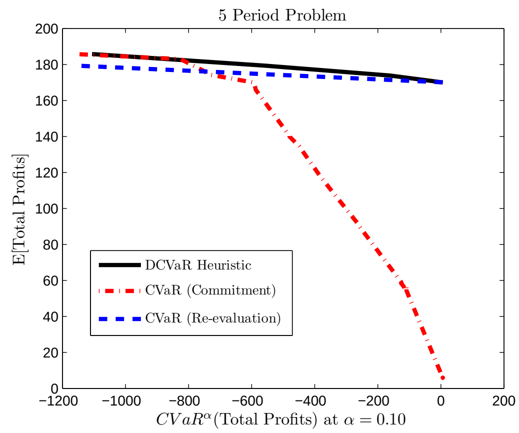

(a) Mean-CVaR tradeoff of terminal wealth W_{N+1}

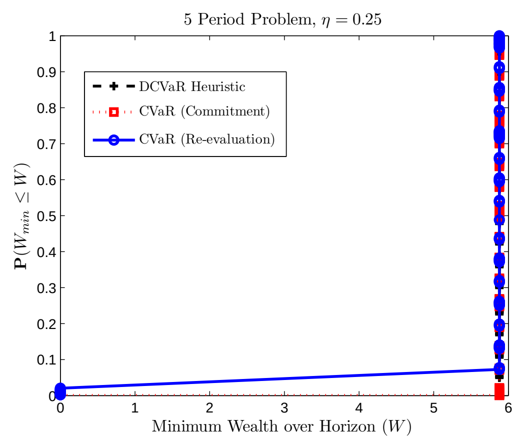

(b) Distribution of minimum wealth over the horizon W_{min}

Figure 3.4: Role of Time Consistency: Performance of different risk-averse objective functions when initial processing margin is positive

In this chapter, we have considered the dynamic operational and financial risk management for a commodity processing firm. We proposed a model to capture the firm's risk aversion using a time consistent, dynamic risk measure based on CVaR. We characterized the optimal operational policies and showed that the procurement and processing decisions in any period are governed by price dependent 'procure up to' and 'process down to' input inventory thresholds. Further, we showed that in the presence of optimal financial hedging, these thresholds can be obtained without knowing the details of the financial hedging decisions themselves. We developed an

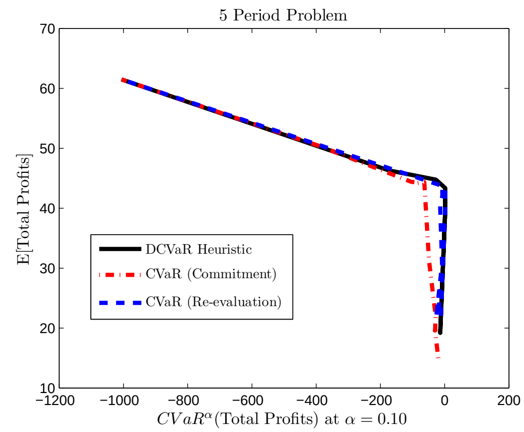

(a) Mean-CVaR tradeoff of terminal wealth W_{N+1}

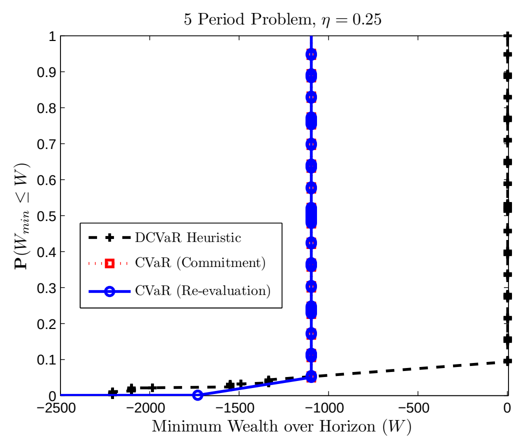

(b) Distribution of minimum wealth over the horizon W_{min}

Figure 3.5: Role of Time Consistency: Performance of different risk-averse objective functions when initial processing margin is zero

efficient heuristic to compute the optimal operational and financial decisions and using numerical studies showed that these heuristics are near optimal. Using numerical studies, we showed that using the optimal policies obtained from a time consistent objective function provide a better mean-CVaR tradeoff for the total profits, compared to those obtained when using the CVaR of the total profits as the objective function. In addition, we found that using a time consistent risk measure also minimizes the probability of extreme losses over the entire horizon.

Our work is one of the few early attempts to model the dynamic risk management

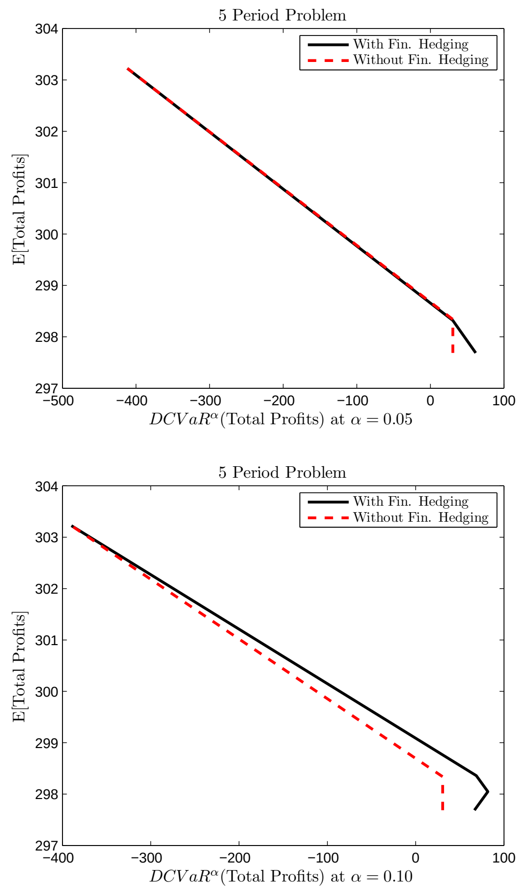

Figure 3.6: Value of Financial Hedging: Mean-DCVaR profiles of total profits with and without financial hedging

problem for firms dealing with storable commodities. While we have considered a firm that maximizes a dynamic risk measure, many commodity trading and processing firms think of risk control in terms of constraints. That is, they are interested in maximizing expected profits over the horizon, while imposing limits on the total risk that can be taken, where risk is measured using a risk measure such as CVaR or variance. This is the approach taken in Kleindorfer and Li (2005), for instance. Exploring the issue of time consistency in such dynamic risk constrained optimization problems will be an interesting area for future research. Further, the context of the problem can be expanded to include multiple input / output commodity sets where the firm has a choice to decide which input to process or what output to produce.

3.8. Appendix: Proofs of Theorems and Lemmas

Proof of Proposition 3.1. Suppose $\mathbf{A}$ and $\mathbf{B}$ are two cash flow streams such that $\mathcal{A}_n(\mathcal{S}_n) \geq \mathcal{B}_n(\mathcal{S}_n)$ and $DCVaR_{n+1}(\mathbf{A}, \mathcal{S}_{n+1}\ ;\eta_{n+1}) \geq DCVaR_{n+1}(\mathbf{B}, \mathcal{S}_{n+1}\ ;\eta_{n+1})$ for all $\mathcal{S}_{n+1} \in \mathcal{H}_{n+1}(\mathcal{S}_n)$ Then, we have

$$\begin{aligned} DCVaR_n(\mathbf{A}, \mathcal{S}_n\ ;\eta_n) &= CVaR_n^{\eta}\Big(\mathcal{A}_n(\mathcal{S}_n) + DCVaR_{n+1}(\mathbf{A}, \mathcal{S}_{n+1}\ ;\eta_{n+1})\,|\mathcal{S}_n\Big) \\ &\geq CVaR_n^{\eta}\Big(\mathcal{B}_n(\mathcal{S}_n) + DCVaR_{n+1}(\mathbf{B}, \mathcal{S}_{n+1}\ ;\eta_{n+1})\,|\mathcal{S}_n\Big) \end{aligned}$$

where the inequality follows from the fact that $CVaR^{\eta}$ is a coherent risk measure and therefore monotonic. Notice that

$$CVaR_n^{\eta}\Big(\mathcal{B}_n(\mathcal{S}_n) + DCVaR_{n+1}(\mathbf{B}, \mathcal{S}_{n+1}\ ;\eta_{n+1})\,|\mathcal{S}_n\Big) = DCVaR_n(\mathbf{B}, \mathcal{S}_n\ ;\eta_n).$$

Thus, $DCVaR(\cdot;\eta)$ given by equation (3.3) is time consistent. □

Proof of Theorem 3.1. From equation (3.7) and the definition of $DCVaR_n^{\eta}$ given in equation (3.3), we have

$$
\begin{aligned}
V_n & (e_n, Q_n, \mathcal{I}_n) \\
= & \max_{\substack{((x_t,m_t,q_t)\in\mathcal{A}_t(e_t,Q_t)),\\ t=n,\ldots,N-1}} \Bigg\{ CVaR_n^{\eta}\Bigg(\Pi_n(x_n, m_n, q_n, \mathcal{I}_n) \\
& \qquad + \; DCVaR_{n+1}^{\eta}\left(\sum_{t=n+1}^{N} \beta^{t-n}\Pi_t(x_t, m_t, q_t, \mathcal{I}_t) \right)\Bigg)\Bigg\} \\
= & \max_{\substack{((x_t,m_t,q_t)\in\mathcal{A}_t(e_t,Q_t)),\\ t=n,\ldots,N-1}} \Bigg\{ \Pi_n(x_n, m_n, q_n, \mathcal{I}_n) + CVaR_n^{\eta}\Bigg(\\
& \qquad \beta DCVaR_{n+1}^{\eta}\left(\sum_{t=n+1}^{N} \beta^{t-(n+1)}\Pi_t(x_t, m_t, q_t, \mathcal{I}_t) \right)\Bigg)\Bigg\} \\
= & \max_{(x_n,m_n,q_n)\in\mathcal{A}_n(e_n,Q_n)} \Bigg\{ \Pi_n(x_n, m_n, q_n, \mathcal{I}_n) + \\
& \max_{\substack{((x_t,m_t,q_t)\in\mathcal{A}_t(e_t,Q_t)),\\ t=n+1,\ldots,N-1}} \Bigg\{ \beta CVaR_n^{\eta} \\
& \qquad \Bigg(DCVaR_{n+1}^{\eta}\left(\sum_{t=n+1}^{N} \beta^{t-(n+1)}\Pi_t(x_t, m_t, q_t, \mathcal{I}_t) \right)\Bigg)\Bigg\}\Bigg\} \\
= & \max_{(x_n,m_n,q_n)\in\mathcal{A}_n(e_n,Q_n)} \Bigg\{ \Pi_n(x_n, m_n, q_n, \mathcal{I}_n) + \\
& \beta CVaR_n^{\eta} \\
& \qquad \Bigg(\max_{\substack{((x_t,m_t,q_t)\in\mathcal{A}_t(e_t,Q_t)),\\ t=n+1,\ldots,N-1}} \\
& \qquad \Bigg\{ DCVaR_{n+1}^{\eta}\left(\sum_{t=n+1}^{N} \beta^{t-(n+1)}\Pi_t(x_t, m_t, q_t, \mathcal{I}_t) \right)\Bigg\}\Bigg)\Bigg\} \\
= & \max_{(x_n,m_n,q_n)\in\mathcal{A}_n(e_n,Q_n)} \Bigg\{ \Pi_n(x_n, m_n, q_n, \mathcal{I}_n) \\
& \qquad + \beta CVaR_n^{\eta}\Big(V_{n+1}(e_{n+1}, Q_{n+1}, \mathcal{I}_{n+1}) \Big)\Bigg\}
\end{aligned}
$$

The second equality follows from the fact that the $CVaR$ is a coherent risk measure and therefore translation invariant and $\Pi(q_n, x_n, \mathcal{I}_n)$ is known with certainty in period n. The interchange between max and $CVaR^\eta$ operations in the fourth step is valid because $CVaR^\eta$ is monotonic. Thus, we can write the optimization problem as a stochastic dynamic program. □

Proof of Lemma 3.1. The proof follows from a standard induction argument. The statement is true for $n = N$. Suppose it is true for all $t = n+1, \ldots, N$. The set

$$\left\{(e_n, Q_n, x_n, m_n, q_n) : e_n \in \mathbb{R}^+, Q_n \in \mathbb{R}, q_n \geq 0, (x_n, m_n) \in \mathcal{A}_n(e_n)\right\}$$

is convex. Further, since $CVaR^\eta$ is a coherent risk measure, it is concave and therefore the maximand in equation (3.8) is concave in $(x_n, m_n, q_n, e_n, Q_n)$. Thus, from Proposition B-4 in Heyman and Sobel (1984), $V_n(e_n, Q_n, \mathcal{I}_n)$ is concave in (e_n, Q_n) for each $\mathcal{I}_n$. The increasing part is straightforward to prove using similar induction arguments. □

We state and prove a lemma that will be used in many of the subsequent proofs.

Lemma 3.8.1. *For a random variable X with finite support $\mathcal{X} = (X^1, \ldots, X^M)$, we have*

$$\begin{aligned} CVaR^\eta(X) &= \min_\psi \sum_{m=1}^M \psi^m X^m \\ &\quad s.t. \quad 0 \leq \psi^m \leq \frac{p^m}{\eta} \text{ for all } m, \; \sum_{m=1}^M \psi^m = 1 \end{aligned} \tag{3.28}$$

where $p^m = \mathbb{P}\{X = X^m\}$ for each m.

Proof. From equation (3.1), we have $CVaR^\eta(X) = \max_\upsilon \left\{\upsilon - \frac{1}{\eta}\mathbb{E}\left[(\upsilon - X)^+\right]\right\}$ which can be written as the following linear program when X has finite support.

$$CVaR^{\eta}(X) = \max_{\upsilon,\mathbf{y}} \left\{ \upsilon - \frac{1}{\eta}\sum_{m=1}^{M} p^m y^m \right\}$$

$$\text{s.t.} \quad y^m \geq \upsilon - X^m,\ y^m \geq 0 \text{ for all } m$$

Writing the dual of the above linear program, where ψ^m is the dual variable corresponding to the m^{th} constraint, we have

$$CVaR^{\eta}(X) = \min_{\psi} \sum_{m=1}^{M} \psi^m X^m$$

$$\text{s.t.} \quad 0 \leq \psi^m \leq \frac{p^m}{\eta} \text{ for all } m,\ \sum_{m=1}^{M} \psi^m = 1$$

and the lemma is proved. □

Proof of Theorem 3.2. Clearly, $V_N^m(e_N, Q_N)$ is linear and thus piecewise linear and concave in (e_N, Q_N), for each $m \leq M_N$. Suppose $V_t^m(e_t, Q_t)$, for $t = n+1, n+2, \ldots, N$ is piecewise linear, concave and continuous in (e_t, Q_t), for each $m \leq M_t$.

It is useful to consider the function $H_n^l(e_{n+1}, Q_{n+1}) \triangleq CVaR_n^l(V_{n+1}(e_{n+1}, Q_{n+1}, \mathcal{I}_{n+1}))$. Let $p^{(l,m)}$ denote the transition probability from $\mathcal{I}_n^l$ to $\mathcal{I}_{n+1}^m$. From Lemma 3.8.1, we can write

$$H_n^l(e_{n+1}, Q_{n+1}) = \min_{\psi^m} \sum_{m=1}^{M_{n+1}} \psi^m V_{n+1}^m(e_{n+1}, Q_{n+1})$$

$$\text{s.t.} \quad \sum_{m=1}^{M_{n+1}} \psi^m = 1 \text{ and } 0 \leq \psi^m \leq \frac{p^{(l,m)}}{\eta} \text{ for all } m$$

Let ψ^{m*} be the optimal solution to the dual problem above for a given value of (e_{n+1}, Q_{n+1}). By the continuity of V_{n+1}^m, there exist $\delta_e > 0$ and $\delta_Q > 0$ such that ψ^{m*} is optimal for all values (e, Q) such that $|e - e_{n+1}| < \delta_e$ and $|Q - Q_{n+1}| < \delta_Q$. Thus, it is easy to see that $H_n^l(e_{n+1}, Q_{n+1})$ is the weighted sum of piecewise linear, concave functions in (e_{n+1}, Q_{n+1}) and therefore H_n^l is piecewise linear, concave and continuous in (e_{n+1}, Q_{n+1}). We have

$$V_n^l(e_n, Q_n) \quad = \quad \max_{(x_n,m_n)\in\mathcal{A}_n(e_n),\ q_n\geq 0} \left\{\beta^{N-n}F_n^l q_n - pm_n - S_n^l x_n + \beta H_n^l(e_{n+1}, Q_{n+1})\right\}$$

Since $H_n^l(e_{n+1}, Q_{n+1})$ is piecewise linear and concave in (e_{n+1}, Q_{n+1}), V_n^l is the solution of a linear program in which e_n appears in the right hand side of the constraint set. Thus, V_n^l is piecewise linear, concave and continuous in (e_n, Q_n), for each l. Thus, by induction, the theorem is true for all $l \leq M_n$, for all n. □

Proof of Lemma 3.2. From equation (3.9), it is clear that $q_{N-1}^* \leq [Q_{N-1}+m_{N-1}^*]^+$. Notice that V_{N-1} is the solution of a linear programming problem and it can be shown that $\frac{\partial V_{N-1}}{\partial Q_{N-1}} \geq \beta F_{N-1}$ for all $Q_{N-1} < 0$.

For a given procurement and processing decision, consider the optimal commitment decision in period n. We can write the optimization problem as

$$\begin{aligned} &\max_{q,q_o} \left\{\beta^{N-n}F_n(q+q_o) + \beta CVaR_n\left(V_{n+1}(e_{n+1}, Q_n + m_n - (q+q_o), \mathcal{I}_{n+1})\right)\right\} \\ \text{s.t.} \quad & 0 \leq q \leq Q_n + m_n,\ q_o \times (Q_n + m_n - q) = 0,\ q_o \geq 0 \end{aligned}$$

where q_o denotes the extent of over-commitment.

If $q < Q_n + m_n$ at optimality, then $q_o = 0$ and there is nothing to prove. Therefore we only consider the situation where $q = Q_n + m_n$ and show that $q_o = 0$ even in this case.

Since V_n is piecewise linear in (e_n, Q_n), there exist $\hat{Q}_n^m > 0$ for all n and $m \in \mathcal{M}_n$ such that $V_n^m(e_n, -Q) = V_n^m(e_n, 0) - \gamma_n^m Q$ for $0 \leq Q \leq \hat{Q}_n^m$. As the induction assumption, let $\gamma_t^m \geq \beta^{N-t}F_t^m$ for all $m \in \mathcal{M}_t$, for all $t > n$. By concavity of V_n^m in Q_n, the slope of the value function for $Q < -\hat{Q}_t^m$ will be less than $-\gamma_t^m$ for all

$m \le M_t$, for all t. Thus proving $q_o = 0$ at optimality by substituting $V_n^m(e_n, -Q) = V_n^m(e_n, 0) - \gamma_n^m Q$ for all $Q \ge 0$ is sufficient to prove the theorem.

The optimal over-commitment can be obtained as a solution to the following linear program

$$\max_{q_o \ge 0,\ \mathbf{y} \ge \mathbf{0},\ \upsilon} \left\{ \beta^{N-n} F_n q_o + \beta\{\upsilon - \frac{1}{\eta} \sum_{m=1}^{M_{n+1}} p^{(l,m)} \times y^m\} \right\}$$
$$\text{s.t.} \quad y^m \ge \upsilon - (V_{n+1}^m(e_{n+1}, 0) - \gamma_{n+1}^m q_o) \text{ for all } m \in \mathcal{M}_{n+1}$$

where $p^{(l,m)}$ is the transition probability from $\mathcal{I}_n^l$ to $\mathcal{I}_{n+1}^m$.

Writing the dual of the linear program above, we have

$$\min_{\psi \ge \mathbf{0}} \beta \sum_{m=1}^{M_{n+1}} \psi^m V_{n+1}^m(e_{n+1}, 0)$$
$$\text{s.t.} \quad 0 \le \psi^m \le \frac{p^{(l,m)}}{\eta} \text{ for all } m \in \mathcal{M}_{n+1},\ \sum_{m=1}^{M_{n+1}} \psi^m = 1,\ \sum_{m=1}^{M_{n+1}} \psi^m \gamma^m \ge \beta^{N-(n+1)} F_n$$

By the condition in the statement of the lemma that the forward prices are unbiased and the induction assumption that $\gamma^m \ge \beta^{N-(n+1)} F_{n+1}^m$, we know that $\psi^m = p^m$ is a feasible solution to the above problem. Thus, the optimal value to the above problem is bounded from above. Noticing that q_o is the lagrange multiplier for the last constraint in the above minimization problem, the optimal solution to the above problem can be written as

$$\max_{q_o \ge 0} \left\{ \min_{\psi} \left\{ \beta \sum_{m=1}^{M_{n+1}} \psi^m V_{n+1}^m(e_{n+1}, 0) - q_o \times (\sum_{m=1}^{M_{n+1}} \psi^m \gamma^m - \beta^{N-(n+1)} F_n) \right\} \right\}$$
$$= \min_{\psi} \left\{ \max_{q_o \ge 0} \left\{ \beta \sum_{m=1}^{M_{n+1}} \psi^m V_{n+1}^m(e_{n+1}, 0) - q_o \times (\sum_{m=1}^{M_{n+1}} \psi^m \gamma^m - \beta^{N-(n+1)} F_n) \right\} \right\}$$

subject to the constraints $0 \le \psi^m \le \frac{p^{l,m}}{\eta}$ for all $m \in \mathcal{M}_{n+1}$, $\sum_{m=1}^{M_{n+1}} \psi^m = 1$. The equality above follows from strong duality.

By linearity of the objective function in q_o, we will always have $q_o = 0$ or $q_o = \infty$ at optimality in the inner maximization problem above. Since the optimal value for the overall problem is finite and bounded, the optimal choice of ψ will be such that $q_o = 0$ at optimality. Hence, over-commitment is never optimal.

Notice that $\sum_{m=1}^{M_{n+1}} \psi^m \gamma_{n+1}^m \geq \beta^{N-(n+1)} F_n$ at optimality. Using optimality conditions for the commitment decision, we can show that $\gamma_n^l \geq \beta^{N-n} F_n^l$ where $V_n^l(e_n, Q_n) = V_n^l(e_n, 0) - \gamma_n^l Q_n$ for $Q_n < 0$, for all $l \leq M_n$. Thus, by induction, it is never optimal to over-commit in any period. $\square$

Proof of Lemma 3.3. In period $N-1$, for $Q_{N-1} \geq 0$, it is clearly optimal to commit all available output inventory for sale; i.e., $q_{N-1}^* = Q_{N-1} + m_{N-1}^*$. Further, we can write $V_{N-1}^m(e_{N-1}, Q_{N-1}) = \beta F_{N-1}^m Q_{N-1} + U_{N-1}^m(e_{N-1})$ where

$$
\begin{aligned}
U_{N-1}^m(e_{N-1}) \quad = \quad & \max_{x,m \in \mathcal{A}_{N-1}(e_{N-1})} \big\{ (\beta F_{N-1}^m - p)m - S_{N-1}x \\
& \qquad + \beta CVaR_{N-1}^m(S_N) \times [e_{N-1} + x - m] \big\}
\end{aligned}
$$

As the induction assumption, for all $t \geq n+1$, let $V_t^m(e_t, Q_t) = \beta^{N-t} F_t^m Q_t + U_t^m(e_t)$ for all $Q_t \geq 0$ and $U_t^m(e_t)$ is such that $U^{m_1}(e_t) \geq U^{m_2}(e_t)$ whenever $F_t^{m_1} \geq F_t^{m_2}$ for all $m_1, m_2 \in \mathcal{M}_t$. From Lemma 3.8.1, we have

$$
\begin{aligned}
CVaR_n^l(V_{n+1}(e_{n+1}, Q_{n+1}, \mathcal{I}_{n+1})) \quad = \quad & \min_{\psi} \sum_{m=1}^{M_{n+1}} \psi^m U^m(e_{n+1}) \\
& \qquad + \sum_{m=1}^{M_{n+1}} \psi^m \beta^{N-(n+1)} F_{n+1}^m Q_{n+1} \\
& \text{s.t.} \quad 0 \leq \psi^m \leq \frac{p^{(l,m)}}{\eta} \text{ for all } m, \ \sum_{m=1}^{M} \psi^m = 1
\end{aligned}
$$

Notice that $\psi^m = p^{(l,m)}$ for all m is a feasible solution to the above problem and by the condition in Lemma 3.3, $\sum_{m=1}^{M_{n+1}} p^{(l,m)} F_{n+1}^m = F_n^l$. By the induction assumption that $U^{m_1}_{n+1} \geq U^{m_2}_{n+1}$ whenever $F^{m_1}_{n+1} \geq F^{m_2}_{n+1}$, an optimal solution to the above problem will always be such that $\sum \psi^m \beta^{N-(n+1)} F^m_{n+1} \leq \beta^{N-(n+1)} F^l_n$ for all values of e_{n+1} and $Q_{n+1} \geq 0$. Therefore, the slope of $\beta CVaR^l_n(V_{n+1}(e_{n+1}, Q_{n+1}))$ with respect to Q_{n+1} is never greater than $\beta^{N-n} F^l_n$.

Now consider the optimal commitment decision in period n. Since the slope of $\beta CVaR^l_n(V_{n+1}(e_{n+1}, Q_{n+1}))$ with respect to Q_{n+1} is never greater than $\beta^{N-n} F_n$ for all $Q_{n+1} \geq 0$, it is optimal to commit all available output inventory in period n. As a result, we can write

$$
\begin{aligned}
V^l_n(e_n, Q_n) &= \max_{x,m \in \mathcal{A}_n(e)} \left\{ \beta^{N-n} F^l_n (Q_n + m_n) - pm - S^l_n x + \beta CVaR^l_n \left(U_{n+1}(e_{n+1}) \right) \right\} \\
&= \beta^{N-n} F^l_n Q_n \\
&\qquad + \underbrace{\max_{x,m \in \mathcal{A}_n(e)} \left\{ (\beta^{N-n} F^l_n - p) m - S^l_n x + CVaR^l_n \left(U_{n+1}(e_{n+1}) \right) \right\}}_{U^l_n(e_n)}
\end{aligned}
$$

By the condition in the lemma that F_{n+1} is stochastically increasing in F_n, the induction assumption about U^m_{n+1} and the monotonicity of $CVaR$, we have $U^{l_1}_n \geq U^{l_2}_n$ whenever $F^{l_1}_n \geq F^{l_2}_n$. Thus, by induction, the lemma is true for all $l \in \mathcal{M}_n$, for all n. □

Proof of Proposition 3.2. Consider the optimization given by equation (3.29). The quantities $\overline{b_n}(l)$ and $\underline{b_n}(l)$ are the optimal ending input inventory levels when the procurement and processing decisions are considered independently and without imposing the capacity constraints.

By concavity of H^l_n, we have $\overline{b_n}(l) \leq \underline{b_n}(l)$ whenever $S^l_n \geq F^l_n - p$. Thus, the starting inventory levels for which positive procurement and processing are optimal are mutually exclusive. Imposing the procurement and processing capacity constraints gives the optimal procurement and processing quantities given by equations (3.19)–(3.20) for this case.

When $\overline{b_n}(l) > \underline{b_n}(l)$, we have $S^l_n < F^l_n - p$. That is, there is value in procuring and processing immediately in the current period. The optimal quantity to process, when there is no processing capacity restriction but a procurement capacity limit of K units would therefore be $(e_n + K - \underline{b_n}(l))^+$. Likewise, the optimal quantity to procure would be $(\overline{b_n}(l) + C - e_n)^+$ when there is no procurement capacity restriction, but a processing capacity restriction of C units. Imposing the procurement and processing capacity limits together leads to the optimal quantities given in equations (3.19)–(3.20) for this case. $\square$

Proof of Theorem 3.3. We prove the theorem by induction. The theorem is true for $V^l_N = S^l_N e_N$, with the set $\mathcal{B}^l_N = \{0\}$ for each $l \in \mathcal{M}_N$.

Assume the theorem is true for $t = n+1, \ldots, N-1, N$. Thus,

$$\mathcal{B}_{n+1} = \{0, D, \ldots, kD, \ldots, [N-(n+1)]aD\}.$$

By construction, the set of break points of $\hat{H}^l_n$, for each $l \in \mathcal{M}_n$ is the same as $\mathcal{B}_{n+1}$. Recall that $C = aD$ and $K = bD$, where a and b are positive integers. The approximate value function V^l_n as given by equation (3.23) is the solution of a linear program where all the decision variables are integral multiples of D. Therefore, $\hat{V}^l_n$ also has changes in slope only at integral multiples of D; i.e., $\mathcal{B}^l_n$ contains only integral multiples of D for each $l \leq M_n$. Further, for $e_n \geq [N-n]aD$ we have $e_{n+1} = e_n - C + K \geq e_n - aD + K \geq [N-(n+1)]aD$. Thus, by the induction assumption, the slope of $\hat{V}^l_n$ with respect to e_n is constant for $e_n \geq [N-n]aD$. By induction, the theorem is therefore true for all n. $\square$

Proof of Theorem 3.4. For a given realization $\mathcal{I}^l_n \in \mathcal{M}_n$ and starting inventory level, the optimal operational decisions are given by the solution to

$$\max_{(x_n,m_n)\in\mathcal{A}_n(e_n)} \left\{ (F_n^l - p)m_n - S_n^l x_n + \beta \max_{\mathbf{H}_{n+1}} \left\{ CVaR_n^l\left(\Theta_{n+1}^{\mathrm{T}}\mathbf{H}_{n+1} + U_{n+1}(e_{n+1}, \mathcal{I}_{n+1})\right)\right\} \right\}$$

Considering the maximization over the financial hedging decisions above, we have

$$\begin{aligned}
\mathcal{C}_n^l(e_{n+1}) &= \max_{\mathbf{H}_{n+1}} \left\{ CVaR_n^l\left(\Theta_{n+1}^{\mathrm{T}}\mathbf{H}_{n+1} + U_{n+1}(e_{n+1}, \mathcal{I}_{n+1})\right)\right\} \\
&= \max_{\mathbf{H}_{n+1},\upsilon} \left\{ \upsilon - \frac{1}{\eta}\mathbb{E}\left[\left(\upsilon - (\Theta_{n+1}^{\mathrm{T}}\mathbf{H}_{n+1} + U_{n+1}(e_{n+1}, \mathcal{I}_{n+1}))\right)^+\right]\right\}
\end{aligned}$$

where the second equality follows from Rockafellar and Uryasev (2000). For a discrete price process, $\mathcal{C}_n^l$ can be written as the solution of a linear program, given by

$$\max_{\mathbf{H}_{n+1},\upsilon,\mathbf{y}} \upsilon - \frac{1}{\eta}\sum_{m=1}^{M_{n+1}} p^{(l,m)} y^m$$

subject to

$$y^m \geq \upsilon - \left(U_{n+1}^m(e_{n+1}) + (\Theta_{n+1}^m)^{\mathrm{T}}\mathbf{H}_{n+1}\right) \text{ for } m = 1, \ldots, M_{n+1}$$

$$y^m \geq 0 \text{ for } m = 1, \ldots, M_{n+1}$$

Writing the dual of the above linear program, we have

$$\mathcal{C}_n^l(e_{n+1}) = \min_{\psi} \sum_{m=1}^{M_{n+1}} \psi^m U_{n+1}^m(e_{n+1})$$

subject to

$$0 \leq \psi^m \leq \frac{p^{(l,m)}}{\eta} \text{ for } m = 1, \ldots, M_{n+1},$$

$$\sum_{m=1}^{M_{n+1}} \psi^m = 1,$$

$$\sum_{m=1}^{M_{n+1}} \psi^m \theta_{j(n+1)}^m = 0 \text{ for } j = 1, \ldots, J$$

where ψ^m for $m = 1, \ldots, M_{n+1}$ is the dual variable corresponding to the m^{th} inequality constraint in the primal problem. Thus, we can obtain the optimal operational decisions without explicitly calculating the financial hedging decisions. Also note that the optimal position in financial instrument j for a given ending inventory level e_{n+1} is given by the dual variable corresponding to the constraint $\sum_{m=1}^{M_{n+1}} \psi^m \theta^m_{j(n+1)} = 0$. $\square$

3.9. Appendix: Break points of Value Function V_n^l

We show there exists a step size such that the break points for V_n^l, for all l, for all n are integral multiples of the step size. To this end, recall that $b_n^l(k)$, for $k = 0, \ldots, \kappa(l)$ are the break points of $H_n^l(e) = \beta CVaR_n^l(V_{n+1}(e, \mathcal{I}_{n+1}))$, for each $l \leq M_n$, for each $n < N$ and $g_n^l(k) = b_n^l(k) - b_n^l(k-1)$ for $k = 1, \ldots, \kappa(l)$. Define $\mathcal{G}_n^l = \bigcup_k g_n^l(k)$, the set of lengths of the linear segments of $H_n^l(e)$. Now, let $\mathcal{G}_n = \bigcup_l \mathcal{G}_n^l$ and $\mathcal{G} = \bigcup_n \mathcal{G}_n$. Thus, $\mathcal{G}$ is the set of lengths of all the linear segments of $H_n^l(\cdot)$, for all l, for all $n < N$. Since the set of possible prices in each period is a finite set, all the elements of $\mathcal{G}$ are rational numbers. Let τ be the greatest common divisor of all the elements of $\mathcal{G} \cup \{K, C\}$. Thus, the break points of H_n^l for each l, for each n are integral multiples of τ. From equation (3.16), we see that V_n^l is the solution of a linear program in which the right hand side coefficients of all the inequality constraints are integral multiples of τ. Thus, V_n^l also has break points at integral multiples of τ for all l, for all n.

Chapter 4

Commodity Operations in Partially Complete Markets

4.1. Introduction

The growth of commodity exchanges provides a wide variety of financial instruments—futures, forwards, options etc.—for firms involved in the commodities business to manage the risk from commodity price uncertainty (Geman, 2005). Thus, firms have an opportunity to combine financial trading along with operational decisions to better manage the commodity price uncertainty. While Chapter 3 considers the role of financial hedging for a risk-averse commodity processing firm, it does not provide specific details of the structure of the financial trading decisions themselves. Further, the analysis therein did not look into the value of operational hedging, i.e., the benefit of having excess procurement or processing capacity to manage the price uncertainty. This is partly because the underlying model for the price uncertainty made no specific assumptions regarding the dynamics of commodity price movements and hence did not lend itself to tractable analysis.

We address these issues in the current chapter. Specifically, we model the uncertainty in commodity prices using a partially complete markets framework. The partially complete markets framework distinguishes between market related and firm specific (private) uncertainty, and was introduced by Smith and Nau (1995) in the context of valuing a stream of uncertain cash flows. We extend $DCVaR$, the time-consistent risk measure introduced in Chapter 3, to this framework and model the operational and financial trading decisions for a risk-averse firm. Specifically,

1. We characterize the optimal financial trading policy for a risk-averse firm maximizing the time-consistent risk measure over the total cash flows from financial trading and any specific operational policy.

2. Contingent on optimal financial trading, we characterize the optimal operational decisions for a commodity processing firm that procures an input commodity from the spot market, processes and sells the output using forward contracts over a multi-period horizon. The firm also earns revenues by trading any remaining input inventory at the end of the horizon. We show that

   (a) The optimal trade policy for selling the output is identical to the risk-neutral optimal commitment policy.

   (b) The value function is concave and piecewise linear, with breakpoints at integral multiples of the greatest common divisor of the procurement and processing capacities.

   (c) The optimal procurement and processing decisions are governed by 'procure up to' and 'process down to' thresholds. These thresholds depend on the realized market and private uncertainties as well as the horizon length. Unlike in the risk-neutral case, these thresholds depend on the firm's subjective probabilities over private uncertainties.

3. We characterize the value of operational hedging. Under the restriction that the worst case expected salvage value is no more than the benefit from processing and selling the output,

   (a) We show that excess processing capacity (relative to procurement capacity) does not provide any value.

(b) On the other hand, we show that excess procurement capacity has value as an operational hedge because it helps the firm to opportunistically procure additional input in periods when the realized spot price is sufficiently low, for processing in future periods. We develop recursive expressions that quantify the value from excess procurement capacity.

The rest of the chapter is organized as follows. In section 4.2, we describe the partially complete market framework and extend the definition of DCVaR to this framework. We also characterize the optimal financial trading policy to hedge any specific stream of random cash flows; e.g., cash flows from operations. Section 4.3 models the commodity processing firm's operations, while section 4.4 characterizes the optimal operational policies. We quantify the value of operational hedging, i.e., the benefit of having excess procurement or processing capacity, in section 4.5. Section 4.6 concludes with directions for future research. Proofs for the various theorems and lemmas are given in the chapter appendix.

4.2. Partially Complete Markets

We use the framework of partially complete markets to model the commodity price uncertainty. We distinguish between market uncertainties, which can be hedged perfectly by trading market securities, and private uncertainties which are specific to the firm. Formally, we assume the period n state of information $\mathcal{I}_n$ can be written as a vector of market and private states of information $\mathcal{I}_n = (\mathcal{I}_n^m, \mathcal{I}_n^p)$. The market is *partially complete* if the following conditions are satisfied (Smith and Nau, 1995):

1. Security prices depend only on the market states and can be written as a function of the market state of uncertainty.

2. The market is *complete* with respect to market uncertainties; i.e., the security prices span the space of cashflows dependent only on the market states.

3. Private events convey no information about future market events; i.e., given $\mathcal{I}_{n-1}^m$, the firm believes that $\mathcal{I}_n^m$ and $\mathcal{I}_{n-1}^p$ are independent.

Conditions 1 and 2 above imply that there exist unique risk-neutral probabilities, $\pi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)$, on the market states $\mathcal{I}_{n+1}^m$ such that $\beta \sum_{\mathcal{I}_{n+1}^m} \pi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)\mathbf{M}_{n+1}(\mathcal{I}_{n+1}^m) = \mathbf{M}_n(\mathcal{I}_n^m)$ where $\mathbf{M}_n(\mathcal{I}_n^m) = (M_n(0, \mathcal{I}_n^m), \ldots, M_n(J, \mathcal{I}_n^m)) \geq \mathbf{0}$ is the vector of the $J + 1$ market security prices and $\beta = \frac{1}{1+r_f}$ is the risk-free discount rate. Also, let $p_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)$ denote the firm's subjective probabilities over the market states of uncertainty $\mathcal{I}_{n+1}^m$, given $\mathcal{I}_n^m$ and $q_n(\mathcal{I}_n^p, \mathcal{I}_n^m)$ denote the subjective probabilities over the private states of uncertainty in period n, given $\mathcal{I}_n^m$.

4.2.1 Dynamic Risk Measure

We extend the dynamic, time-consistent risk measure DCVaR defined in Chapter 3 to the partially complete market setting. Let $\mathbf{X} = (X_1(\mathcal{I}_1^m, \mathcal{I}_1^p), \ldots, X_N(\mathcal{I}_N^m, \mathcal{I}_N^p))$ be a stream of random cashflows over the horizon, with the cashflow in any specific period n dependent on the realized market and private uncertainties. We define the dynamic risk measure, $DCVaR$, as follows.

$$DCVaR_n(\mathbf{X}; \eta_n) = \begin{cases} CVaR_N^m\Big(CVaR_N^p\Big(X_N(\mathcal{I}_N^m, \mathcal{I}_N^p) \mid \mathcal{I}_N^m\Big)\Big) & \text{for } n = N \\ CVaR_n^m\Big(CVaR_n^p\Big(X_n(\mathcal{I}_n^m, \mathcal{I}_n^p) + \beta DCVaR_{n+1}(\mathbf{X}; \eta_{n+1}) \mid \mathcal{I}_n^m\Big)\Big) & \text{for } n < N \end{cases} \tag{4.1}$$

where $CVaR_n^p(\cdot|\mathcal{I}_n^m)$ denotes the conditional value-at-risk evaluated over the private states of uncertainty in period n, conditional on $\mathcal{I}_n^m$ and $CVaR^m(\cdot)$ is the conditional value-at-risk evaluated over the market states of uncertainty. Both $CVaR_n^p$

and $CVaR_n^m$ are evaluated at the level $0 < \eta_n \leq 1$, using the firm's subjective probabilities over the different market and private states of uncertainty. $DCVaR_n(\mathbf{X}; \eta_n)$ defined in equation (4.1) represents a conditional risk mapping (cf. Ruszczynski and Shapiro, 2006) and is time-consistent.

We end this section by proving an useful property of $DCVaR_n(\mathbf{X}; \eta_n)$, which will be used in the subsequent analysis.

Lemma 4.1. *For all $n < N$, the risk measure $DCVaR_{n+1}(\mathbf{X}; \eta_{n+1})$ is measurable with respect to the σ-algebra generated by the market states of uncertainty $\mathcal{I}_n^m$ and*

$$\begin{aligned} DCVaR_n(\mathbf{X}; \eta_n) \quad = \quad & CVaR_n^m \Bigg(CVaR_n^p \Big(X_n(\mathcal{I}_n^m, \mathcal{I}_n^p) \mid \mathcal{I}_n^m \Big) \\ & + \beta DCVaR_{n+1}(\mathbf{X}; \eta_{n+1}) \Bigg) \end{aligned} \tag{4.2}$$

In the next section, we consider the problem of a firm using market securities to hedge the uncertainty in a cashflow stream $\mathbf{X}$, with the objective of maximizing $DCVaR$ of the total cashflows.

4.2.2 Hedging using Market Securities

We consider a firm that owns a project which generates a cash flow stream $\mathbf{X}$, with the cash flow in any period n depending on the realized private and market uncertainty. The firm is interested in hedging the uncertainty in the cashflow stream by trading market securities. Let $\alpha_n = (\alpha_n(0), \ldots, \alpha_n(J))$ be the firm's position in the various market securities, after observing the state of market uncertainty in period n. Let $\mathcal{B} = (\alpha_1, \ldots, \alpha_N)$ be the trading policy for market securities.We can model the firm's problem as

$$\max_{\mathcal{B}} DCVaR_1 \left(\mathbf{X} + \mathbf{F}(\mathcal{B}); \eta_1 \right)$$

where $\mathbf{F}(\mathcal{B}) = (F_1(\mathcal{B}, \mathcal{I}_1^m), \ldots, F_N(\mathcal{B}, \mathcal{I}_N^m))$ is the stream of cash flows from trading the market securities. Specifically,

$$F_n(\mathcal{B}, \mathcal{I}_n^m) = [\alpha_{n-1} - \alpha_n]^{\mathrm{T}} \mathbf{M}_n(\mathcal{I}_n^m) \tag{4.3}$$

The firm's optimal trading policy can be determined by solving the following stochastic dynamic program.

$$\begin{aligned} V_n(\mathbf{X}, \alpha_{n-1}, \mathcal{I}_n^m) &= \max_{\alpha_n} \Bigg\{ [\alpha_{n-1} - \alpha_n]^{\mathrm{T}} \mathbf{M}_n(\mathcal{I}_n^m) \\ &\quad + CVaR_n^p\Big(X_n(\mathcal{I}_n^m, \mathcal{I}_n^p) \mid \mathcal{I}_n^m\Big) \\ &\quad + CVaR_n^m\Bigg(\beta V_{n+1}(\mathbf{X}, \alpha(n), \mathcal{I}_{n+1}^m)\Bigg)\Bigg\} \text{ for } n < N \end{aligned} \tag{4.4}$$

$$\begin{aligned} V_N(\mathbf{X}, \alpha_{n-1}, \mathcal{I}_N^m) &= (\alpha_{N-1})^{\mathrm{T}} \mathbf{M}_N(\mathcal{I}_N^m) \\ &\quad + CVaR_N^p\Big(X_N(\mathcal{I}_N^m, \mathcal{I}_N^p) \mid \mathcal{I}_N^m\Big) \end{aligned} \tag{4.5}$$

The next theorem states, under some mild restrictions on the firm's preferences and subjective probabilities, that the value function $V_n(\cdot)$ of the above maximization problem has a particularly simple form. Theorem 4.1 parallels the result obtained by Smith and Nau (1995), who impose the restriction that the firm's risk aversion is represented by an additive exponential utility function over net cashflows.

Theorem 4.1. *When the firm's subjective probabilities $p_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)$ over market uncertainty and η_n are such that $\frac{p_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)}{\eta_n} \geq \pi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)$ for each $(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)$ for all $n < N$, then*

$$\begin{aligned} V_n(\mathbf{X}, \alpha_{n-1}, \mathcal{I}_n^m) &= (\alpha_{n-1})^{\mathrm{T}} \mathbf{M}_n(\mathcal{I}_n^m) \\ &\quad + \mathbb{E}_{\pi_n}\left[\sum_{\tau=n}^{N} \beta^{\tau-n} CVaR_\tau^p\Big(X_\tau(\mathcal{I}_\tau^m, \mathcal{I}_\tau^p) \mid \mathcal{I}_\tau^m\Big)\right] \end{aligned} \tag{4.6}$$

where $\mathbb{E}_{\pi_n}[\cdot]$ *denotes expectation over* $\mathcal{I}^m_{n+1}$ *under the risk-neutral probability measure, conditional on* $\mathcal{I}^m_n$*.*

The above result implies that under an optimal financial trading policy, the firm's subjective probabilities over the states of market uncertainty does not affect the optimization. However, the firm's subjective probabilities over the private states affects the optimization through $CVaR^p_n(\cdot)$. Notice that for a risk-neutral firm, $\eta_n = 1$ for each n and the condition, $\frac{p_n(\mathcal{I}^m_{n+1}, \mathcal{I}^m_n)}{\eta_n} \geq \pi_n(\mathcal{I}^m_{n+1}, \mathcal{I}^m_n)$ reduces to $p_n(\mathcal{I}^m_{n+1}, \mathcal{I}^m_n) = \pi_n(\mathcal{I}^m_{n+1}, \mathcal{I}^m_n)$; i.e., the risk-neutral firm's subjective probabilities over market uncertainties coincide with the risk-neutral probabilities implied by the security prices.

Now, consider the situation where the project cashflows, $\mathbf{X}$, depend on an operational policy, γ, and the firm is interested in maximizing $DCVaR$ of total cash flows, by optimizing over the joint operational and financial policy. The above result implies that an operational policy γ^* is an optimal policy, if it maximizes $\mathbb{E}_{\pi_1}\left[\sum_{\tau=1}^{N} \beta^{\tau-1} CVaR^p_\tau\Big(X_\tau(\gamma_\tau, \mathcal{I}^m_\tau, \mathcal{I}^p_\tau) \mid \mathcal{I}^m_\tau\Big)\right]$. Subsequently, the optimal financial trading policy is the $CVaR^p_\tau\Big(X_\tau(\gamma^*_\tau, \mathcal{I}_\tau), \mathcal{I}^m_\tau\Big) \triangleq CVaR^p_\tau\Big(X_\tau(\gamma^*_\tau, \mathcal{I}^m_\tau, \mathcal{I}^p_\tau) \mid \mathcal{I}^m_\tau\Big)$ replicating portfolio. In the next section, we apply these results to a commodity processing firm's operations and analyze the optimal operational policies.

4.3. Model Description and Analysis

We consider a commodity processor who procures, processes and trades commodities over a finite horizon. The time periods are indexed by n, with $n = 1$ denoting the first period and $n = N$ the last period. In each period, the firm procures the input commodity from a spot market, where $S_n(\mathcal{I}_n)$ denotes the spot price in period n. While market instruments exist, and are used extensively, to manage commodity price uncertainty, not all of the uncertainty in the commodity spot price can be hedged using such instruments. The difference between the price of the market instrument (say, futures contract) and spot price is termed *basis risk* and exists because of factors such as timing, location and quality discrepancies between the physical commodity and the commodity underlying the market instrument (see Lapan and Moschini, 1994; Paroush and Wolf, 1989; Moschini and Lapan, 1995, for instance). For this reason, we model the commodity spot price as a function of both the market and private uncertainty.

The firm earns revenues by processing the input and selling the output commodity (processed product) using forward contracts of different maturities. The forward price on contract ℓ is denoted by $F_n^\ell(\mathcal{I}_n^m)$, where N_ℓ, the maturity period of contract ℓ, is greater than n. We use $\mathbf{F}_n(\mathcal{I}_n^m) = (F_n^\ell(\mathcal{I}_n^m), \ldots, F_n^L(\mathcal{I}_n^m))$ to denote the vector of forward prices of all contracts yet to mature as of period n. We assume the forward prices are pegged to the price of actively traded futures instruments on the output commodity and therefore, the forward prices depend only on the state of market uncertainty. Output quality differences, i.e., difference between the quality of output produced and quality of the output underlying the futures instrument, can be accounted for by assuming the firm only produces output of the required quality. Further, the uncertainty in output quality can be incorporated by modeling the input spot price as the per-unit price of input required to produce a unit of the output of the specified quality and modeling the input spot price to depend on market and private uncertainties.

The delivery period for all commitments made against contract ℓ is N_ℓ and period $N_\ell - 1$ is the last period in which the firm can commit to sell the output commodity using the forward contract ℓ. In addition to the output commodity sales, the firm can also earn revenues by trading the input commodity with other processors. For ease of exposition, we assume that all input commodity trading occurs at the end of the horizon, at the trade (salvage) price of $S_N(\mathcal{I}_N)$. Let $\mathcal{I}_n = (\mathcal{I}_n^m, \mathcal{I}_n^p)$ denote the relevant information available to the firm at the beginning of period n regarding the various commodity prices.

On the operational side, the firm has a per-period procurement and processing capacity restriction of K and C units respectively. The firm incurs a variable cost of p to process one unit of input into the output commodity. For simplicity, we assume all physical holding costs for the various commodities are negligible.

At the beginning of each period n, the firm observes the input spot price, S_n, and the output forward prices, $\mathbf{F}_n$, for the period. Let e_n and Q_n denote the input and output commodity inventories respectively at the beginning of period n. Let $\mathbf{R}_n = (R_n^\ell, \ldots, R_n^L)$ denote the vector of cumulative commitments against all forward contracts yet to mature. Based on this information, the firm makes the following decisions in each period: 1) the quantity of input commodity to procure, x_n, 2) the quantity to process, m_n and 3) the quantity of the output commodity to commit to sale against the forward contracts, $\mathbf{q}_n = (q_n^\ell, \ldots, q_n^L)$.

The procurement and processing decisions in any period are subject to capacity and inventory availability constraints and the feasible set of actions in period n is given by $\mathcal{A}_n(e_n)$ where

$$\mathcal{A}_n(e_n) = \{(x_n, m_n) : 0 \le x_n \le K,\ 0 \le m_n \le \min\{C, e_n + x_n\}\} \tag{4.7}$$

The output commodity sale commitments are not reversible; i.e., $\mathbf{q}_n \ge \mathbf{0}$. It is not necessary that $\sum_{l=\ell}^{L} R_n^l + q_n^l \le Q_n + m_n$, as the firm can commit to sell more output than is available on-hand as long as all the output committed for sale against a forward contract is delivered on the delivery date specified in the forward contract. That is, for each ℓ we require $R_{N_\ell-1}^\ell + q_{N_\ell-1}^\ell \le Q_{N_\ell-1} + m_{N_\ell-1}$. We denote the set of feasible commitment vectors in period n by $\mathcal{Q}(\mathbf{R}_n, Q_n, m_n)$.

The state transition equations are given by

$$\begin{aligned} e_{n+1} &= e_n + x_n - m_n & (4.8)\\ \mathbf{R}_{n+1} &= \mathbf{R}_n + \mathbf{q}_n & (4.9)\\ Q_{n+1} &= \begin{cases} Q_n + m_n & n \neq N_\ell - 1 \text{ for any } \ell \\ Q_n + m_n - R_n^\ell - q_n^\ell & n = N_\ell - 1 \end{cases} & (4.10) \end{aligned}$$

The profits realized by the firm in period n are given by

$$\Pi_n(x_n, m_n, \mathbf{q}_n, \mathcal{I}_n) = \begin{cases} \sum_{l=\ell}^{L} \beta^{N_l - n} F_n^l q_n^l - S_n x_n - p m_n & n < N \\ S_N e_N & n = N \end{cases} \quad (4.11)$$

where β is the risk-free discount factor.

The profit function in equation (4.11) above recognizes revenues from output sales at the time of commitment rather than at delivery. Since commodity sale commitments are not reversible and we assume no counter party risk is present, recognizing revenue at the time of commitment rather than at delivery is without loss of generality.

We model the risk-averse firm's objective function by the time-consistent risk measure $DCVaR(\cdot;\eta)$, defined in equation (4.1), on the stream of discounted cashflows. To keep the exposition simple, we assume that the levels η are the same across all periods (we also suppress η in the following, in the interests of keeping the notation manageable). Let $\mathbf{CF}_n^N = \Big(\pi_n(\cdot)+[\alpha_{n-1}-\alpha_n]^{\mathrm{T}}\mathbf{M}_n(\cdot), \ldots, \pi_N(\cdot)+(\alpha_{N-1})^{\mathrm{T}}\mathbf{M}_N(\cdot)\Big)$ and $\mathcal{B}_n^N = (\alpha_n, \ldots, \alpha_{N-1})$. Using the short hand notation $DCVaR_n^\eta(\cdot) \triangleq DCVaR_n(\cdot\,;\eta_n)$, the firm's joint operational and financial risk management problem in period n can then be written as

$$V_n(e_n, Q_n, \mathbf{R}_n, \alpha_{n-1}, \mathcal{I}_n) = \max_{\substack{(x_\tau, m_\tau)\in\mathcal{A}_\tau(e_\tau)\\ q_\tau\in\mathcal{Q}_\tau(\mathbf{R}_\tau, Q_\tau, m_\tau)\\ \tau=n,\ldots,N-1}} \left\{ \max_{\mathcal{B}_n^N} \left\{ DCVaR_n^\eta\Big(\mathbf{CF}_n^N\Big) \right\} \right\} \quad (4.12)$$

Assuming the conditions of Theorem 4.1 hold and assuming the firm uses the optimal financial trading policy, we can use the result therein to determine the optimal operational policy by solving the following stochastic dynamic program

$$\begin{aligned} V_n(e_n, Q_n, \mathbf{R}_n, \mathcal{I}_n) &= \max_{\substack{(x_n, m_n) \in \mathcal{A}_n(e_n), \\ \mathbf{q}_n \in \mathcal{Q}_n(\mathbf{R}_n, Q_n, m_n)}} \Bigg\{ \Pi_n(x_n, m_n, \mathbf{q}_n, \mathcal{I}_n) + \\ &\qquad \beta \mathbb{E}_{\pi_n} \left[CVaR^p_{n+1}\Big(V_{n+1}(\cdot), \mathcal{I}^m_{n+1}\Big)\right] \Bigg\} \quad (4.13) \\ &\qquad \text{for } n = 1, \ldots, N-1 \text{ and} \\ V_N(e_N, Q_N, \mathbf{R}_N, \mathcal{I}_N) &= S_N e_N \quad (4.14) \end{aligned}$$

Notice that the objective function in equation (4.13) is very similar to the objective function in SDP formulations for expected value maximization. However, there are crucial differences. Specifically, the expectation is taken with respect to the risk-neutral probability measure and only over the states of market uncertainty. Further, the argument for the expectation operator is $CVaR^p_{n+1}\Big(V_{n+1}(\cdot), \mathcal{I}^m_{n+1}\Big)$, the CVaR of V_{n+1} evaluated over the states of private uncertainty for each $\mathcal{I}^m_{n+1}$.

In the next section, we derive the structure of the optimal operational policy. Later, we use these results to quantify the value from operational flexibility.

4.4. Optimal Operational Policy

4.4.1 Optimal Commitment Policy

4.4.1.1 Single forward contract.

We start our analysis by considering the case when only a single forward contract is available to sell the output. This would be the situation for n such that $N_{L-1} \leq n < N_L$. Without loss of generality, we assume $N_L = N$. In period $N-1$, the firm solves the following optimization problem to determine the optimal actions

$$
\begin{aligned}
V_{N-1}(e_{N-1}, Q_{N-1}, R^L_{N-1}, \mathcal{I}_{N-1}) \quad = \quad & \max\Big\{\pi_{N-1}(x_{N-1}, m_{N-1}, q^L_{N-1}) + \\
& \beta\mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_N\Big((S_N \times e_N), \mathcal{I}^m_N\Big)\Big]\Big\} \\
\text{s.t.} \\
& 0 \le q^L_{N-1} \le Q_{N-1} + m_{N-1} - R^L_{N-1} \\
& 0 \le e_N = e_{N-1} + x_{N-1} - m_{N-1} \\
& 0 \le m_{N-1} \le C \\
& 0 \le x_{N-1} \le K
\end{aligned}
$$

It is optimal to have $q^L_{N-1} = Q_{N-1} + m_{N-1} - R^L_{N-1}$, as any left over output inventory does not earn any revenue in the next period. Substituting the optimal commitment quantity, we can write

$$
V_{N-1}(e_{N-1}, Q_{N-1}, C^L_{N-1}, \mathcal{I}_{N-1}) \quad = \quad = \beta F^L_{N-1}(Q_{N-1} - R^L_{N-1}) + U_{N-1}\Big(e_{N-1}, \mathcal{I}_{N-1}\Big)
$$

where

$$
\begin{aligned}
U_{N-1}\Big(e_{N-1}, \mathcal{I}_{N-1}\Big) \quad = \quad & \max\Big\{(\beta F^L_{N-1} - p)m_{N-1} - S_{N-1}x_{N-1} \\
& +\beta\mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_N\Big(S_N, \mathcal{I}^m_N\Big)\Big] \times e_N\Big\} \\
\text{s.t.} \\
& 0 \le e_N = e_{N-1} + x_{N-1} - m_{N-1} \\
& 0 \le m_{N-1} \le C \\
& 0 \le x_{N-1} \le K
\end{aligned}
$$

As $U_{N-1}(\cdot)$ is the solution of a linear program where e_{N-1} is on the right hand side of the constraints, it is piecewise linear and concave in e_{N-1} for each realization of

$\mathcal{I}_{N-1}$. Also notice that V_{N-1} is separable in Q_{N-1} and e_{N-1} and furthermore, linear in Q_{N-1} with the coefficient of Q_{N-1} dependent only on $\mathcal{I}^m_{N-1}$.

We have $\beta^{N-(N-2)} F^L_{N-2}(\mathcal{I}^m_{N-2}) = \beta \mathbb{E}_{\pi_{N-2}} \left[\beta^{N-(N-1)} F^L_{N-1}(\mathcal{I}^m_{N-1}) \right]$ for each $\mathcal{I}^m_{N-2}$; i.e., the revenue from making a commitment to sell the output in period $N-2$ is equal to the discounted expected revenue from postponing and making the commitment in period $N-1$. Thus, it is without loss of optimality to postpone all commitments against contract L to period $N-1$. Thus, under an optimal commitment policy we would have $R^L_n = 0$ for all $n < N$. Using an induction argument, we can show that

$$V_n(Q_n, e_n, \mathcal{I}_n) = \Delta^m_n Q_n + U_n(e_n, \mathcal{I}_n) \text{ for } N_{L-1} \leq n \leq N$$

where $\Delta^m_n = \beta^{N-n} F^L_n(\mathcal{I}^m_n)$ with the superscript m indicating that it depends only on $\mathcal{I}^m_n$. Further, $U_n(\cdot)$ is piecewise linear and concave in e_n for each $\mathcal{I}_n$.

4.4.1.2 Multiple forward contracts.

Now consider the commitment decision at $n = N_{L-1} - 1$. The firm can commit against contract $L-1$ and earn a revenue of $\beta FL - 1_{N_{L-1}-1}$. By not committing, the risk-adjusted value of uncommitted inventory is $\beta \mathbb{E}_{\pi_{N_{L-1}-1}} \left[\Delta^m_{N_{L-1}} \right]$. Thus, the optimal commitment quantity against $L-1$ is equal to $Q_{N_{L-1}-1} + m_{N_{L-1}-1} - R^{L-1}_{N_{L-1}-1}$ if $F^{L-1}_{N_{L-1}-1} \geq \mathbb{E}_{\pi_{N_{L-1}-1}} \left[\Delta^m_{N_{L-1}} \right]$ and zero otherwise. Thus, we have

$$V_n(Q_n, e_n, \mathcal{I}_n) = \beta \max \left\{ F^{L-1}_n, \mathbb{E}_{\pi_n} \left[\Delta^m_{n+1} \right] \right\} (Q_n - R^{L-1}_n) + U_n(e_n, \mathcal{I}_n)$$

for $n = N_{L-1} - 1$.

For any $n < N_{L-1} - 1$, the benefit from committing against contract $L-1$ is $\beta^{N_{L-1}-n} F^{L-1}_n$. However, the benefit from not committing and postponing the commitment decision to period $N_{L-1} - 1$ is $\beta^{N_{L-1}-1-n} \mathbb{E}_{\pi_n} \left[\beta \max \left\{ F^{L-1}_n, \mathbb{E}_{\pi_n} \left[\Delta^m_{n+1} \right] \right\} \right]$ which is greater than $\beta^{N_{L-1}-n} F^{L-1}_n$ because the markets are complete with respect to market uncertainties and the forward prices are a martingale under the risk-neutral probabilities. Using induction, we can prove

Lemma 4.2. *It is optimal to postpone commitment against any specific contract ℓ to period $N_\ell - 1$ and the optimal commitment decision in period $N_\ell - 1$ is given by*

$$q^{(\ell *)}_{N_\ell - 1} = \begin{cases} Q_{N_\ell - 1} + m_{N_\ell - 1} - R^\ell_1 & \text{if } F^\ell_{N_\ell - 1} \geq \mathbb{E}_{\pi_{N_\ell - 1}}\left[\Delta^m_{N_\ell}\right] \\ 0 & \text{otherwise} \end{cases} \tag{4.15}$$

where

$$\Delta^m_n = \begin{cases} 0 & n = N \\ \beta \max\{F^\ell_n, \mathbb{E}_{\pi_n}\left[\Delta^m_{n+1}\right]\} & \text{if } n = N_\ell - 1 \text{ for } \ell = 1, \ldots, L \\ \beta \mathbb{E}_{\pi_n}\left[\Delta^m_{n+1}\right] & \text{otherwise} \end{cases} \tag{4.16}$$

Under an optimal commitment policy, $R^\ell_n = R^\ell_1$ for all $n < N_\ell$ and it does not affect the value function. Therefore, we do not need to keep track of the vector $\mathbf{C}_n$ as part of the state variable. As a consequence of the optimal commitment policy, we can write

$$V_n(Q_n, e_n, \mathcal{I}_n) = \Delta^m_n Q_n + U_n(e_n, \mathcal{I}_n) \tag{4.17}$$

for all $n \leq N$, where

$$U_n(e_n, \mathcal{I}_n) = \max_{\substack{0 \leq m_n \leq \min\{C, e_n + x_n\}, \\ 0 \leq x_n \leq K}} \Big\{ (\Delta^m_n - p) m_n - S_n x_n + \beta \mathbb{E}_{\pi_n}\Big[CVaR^p_{n+1}\Big(U_{n+1}(\cdot), \mathcal{I}^m_{n+1} \Big) \Big] \Big\} \tag{4.18}$$

for $n < N$ and $U_N(e_N, \mathcal{I}_N) = S_N e_N$.

4.4.2 Optimal Procurement and Processing Policy

Using an induction argument, similar to the one used in Chapter 3, we can prove that $U_n(e_n, \mathcal{I}_n)$ defined in equation (4.18) is continuous, concave and piecewise linear in e_n for each realization of $\mathcal{I}_n$, for all $n \leq N$. Consequently, we can also prove that $\mathcal{H}_n(e_{n+1}, \mathcal{I}_n) \triangleq \beta \mathbb{E}_{\pi_n}\left[CVaR^p_{n+1}\Big(U_{n+1}(e_{n+1}, \mathcal{I}_{n+1}), \mathcal{I}^m_{n+1}\Big)\right]$ is concave and piecewise linear in e_{n+1} for all $n < N$.

Let $0 = b_n(0) < b_n(1) < \ldots < b_n(k) < b_n(k+1) < \ldots < b_n(\kappa_n) < \infty$ denote the break points for $\mathcal{H}_n(e_{n+1}, \mathcal{I}_n)$; these are the points at which there is a change in slope of $\mathcal{H}_n(e_{n+1}, \mathcal{I}_n)$ (to keep the notation simple, we do not show the dependence of the breakpoints $b_n(k)$ on $\mathcal{I}_n$, but the reader should be aware of this dependence and that the breakpoints are not necessarily the same for different realizations of $\mathcal{I}_n$). As the number of possible price realizations are finite in each period, we can use an induction argument to prove that both the number of break points $\kappa_n + 1$ as well as the magnitude $b_n(\kappa_n)$ are finite.

For $k = 1, \ldots, \kappa_n$ let $g_n^{(k)} = b_n(k) - b_n(k-1)$ and $\Upsilon_n^{(k)} = \frac{\mathcal{H}_n(b_n(k), \mathcal{I}_n) - \mathcal{H}_n(b_n(k-1), \mathcal{I}_n)}{g_n^{(k)}}$ and $\Upsilon_n^{(\kappa_n+1)}$ is the slope of $\mathcal{H}_n$ for $e_{n+1} > b_n(\kappa_n)$. By concavity of $\mathcal{H}_n$, we have $\Upsilon_n^{(k+1)} < \Upsilon_n^{(k)}$ for all $k \leq \kappa_n$. Using arguments similar to those in Chapter 3 we can prove that

Proposition 4.1. *In any period n, for a realization $\mathcal{I}_n$ of the prices, there exist two input inventory levels, $\underline{b_n}$ and $\overline{b_n}$ such that*

$$\underline{b_n} = \begin{cases} b_n(k) & \text{if } \exists\ k \leq \kappa_n \text{ s.t. } \Upsilon_n^{(k)} > \Delta_n^m - p \geq \Upsilon_n^{(k)} \\ 0 & \text{if } \Upsilon_n^{(1)} \leq \Delta^m - p \\ \infty & \text{if } \Upsilon_n^{(\kappa_n+1)} > \Delta_n^m - p \end{cases} \tag{4.19}$$

$$\overline{b_n} = \begin{cases} b_n(k) & \text{if } \exists\ k \leq \kappa_n \text{ s.t. } \Upsilon_n^{(k)} > S_n \geq \Upsilon_n^{(k+1)} \\ 0 & \text{if } \Upsilon_n^{(1)} \leq S_n \\ \infty & \text{if } \Upsilon_n^{(\kappa_n+1)} > S_n \end{cases} \tag{4.20}$$

The optimal procurement and processing quantities (x_n^*, m_n^*) *are then given by*

$$x_n^* = \begin{cases} \min\{K, (\overline{b_n} - e_n)^+\} & \text{if } \overline{b_n} \leq \underline{b_n} \\ \min\{K, (\overline{b_n} + C - e_n)^+\} & \text{if } \overline{b_n} > \underline{b_n} \end{cases} \tag{4.21}$$

$$m_n^* = \begin{cases} \min\{C, (e_n - \underline{b_n})^+\} & \text{if } \overline{b_n} \leq \underline{b_n} \\ \min\{C, (e_n + K - \underline{b_n})^+\} & \text{if } \overline{b_n} > \underline{b_n} \end{cases} \tag{4.22}$$

4.5. Operational Hedging

As seen from the analysis in the previous section, trading in the financial markets helps the firm hedge the uncertainty in revenues from output sales and the firm's risk aversion does not affect the value of output inventory. However, financial trading does not help the firm hedge against private uncertainties that affect input prices, and therefore it also needs operational levers such as excess procurement or processing capacity to manage the input price uncertainty. In this section, we explore how the firm's choice of procurement and processing capacities, affects the value from operations.

As the firm cannot completely hedge the uncertainty in input commodity prices, it is more likely that the firm will process and carry output inventory rather than have unprocessed input inventory. However, having excess procurement capacity relative to processing capacity allows the firm to opportunistically procure more in periods when the realized input price is sufficiently low and process in later periods. To determine the value of excess procurement capacity, we first consider the case where the firm has equal procurement and processing capacity.

4.5.1 Equal Procurement and Processing Capacity

With identical procurement and processing capacities (i.e., $K = C$), consider the firm's problem in period $N-1$. We have

$$U_{N-1}(e_{N-1}, \mathcal{I}_{N-1}) = \max_{\substack{0 \le m \le \min\{C, e_{N-1}+x\}, \\ 0 \le x \le C}} \Big\{ (\Delta^m_{N-1} - p)m - S_{N-1}x + \beta \mathbb{E}_{\pi_{N-1}} \left[CVaR^p_N(S_N, \mathcal{I}^m_N) \right] (e_{N-1} + x - m) \Big\}$$

Suppose $\beta \mathbb{E}_{\pi_{N-1}} \left[CVaR^p_N(S_N, \mathcal{I}^m_N) \right] < \Delta^m_{N-1} - p$. Then, it is optimal to process all available input inventory, up to the processing capacity; i.e., $m^* = \min\{C, e_{N-1} + x\}$. The optimal procurement quantity is then given by

$$x^* = \begin{cases} C & \text{if } S_{N-1} \le \beta \mathbb{E}_{\pi_{N-1}} \left[CVaR^p_N(S_N, \mathcal{I}^m_N) \right] \\ (C - e_{N-1})^+ & \text{if } \beta \mathbb{E}_{\pi_{N-1}} \left[CVaR^p_N(S_N, \mathcal{I}^m_N) \right] < S_{N-1} \le \Delta^m_{N-1} - p \\ 0 & \text{if } \Delta^m_{N-1} - p < S_{N-1} \end{cases}$$

Substituting the optimal procurement and processing quantities, we get

$$U_{N-1}(e_{N-1}, \mathcal{I}_{N-1}) = \begin{cases} \max\Big\{ \beta \mathbb{E}_{\pi_{N-1}} \left[CVaR^p_N(S_N, \mathcal{I}^m_N) \right], \min\{S_{N-1}, \Delta^m_{N-1} - p\} \Big\} e_{N-1} + \Lambda^1_{N-1} & e_{N-1} \in [0, C] \\ \beta \mathbb{E}_{\pi_{N-1}} \left[CVaR^p_N(S_N, \mathcal{I}^m_N) \right] e_{N-1} + \Lambda^2_{N-1} & e_{N-1} > C \end{cases}$$

where Λ^1_{N-1} and Λ^2_{N-1} denote constant terms independent of e_{N-1}.

Notice that the marginal, risk-adjusted value of input inventory $\frac{\partial U_{N-1}}{\partial e_{N-1}}$ is less than $\Delta^m_{N-1} - p$ for all e_{N-1}. The next theorem states that this is true for all $n < N$, provided the firm's subjective probabilities on the salvage value S_N satisfy the condition that $CVaR^p_N(S_N, \mathcal{I}^m_N) < F^L_N - p$ for each $\mathcal{I}^m_N$.

Theorem 4.2. *Let the firm's subjective probabilities on salvage value S_N be such that*

$$CVaR_N^p(S_N, \mathcal{I}_N^m) < F_N^L - p \quad \textit{for each } \mathcal{I}_N^m \tag{4.23}$$

Then, the marginal risk-adjusted value of input inventory in any period $n < N$ is no more than $\Delta_n^m - p$ for all $e_n \geq 0$, for each realization of $\mathcal{I}_n$ and $m_n^* = \min\{C, e_n + x_n^*\}$.

The condition implied by equation (4.23) says that the expected salvage value in the worst η fraction of cases is less than the value from processing the input and selling the output, for a given realization of market uncertainty. This is a reasonable assumption for a risk-averse firm that would prefer the certain revenue from processing and selling the output rather than the possibly higher but uncertain revenue from salvaging. In fact, for a risk-averse firm it is reasonable to assume that $\mathbb{E}_{q_N}[S_N] < F_N^L - p$; i.e., the expected salvage value based on the firm's subjective probabilities is no more than the value from processing and selling the output. As $CVaR_N^p(S_N) \leq \mathbb{E}_{q_N}[S_N]$ for all $\eta \in (0, 1]$, the condition in equation (4.23) is less restrictive.

A consequence of the above result is that the firm prefers to process any input available, and carries input inventory into the next period only when constrained by the processing capacity. Thus, when $K = C$ and $e_1 = 0$, we will always have $e_n = 0$ for $n > 1$ under an optimal processing policy and the value function can be written as

$$\begin{aligned} U_n(0, \mathcal{I}_n) &= (\Delta_n^m - p - S_n)^+ \times C \\ &\qquad + \mathbb{E}_{\pi_n}\left[\sum_{\tau=n+1}^{N-1} \beta^{\tau-n} CVaR_\tau^p\left((\Delta_\tau^m - p - S_\tau)^+, \mathcal{I}_\tau^m\right)\right] \times C \\ &= (\Delta_n^m - p - S_n)^+ \times C \\ &\qquad + \beta\mathbb{E}_{\pi_n}\left[CVaR_{n+1}^p\Big(U_{n+1}(0, \mathcal{I}_{n+1}), \mathcal{I}_{n+1}^m\Big)\right] \end{aligned} \tag{4.24}$$

Notice that the analysis remains the same for all $K < C$, with the only change that C in equation (4.24) is replaced by K; i.e., any processing capacity in excess of the procurement capacity does not provide additional value to the firm. This is because the firm's preferred action in each period is to process all available input inventory, up to the processing capacity. When $K < C$, the processing capacity constraint is never binding. However, excess procurement capacity can be valuable to procure input in periods when the realized spot price is sufficiently low. We now consider the value from operations when the firm has excess procurement capacity available.

4.5.2 Excess Procurement Capacity

We consider the scenario where the firm has excess procurement capacity relative to the processing capacity. Specifically, the procurement capacity in each period, K, is such that $K > C$.

Consider period $T-1$. From the analysis in the earlier section, we have $m^*_{N-1} = \min\{e_{N-1} + x_{N-1}, C\}$ for any procurement quantity x in period $T-1$ and

$$U_{N-1}(e_{N-1}, \mathcal{I}_{N-1}) = \max_{0 \le x \le K} \Big\{ (\Delta^m_{N-1} - p)m^*_{N-1} - S_{N-1}x + \beta \mathbb{E}_{\pi_{N-1}} \Big[CVaR^p_N \Big(S_N, \mathcal{I}^m_N \Big) \Big] \times (e_{N-1} + x - m^*_{N-1}) \Big\}$$

It is easy to see that the optimal procurement quantity, x^*_{N-1} is given by

$$x^*_{N-1} = \begin{cases} K & \text{if } S_{N-1} \le \beta \mathbb{E}_{\pi_{N-1}} \Big[CVaR^p_N \Big(S_N, \mathcal{I}^m_N \Big) \Big] \\ (C - e_{N-1})^+ & \text{if } \beta \mathbb{E}_{\pi_{N-1}} \Big[CVaR^p_N \Big(S_N, \mathcal{I}^m_N \Big) \Big] < S_{N-1} \le \Delta^m_{N-1} - p \\ 0 & \text{if } \Delta^m_{N-1} - p < S_{N-1} \end{cases}$$

Substituting the optimal procurement quantity above, we get

$$\begin{aligned} U_{N-1}(e_{N-1},\mathcal{I}_{N-1}) &= \beta\mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_N\Big(S_N,\mathcal{I}^m_N\Big)\Big]e_{N-1} \\ &+\Big(\Delta^m_{N-1}-p-\beta\mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_N\Big(S_N,\mathcal{I}^m_N\Big)\Big]\Big)C \\ &+\Big(\beta\mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_N\Big(S_N,\mathcal{I}^m_N\Big)\Big]-S_{N-1}\Big)^+K \quad (4.25)\end{aligned}$$

when $e_{N-1} \geq C$ and

$$\begin{aligned} U_{N-1}(e_{N-1},\mathcal{I}_{N-1}) &= \max\Big\{\beta\mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_N\Big(S_N,\mathcal{I}^m_N\Big)\Big], \\ &\qquad \min\big\{S_n,\Delta^m_{N-1}-p\big\}\Big\}e_{N-1} \\ &+(\Delta^m_{N-1}-p-S_{N-1})^+C \\ &-\Big(\beta\mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_N\Big(S_N,\mathcal{I}^m_N\Big)\Big]-S_{N-1}\Big)^+C \\ &+\Big(\beta\mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_N\Big(S_N,\mathcal{I}^m_N\Big)\Big]-S_{N-1}\Big)^+K \quad (4.26)\end{aligned}$$

when $0 \leq e_{N-1} < C$.

Notice that U_{N-1} is piecewise linear and concave in e_{N-1} and the marginal value of input inventory is never more than $\Delta^m_{N-1} - p$ for a given realization of $\mathcal{I}^m_{N-1}$. The next theorem shows that these properties are true for any general n and further-more, the breakpoints of the value function are integral multiples of the greatest common divisor of the procurement and processing capacities.

Theorem 4.3. *Let the firm's subjective probabilities over the salvage values satisfy equation (4.23) and D denote the greatest common divisor of the procurement and processing capacities. Then, the value function $U_n(e_n, \mathcal{I}_n)$ is piecewise linear, concave and continuous in e_n with break points at integral multiples of D.*

Let the marginal risk-adjusted value of input inventory in period t when $e_n \in$

$[(k-1)D, kD)$ *for* $k = 1, 2, \ldots$ *be denoted by* $\Theta_n^{(k)}$*. Then,* $\Theta_N^{(k)} = S_N$ *and*

$$\Theta_n^{(k)} = \max\left\{\Upsilon_{n+1}^{(k+b-a)}, \min\left\{\Delta_n^m - p, \Upsilon_{n+1}^{(k-a)}\right\}\right\} \text{ for } n < N \text{ and } k \geq 1 \tag{4.27}$$

where a *and* b *are integers such that* $C = aD$ *and* $K = bD$ *with* $b > a$*, and*

$$\Upsilon_n^{(k)} = \begin{cases} \beta \mathbb{E}_{\pi_n}\left[CVaR_{n+1}^p\Big(\Theta_{n+1}^{(k)}, \mathcal{I}_{n+1}^m\Big)\right] & \text{for } n = N-1 \text{ and } k \geq 1 \\ \\ \beta \mathbb{E}_{\pi_n}\left[-CVaR_{n+1}^p\Big(-\Theta_{n+1}^{(k)}, \mathcal{I}_{n+1}^m\Big)\right] & \text{for } n < N-1 \text{ and } k \geq 1 \end{cases} \tag{4.28}$$

with $\Upsilon_n^{(k)} \triangleq \infty$ *for* $k \leq 0$*, for all* n*.*

Theorem 4.3 provides a way to recursively determine the total risk-adjusted value from operations for a given set of procurement and processing capacities. We can use this to determine the benefit from having excess procurement capacity. More specifically, let $\nu_n(K, C)$ denote the value of the excess procurement capacity, when the excess capacity is available from periods n through $N-1$. (Clearly, ν_n also depends on $\mathcal{I}_n$. In order to keep the notation simple, we do not explicitly show this dependency.)

In periods $1, \ldots, n-1$, the procurement capacity is equal to C. From the analysis in Section 4.5.1, we know that $e_n = 0$ under an optimal policy. Therefore,

$$\nu_n(K, C) = U_n(0, \mathcal{I}_n) - \hat{U}_n(0, \mathcal{I}_n) \tag{4.29}$$

where $\hat{U}_n(0, \mathcal{I}_n)$ is given by equation (4.24). Proposition 4.2 provides a recursive expression for ν_n.

Proposition 4.2. *The risk-adjusted value of having excess procurement capacity* $(K-$

C) from periods n through $N-1$, where $K > C$, is equal to

$$\nu_n(K,C) = \beta \mathbb{E}_{\pi_n}\left[CVaR^p_{n+1}\Big(\nu_{n+1}, \mathcal{I}^m_{n+1}\Big)\right] + \sum_{j=1}^{b-a}\left(\Upsilon^{(j)}_n - S_n\right)^+ D \tag{4.30}$$

with $\nu_N(K,C) \triangleq 0$.

4.6. Conclusions

In this chapter, we considered the dynamic financial and operational decisions for a commodity processing firm operating in a partially complete financial market. We extended the time-consistent risk measure introduced in Chapter 3 to the partially complete market framework and characterized the optimal financial trading and operational policies. Specifically, we showed that the optimal financial portfolio replicates the $CVaR$ over states of private uncertainty of the operational cashflows for each market state. We also showed that the optimal output commitment policy is identical to the optimal commitment policy for a risk-neutral firm, which is a consequence of the fact that the uncertainty in revenue from output sales depend only on market uncertainties. Similar to the risk-neutral case, the optimal procurement and processing decisions in any period are governed by 'procure up to' and 'process down to' thresholds. However, unlike the risk-neutral case, these thresholds are also dependent on the firm's risk aversion and subjective probabilities over the states of private uncertainty. Under a mild restriction on the salvage value of input inventory at the end of the horizon, we showed that excess processing capacity does not provide any benefit, while excess procurement capacity provides an additional lever for the firm to manage input price uncertainty.

This work extends our analysis in Chapter 3 and provides additional insights into the value of financial and operational hedging for a commodity processing firm in a dynamic setting. While we characterize the value of operational hedging analytically, we did not perform any comparative statics on how this benefit varies as a function of factors such as degree of risk aversion, horizon length or amount of additional capacity. It will be useful to explore these analytically and/or numerically as part of future research. An analytical characterization of the comparative statics will presumably require more specific assumptions on the dynamics of how the market uncertainties evolve over time. It would also be worthwhile to consider the other extensions suggested in Chapter 3, e.g., multiple input / output commodities, time consistent risk constraints, under a partially complete market setting and quantify the value of financial and operational hedging.

4.7. Appendix: Proofs of Theorems and Lemmas

Proof of Lemma 4.1. Notice that $CVaR_N^m\left(CVaR_N^p\left(X_N(\mathcal{I}_N^m, \mathcal{I}_N^p) \mid \mathcal{I}_N^m\right)\right)$ is measurable with respect to $\mathcal{I}_{N-1}^m$, because $\mathcal{I}_{N-1}^p$, the state of private uncertainty in period $N-1$, does not convey any information about future market uncertainty and the probability distribution, $p_{N-1}(\mathcal{I}_N^m)$, over market states of uncertainty in period N is completely determined by $\mathcal{I}_{N-1}^m$. Suppose $DCVaR^{n+1}(\mathbf{X}; \eta_{n+1})$ is measurable with respect to $\mathcal{I}_n^m$. Then, from equation (4.1) we have

$$\begin{aligned} DCVaR_n(\mathbf{X}; \eta_n) &= CVaR_n^m\left(CVaR_n^p\left(X_n(\mathcal{I}_n) + \beta DCVaR_{n+1}(\mathbf{X}; \eta_{n+1}) \mid \mathcal{I}_n^m\right)\right) \\ &= CVaR_n^m\left(CVaR_n^p\left(X_n(\mathcal{I}_n) \mid \mathcal{I}_n^m\right) + \beta DCVaR_{n+1}(\mathbf{X}; \eta_{n+1})\right) \end{aligned}$$

where the second equality follows from the fact that $CVaR$ is a coherent risk measure and satisfies the property of translation invariance. The right hand side expression in the second equality above is measurable with respect to $\mathcal{I}_{n-1}^m$, thus completing the proof. □

Proof of Theorem 4.1. The theorem is clearly true for period N. Suppose it is true for periods $n+1, \ldots, N$. Let

$$\mathcal{C}_{n+1}(\mathcal{I}_{n+1}^m) \triangleq \mathbb{E}_\pi\left[\sum_{\tau=n+1}^{T} \beta^{\tau-(n+1)} CVaR_\tau^p\left(X_\tau(\mathcal{I}_\tau^m, \mathcal{I}_\tau^p) \mid \mathcal{I}_\tau^m\right) \middle| \mathcal{I}_{n+1}^m\right]$$

From equation (4.4), we have

$$
\begin{aligned}
V_n(\mathbf{X}, \alpha_{n-1}, \mathcal{I}_n^m) &= \max_{\alpha_n} \Big\{ [\alpha_{n-1} - \alpha_n]^{\mathrm{T}} \mathbf{M}_n(\mathcal{I}_n^m) + CVaR_n^p \Big(X_n(\mathcal{I}_n^m, \mathcal{I}_n^p) \mid \mathcal{I}_n^m \Big) \\
&\quad + CVaR_n^m \Big(\beta \alpha_n^{\mathrm{T}} \mathbf{M}_{n+1}(\mathcal{I}_{n+1}^m) + \beta \mathcal{C}_{n+1}(\mathcal{I}_{n+1}^m) \Big) \Big\} \\
&= (\alpha_{n-1})^{\mathrm{T}} \mathbf{M}_n(\mathcal{I}_n^m) + CVaR_n^p \Big(X_n(\mathcal{I}_n^m, \mathcal{I}_n^p) \mid \mathcal{I}_n^m \Big) \\
&\quad + \max_{\alpha_n} \Big\{ CVaR_n^m \Big((\alpha_n)^{\mathrm{T}} [\beta \mathbf{M}_{n+1}(\mathcal{I}_{n+1}^m) - \mathbf{M}_n(\mathcal{I}_n^m)] \\
&\qquad + \beta \mathcal{C}_{n+1}(\mathcal{I}_{n+1}^m) \Big) \Big\}
\end{aligned}
$$

The maximization over α_n can be written as the following linear program

$$
\begin{aligned}
&\max_{\alpha_n, \upsilon, z(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)} \upsilon - \frac{1}{\eta_n} \sum_{\mathcal{I}_{n+1}^m} p_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m) z(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m) \\
&\text{s.t.} \\
&z(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m) \geq \upsilon - \big((\alpha_n)^{\mathrm{T}} [\beta \mathbf{M}_{n+1}(\mathcal{I}_{n+1}^m) - \mathbf{M}_n(\mathcal{I}_n^m)] + \beta \mathcal{C}_{n+1}(\mathcal{I}_{n+1}^m) \big) \ \forall\, \mathcal{I}_{n+1}^m \\
&z(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m) \geq 0 \ \forall\, \mathcal{I}_{n+1}^m
\end{aligned}
$$

The dual of the above linear program is then

$$
\begin{aligned}
&\min_{\psi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)} \sum_{\mathcal{I}_{n+1}^m} \psi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m) \beta \mathcal{C}(n+1, \mathcal{I}_{n+1}^m) \\
&\text{s.t.} \\
&0 \leq \psi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m) \leq \frac{p_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)}{\eta_n} \ \forall\, \mathcal{I}_{n+1}^m \\
&\sum_{\mathcal{I}_{n+1}^m} \psi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m) = 1 \\
&\sum_{\mathcal{I}_{n+1}^m} \psi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m) \beta M_{n+1}(j, \mathcal{I}_{n+1}^m) = M_n(j, \mathcal{I}_n^m) \ \forall\, j
\end{aligned}
$$

By the partial markets assumption, there is a unique solution, namely the risk-neutral probabilities $\pi_n(\mathcal{I}_{n+1}^m, \mathcal{I}_n^m)$, which satisfy the set of linear equalities in the above minimization. By the conditions of the theorem, the risk-neutral probabilities also satisfy the inequalities in the above problem. Substituting this, we get

$$\begin{aligned} V_n(\mathbf{X}, \alpha_{n-1}, \mathcal{I}_n^m) &= (\alpha_{n-1})^{\mathrm{T}} \mathbf{M}_n(\mathcal{I}_n^m) + CVaR_n^p\Big(X_n(\mathcal{I}_n^m, \mathcal{I}_n^p) \mid \mathcal{I}_n^m\Big) \\ &\qquad + \beta \mathbb{E}_\pi \left[\mathcal{C}_{n+1}(\mathcal{I}_{n+1}^m) | \mathcal{I}_n^m\right] \\ &= (\alpha_{n-1})^{\mathrm{T}} \mathbf{M}_n(\mathcal{I}_n^m) + \mathbb{E}_\pi \left[\sum_{\tau=n}^{N} \beta^{\tau-n} CVaR_\tau^p\Big(X_\tau(\mathcal{I}_\tau^m, \mathcal{I}_\tau^p) \mid \mathcal{I}_\tau^m\Big) \Big| \mathcal{I}_n^m \right] \end{aligned}$$

□

Proof of Theorem 4.2. The inequality (4.23) implies $\beta \mathbb{E}_{\pi_{N-1}} \left[CVaR_N^p(S_N, \mathcal{I}_N^m)\right] < \Delta_{N-1}^m - p$ and hence the theorem is true for $n = N - 1$. Suppose the theorem is true for $n + 1, \ldots, N - 1$.

Now, $\frac{\partial U_{n+1}}{\partial e_{n+1}} \leq \Delta_{n+1}^m - p$ for each $\mathcal{I}_{n+1}$ implies that $\frac{\partial CVaR_{n+1}^p\left(U_{n+1}, \mathcal{I}_{n+1}^m\right)}{\partial e_{n+1}} \leq \Delta_{n+1}^m - p$ and hence $\Upsilon_n(k)$, the slope of $\beta \mathbb{E}_{\pi_n} \left[CVaR_{n+1}^p\Big(U_{n+1}, \mathcal{I}_{n+1}^m\Big)\right]$ is $\leq \beta \mathbb{E}_{\pi_n} \left[\Delta_{n+1}^m - p\right] = \Delta_n^m - p$ for $k = 1, \ldots, \kappa_n + 1$. From equation (4.19), we have $\underline{b_n} = 0$ and the optimal processing quantity is given by $m_n^* = \min\{C, e_n + x_n\}$ for a given procurement quantity x_n.

By the concavity and piecewise linear nature of $\mathcal{H}_n(e_{n+1})$ in e_{n+1}, we can write

$$\begin{aligned} \mathcal{H}_n(e_{n+1}) &= \max_{\delta e_{n+1}^{(k)}} \mathcal{H}_n(0) + \sum_{k=1}^{\kappa_n+1} \Upsilon_n^{(k)} \delta e_{n+1}^{(k)} \\ &\quad \text{s.t.} \\ &\quad 0 \leq \delta e_{n+1}^{(k)} \leq g_n^{(k)} \;\; k = 1, \ldots, \kappa_n \\ &\quad \delta e_{n+1}^{(\kappa_n+1)} \geq 0 \\ &\quad \sum_{k=1}^{\kappa_n+1} \delta e_{n+1}^{(k)} = e_{n+1} \end{aligned}$$

for any $e_{n+1} \geq 0$ and for each $\mathcal{I}_n$

Using the above representation for $\mathcal{H}_n$, we have

$$U_n(e_n, \mathcal{I}_n) = \max_{0 \le x_n \le (C-e_n)} \{(\Delta_n^m - p)(e_n + x_n) - S_n x_n + \mathcal{H}_n(0)\}$$

if $0 \le e_n \le C$ and $S_n > \Upsilon_n^{(1)}$ and

$$\begin{aligned} U_n(e_n, \mathcal{I}_n) \;\; = \;\; & \max_{(C-e_n)^+ \le x_n \le C} \left\{ (\Delta_n^m - p)C - S_n x_n + \mathcal{H}_n(0) + \sum_{k=1}^{\kappa_n+1} \Upsilon_n^{(k)} \delta e_{n+1}^{(k)} \right\} \\ & \text{s.t.} \\ & 0 \le \delta e_{n+1}^{(k)} \le g_n^{(k)} \;\; k = 1, \ldots, \kappa_n \\ & \delta e_{n+1}^{(\kappa_n+1)} \ge 0 \\ & \sum_{k=1}^{\kappa_n+1} \delta e_{n+1}^{(k)} = e_n + x_n - C \end{aligned}$$

if $e_n > C$ or $S_n \le \Upsilon_n^{(1)}$.

In the first scenario, the marginal risk-adjusted value of input inventory, i.e., the slope of U_n with respect to e_n is equal to $\min\{\Delta_n^m - p, S_n\} \le \Delta_n^m - p$. In the second scenario, the slope of U_n with respect to e_n is equal to $\max\{S_n, \Upsilon_n^{(\bar{k})}\}$ where $\bar{k}$ is such that $e_n + C \in [b_n(\bar{k}-1), b_n(\bar{k}))$. By concavity of $\mathcal{H}_n$ and the fact that $S_n \le \Upsilon_n^{(1)}$, we have $\max\{S_n, \Upsilon_n^{(\bar{k})}\} \le \Upsilon_n^{(1)} \le \Delta_n^m - p$ where the second inequality follows from the induction hypothesis. Thus, for all e_n and all $\mathcal{I}_n$, we have $\frac{\partial U_n}{\partial e_n} \le \Delta_n^m - p$. □

Proof of Theorem 4.3. We prove the theorem by induction. Clearly, the theorem is true for period $N-1$ and we can write

$$U_{N-1}(e_{N-1}, \mathcal{I}_{N-1}) = \Theta_{N-1}^{(k)} e_{N-1} + \Lambda_{N-1}^{(k)} \text{ for } e_{N-1} \in [(k-1)D, kD)$$

We have

$$\begin{aligned}
CVaR^p_{N-1} & \Big(U_{N-1}(e_{N-1},\mathcal{I}_{N-1}),\mathcal{I}^m_{N-1}\Big) \\
= \quad & \min_{\psi(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1})} \sum_{\mathcal{I}^p_{N-1}} \psi(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1}) U_{N-1}(e_{N-1},\mathcal{I}_{N-1}) \\
& \text{s.t.} \\
& 0 \le \psi(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1}) \le \frac{q_{N-1}(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1})}{\eta} \ \forall\, \mathcal{I}^p_{N-1} \\
& \sum_{\mathcal{I}^p_{N-1}} \psi(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1}) = 1
\end{aligned}$$

Let $\psi^*(\cdot)$ be the optimal solution to the above problem so that

$$\begin{aligned}
CVaR^p_{N-1} & \Big(U_{N-1}(e_{N-1},\mathcal{I}_{N-1}),\mathcal{I}^m_{N-1}\Big) \\
= \quad & \sum_{\mathcal{I}^p_{N-1}} \psi^*(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1}) U_{N-1}(e_{N-1},\mathcal{I}_{N-1}) \\
= \quad & \sum_{\mathcal{I}^p_{N-1}} \psi^*(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1}) \Theta^{(k)}_{N-1} e_{N-1} + \sum_{\mathcal{I}^p_{N-1}} \psi^*(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1}) \Lambda^{(k)}_{N-1} \\
= \quad & -\sum_{\mathcal{I}^p_{N-1}} \psi^*(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1}) (-\Theta^{(k)}_{N-1}) e_{N-1} \\
& \qquad\qquad + \sum_{\mathcal{I}^p_{N-1}} \psi^*(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1}) \Lambda^{(k)}_{N-1}
\end{aligned}$$

For a given $\mathcal{I}^m_{N-1}$, notice that $\Theta^{(k)}_{N-1} \le \Delta^m_{N-1} - p$. Let $\mathcal{I}^p_{N-1}(i)$ and $\mathcal{I}^p_{N-1}(j)$ be such that $S_{N-1}(\mathcal{I}^m_{N-1},\mathcal{I}^p_{N-1}(i)) > S_{N-1}(\mathcal{I}^m_{N-1},\mathcal{I}^p_{N-1}(j))$ (In the following, we use the shorthand notation $g_n(\cdot,j)$ to denote $g_n(\cdot,\mathcal{I}^m_n,\mathcal{I}^p_n(j))$ in order to simplify the notation). From equations (4.25) and (4.26) we can verify that $U_{N-1}(e_{N-1},i) \le U_{N-1}(e_{N-1},j)$ and $\Theta^{(k)}_{N-1}(i) \ge \Theta^{(k)}_{N-1}(j)$ for all $e_{N-1} \in [(k-1)D, kD)$, for each k. Also, the two inequalities together imply that $\Lambda^{(k)}_{N-1}(i) \le \Lambda^{(k)}_{N-1}(j)$ for each k. As a result, we have $CVaR^p_{N-1}\Big(\Lambda^{(k)}_{N-1},\mathcal{I}^m_{N-1}\Big) = \sum_{\mathcal{I}^p_{N-1}} \psi^*(\mathcal{I}^p_{N-1},\mathcal{I}^m_{N-1})\Lambda^{(k)}_{N-1}$ and $CVaR^p_{N-1}\Big(-$

$\Theta^{(k)}_{N-1}, \mathcal{I}^m_{N-1}\Big) = \sum_{\mathcal{I}^p_{N-1}} \psi^*(\mathcal{I}^p_{N-1}, \mathcal{I}^m_{N-1})(-\Theta^{(k)}_{N-1})$ and thus

$$\begin{aligned} CVaR^p_{N-1}\Big(U_{N-1}(e_{N-1}, \mathcal{I}_{N-1}), \mathcal{I}^m_{N-1}\Big) \;&=\; CVaR^p_{N-1,l}\Big(\Lambda^{(k)}_{N-1}, \mathcal{I}^m_{N-1}\Big) \\ &\quad - CVaR^p_{N-1}\Big(-\Theta^{(k)}_{N-1}, \mathcal{I}^m_{N-1}\Big) \times e_{N-1} \end{aligned}$$

and $\beta \mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_{N-1}\Big(U_{N-1}(e_{N-1}, \mathcal{I}_{N-1}), \mathcal{I}^m_{N-1}\Big)\Big] = \Psi^{(k)}_{N-1} + \Upsilon^{(k)}_{N-1} e_{N-1}$ for $e_{N-1} \in [(k-1)D, kD)$, where $\Psi^{(k)}_{N-1} \triangleq \beta \mathbb{E}_{\pi_{N-1}}\Big[CVaR^p_{N-1}\Big(\Lambda^{(k)}_{N-1}, \mathcal{I}^m_{N-1}\Big)\Big]$.

Suppose the above hold for periods $n+1, \ldots, N-1$. In period $n+1$, we have

$$U_{n+1}(e_{n+1}, \mathcal{I}_{n+1}) = \Theta^{(k)}_{n+1} e_{n+1} + \Lambda^{(k)}_{n+1} \text{ for } e_{n+1} \in [(k-1)D, kD)$$

where $\Delta^m_{n+1} - p \geq \Theta^{(k)}_{n+1} \geq \Theta^{(k+1)}_{n+1}$. Further, for i and j such that $S_{n+1}(i) > S_{n+1}(j)$, we have $U_{n+1}(e_{n+1}, i) \leq U_{n+1}(e_{n+1}, j)$ and $\Theta^{(k)}_{n+1}(i) \geq \Theta^{(k)}_{n+1}(j)$ for all k for all $e_{n+1} \in [(k-1)D, kD)$, for a given $\mathcal{I}^m_{n+1}$.

In period n, we have

$$U_n(e_n, \mathcal{I}_n) = \max_{0 \leq x \leq K} \left\{ \max_{0 \leq m \leq \min\{C, e_n + x\}} \{(\Delta^m_n - p)m + \mathcal{H}_n(e_{n+1}, \mathcal{I}_n)\} - S_n x \right\}$$

where

$$\mathcal{H}_n(e_{n+1}, \mathcal{I}_n) = \beta \mathbb{E}_{\pi_n}\Big[CVaR^p_{n+1}\Big(U_{n+1}(e_{n+1}, \mathcal{I}_{n+1}), \mathcal{I}^m_{n+1}\Big)\Big] = \Psi^{(k)}_n + \Upsilon^{(k)}_n e_{n+1}$$

for $e_{n+1} \in [(k-1)D, kD)$, for each $k \geq 1$.

As $\Delta^m_{n+1} - p \geq \Theta^{(k)}_{n+1}$ for each k, we have

$$\begin{aligned} & CVaR^p_{n+1}\Big(-\Theta^{(k)}_{n+1}, \mathcal{I}^m_{n+1}\Big) \geq -(\Delta^m_{n+1} - p) \\ \Rightarrow \quad & -CVaR^p_{n+1}\Big(-\Theta^{(k)}_{n+1}, \mathcal{I}^m_{n+1}\Big) \leq (\Delta^m_{n+1} - p) \\ \Rightarrow \quad & \beta \mathbb{E}_{\pi_n}\Big[-CVaR^p_{n+1}\Big(-\Theta^{(k)}_{n+1}, \mathcal{I}^m_{n+1}\Big)\Big] = \Upsilon^{(k)}_n \leq \beta \mathbb{E}_{\pi_n}\big[\Delta^m_{n+1} - p\big] = \Delta^m_n - p \end{aligned}$$

and therefore the optimal processing quantity in period n is equal to $\min\{C, e_n + x\}$ for each $\mathcal{I}_n$.

Consider the case when $e_n \in [(k-1)D, kD)$ with $k \geq a+1$; i.e., $e_n \geq C$. In this case, $m_n^* = C$ and we have

$$U_n(e_n, \mathcal{I}_n) = (\Delta_n^m - p)C + \max_{0 \leq x \leq K} \left\{ \Upsilon_n^{(s-a)} \times (e_n + x - C) - S_n x + \Psi_n^{(s-a)} \right\}$$

where s is an integer such that $e_n + x \in [(s-1)D, sD)$. By concavity of $\mathcal{H}_n$, we can determine the optimal procurement quantity, x_n^*, as follows

$$x_n^* = \begin{cases} K & \text{if } \Upsilon_n^{(k+b-a)} \geq S_n \\ \hat{r}_n D + C - e_n & \text{if } \Upsilon_n^{(k-a)} \geq S_n > \Upsilon_n^{(k+b-a)} \text{ and } \hat{r}_n \text{ s.t. } \Upsilon_n^{\hat{r}_n} \geq S_n > \Upsilon_n^{(\hat{r}_n+1} \\ 0 & \text{if } S_n > \Upsilon_n^{(k-a)} \end{cases}$$

Substituting x_n^*, we can therefore write

$$U_n(e_n, \mathcal{I}_n) = \max\left\{ \Upsilon_{n+1}^{(k+b-a)}, \min\left\{ S_n, \Upsilon_{n+1}^{(k-a)} \right\} \right\} e_n + \Lambda_n^{(k)}$$

where $\Lambda_n^{(k)}$ represents constant terms not involving e_n, when $e_n \in [(k-1)D, kD)$ for $k \geq a+1$.

For a given $\mathcal{I}_n^m$, suppose $\mathcal{I}_n^p(i)$, with $i = 1, 2, 3$ are such that $S_n(1) > \Upsilon_n^{(k-a)} \geq$

$S_n(2) > \Upsilon_n^{(k+b-a)} \geq S_n(3)$. Substituting x_n^*, we get

$$
\begin{aligned}
U_n(e_n, 3) &= (\Delta_n^m - p)C + \Upsilon_n^{(k+b-a)} e_n + \Psi_n^{(k+b-a)} - \Upsilon_n^{(k+b-a)} C \\
&\qquad + \left[\Upsilon_n^{(k+b-a)} - S_n(3)\right] K \\
&\geq (\Delta_n^m - p)C + \Upsilon_n^{(k+b-a)} e_n + \Psi_n^{(k+b-a)} - \Upsilon_n^{(k+b-a)} C \\
&\qquad + \left[\Upsilon_n^{(k+b-a)} - S_n(3)\right] (\hat{r}_n D + C - e_n) \\
&\qquad\qquad (\text{since } K > \hat{r}_n D + C - e_n) \\
&= (\Delta_n^m - p)C + \Upsilon_n^{(k+b-a)} \hat{r}_n D + \Psi_n^{(k+b-a)} \\
&\qquad + S_n(3) e_n - S_n(3)\hat{r}_n D - S_n(3) C \\
&\geq (\Delta_n^m - p)C + \Upsilon_n^{(\hat{r}_n)} \hat{r}_n D + \Psi_n^{(\hat{r}_n)} - S_n(3)(\hat{r}_n D + C - e_n) \\
&\qquad\qquad (\text{by concavity of } \mathcal{H}_{n+1}) \\
&\geq (\Delta_n^m - p)C + \Upsilon_n^{(\hat{r}_n)} \hat{r}_n D + \Psi_n^{(\hat{r}_n)} - S_n(2)(\hat{r}_n D + C - e_n) \\
&\qquad\qquad (\text{since } S_n(3) < S_n(2)) \\
&= U_n(e_n, 2) \\
&\geq (\Delta_n^m - p)C + \Upsilon_n^{(\hat{r}_n)} (e_n - C) + \Psi_n^{(\hat{r}_n)} - S_n(2)(e_n - C + C - e_n) \\
&\qquad\qquad (\text{since } \hat{r}_n D \geq e_n - C) \\
&= (\Delta_n^m - p)C + \Upsilon_n^{(\hat{r}_n)} (e_n - C) + \Psi_n^{(\hat{r}_n)} \\
&\geq (\Delta_n^m - p)C + \Upsilon_n^{(k-a)} (e_n - C) + \Psi_n^{(k-a)} \qquad (\text{by concavity of } \mathcal{H}_{n+1}) \\
&= U_n(e_n, 1)
\end{aligned}
$$

Also, $\Theta_n^{(k)}(1) = \Upsilon_n^{(k-a)} \geq \Theta_n^{(k)}(2) = S_n(2) > \Theta_n^{(k)}(3) = \Upsilon_n^{(k+b-a)}$. As a result, we have

$$
CVaR_n^p\Big(U_n(e_n, \mathcal{I}_n), \mathcal{I}_n^m\Big) = CVaR_n^p\Big(\Lambda_n^{(k)}, \mathcal{I}_n^m\Big) - CVaR_n^p\Big(-\Theta_n^{(k)}, \mathcal{I}_n^m\Big) e_n
$$

for $e_n \in [(k-1)D, kD)$ and $k \geq a+1$.

For the case $e_n \in [(k-1)D, kD)$ with $k \leq a$; i.e., $e_n < C$, the optimal procurement quantity is given by

$$x_n^* = \begin{cases} K & \text{if } \Upsilon_n^{(k+b-a)} \geq S_n \\ \\ \hat{r}_n D + C - e_n & \text{if } \Delta_n^m - p \geq S_n > \Upsilon_n^{(k+b-a)} \text{ and } \hat{r}_n \text{ s.t. } \Upsilon_n^{\hat{r}_n} \geq S_n > \Upsilon_n^{(\hat{r}_n+1)} \\ \\ 0 & \text{if } S_n > \Delta_n^m - p \end{cases}$$

Substituting the optimal procurement quantity in the optimization, we get

$$U_n(e_n, \mathcal{I}_n) = \max\left\{\Upsilon_n^{(k+b-a)}, \min\left\{S_n, \Delta_n^m - p\right\}\right\} e_n + \Lambda_n^{(k)}$$

where $\Lambda_n^{(k)}$ represents constant terms not involving e_n, when $e_n \in [(k-1)D, kD)$ for $k \leq a$.

Using arguments similar to the case when $e_n \geq C$, we can show that $U_n(e_n, 1) \leq U_n(e_n, 2) \leq U_n(e_n, 3)$ and $\Theta_n^{(k)}(1) \geq \Theta_n^{(k)}(2) \geq \Theta_n^{(k)}(3)$ where $\mathcal{I}_n^p(i)$ are such that $S_n(1) > S_n(2) > S_n(3)$, for a given $\mathcal{I}_n^m$. Thus, by induction, $U_n(e_n, \mathcal{I}_n)$ is concave, continuous and piecewise linear in e_n with break points at integral multiples of D, and the marginal risk-adjusted value of input inventory is given by equation (4.27). □

Proof of Proposition 4.2. We know that $U_n(e_n, \mathcal{I}_n) = \Theta_n^{(k)} e_n + \Lambda_n^{(k)}$ for $e_n \in [(k-1)D, kD)$ for each $k \geq 1$. Thus, $U_n(0, \mathcal{I}_n) = \Lambda_n^{(1)}$.

We have

$$\begin{aligned} \Lambda_n^{(1)} \triangleq U_n(0, \mathcal{I}_n) &= (\Delta_n^m - p - S_n)^+ C + \Upsilon_n^{(\hat{r}_n)} \times (\hat{r}_n D) + \Psi_n^{(\hat{r}_n)} - S_n \hat{r}_n D \\ &= (\Delta_n^m - p - S_n)^+ C + \Psi_n^{(1)} + \sum_{j=1}^{\hat{r}_n} \Upsilon_n^{(j)} D - S_n \hat{r}_n D \\ &= (\Delta_n^m - p - S_n)^+ C + \Psi_n^{(1)} + \sum_{j=1}^{b-a} \left(\Upsilon_n^{(j)} - S_n\right)^+ D \end{aligned}$$

where $\hat{r}_n$ is such that $\Upsilon_n^{(\hat{r}_n)} \geq S_n > \Upsilon_n^{(\hat{r}_n+1)}$ or $(b-a)$, which ever is smaller. The first equality follows from the continuity of $\mathcal{H}_n$ whereby $\Psi_n^{(k)} = \Psi_n^{(1)} + \sum_{j=1}^{k-1} \left(\Upsilon_n^{(j)} - \Upsilon_n^{(k)}\right) D$ for each k. The second equality follows from the concavity of $\mathcal{H}_n$.

Notice, $\Lambda_n^{(1)}$ and $\hat{U}_n(0, \mathcal{I}_n)$ are both decreasing in S_n for each $n < N$. Further,

$$\begin{aligned}
\nu_n(K, C) &= \Lambda_n^{(1)} - \hat{U}_n(0, \mathcal{I}_n) \\
&= (\Delta_n^m - p - S_n)^+ C + \Psi_n^{(1)} + \sum_{j=1}^{b-a} \left(\Upsilon_n^{(j)} - S_n\right)^+ D \\
&\quad - (\Delta_n^m - p - S_n)^+ C - \beta \mathbb{E}_{\pi_n} \left[CVaR_{n+1}^p \left(\hat{U}_{n+1}(0, \mathcal{I}_{n+1}), \mathcal{I}_{n+1}^m \right) \right] \\
&= \Psi_n^{(1)} + \sum_{j=1}^{b-a} \left(\Upsilon_n^{(j)} - S_n\right)^+ D - \beta \mathbb{E}_{\pi_n} \left[CVaR_{n+1}^p \left(\hat{U}_{n+1}(0, \mathcal{I}_{n+1}), \mathcal{I}_{n+1}^m \right) \right]
\end{aligned}$$

is also decreasing in S_n, for all $n < N$. Thus, we have

$$CVaR_n^p\left(\nu_n(K, C), \mathcal{I}_n^m\right) = CVaR_n^p\left(\Lambda_n^{(1)}, \mathcal{I}_n^m\right) - CVaR_n^p\left(\hat{U}_n(0, \mathcal{I}_n), \mathcal{I}_n^m\right)$$

for all $n < N$. Finally,

$$\begin{aligned}
\nu_n(K, C) &= \Psi_n^{(1)} + \sum_{j=1}^{b-a} \left(\Upsilon_n^{(j)} - S_n\right)^+ D - \beta \mathbb{E}_{\pi_n} \left[CVaR_{n+1}^p \left(\hat{U}_{n+1}(0, \mathcal{I}_{n+1}), \mathcal{I}_{n+1}^m \right) \right] \\
&= \beta \mathbb{E}_{\pi_n} \left[CVaR_{n+1}^p \left(\Lambda_{n+1}^{(1)}, \mathcal{I}_{n+1} \right) \right] + \sum_{j=1}^{b-a} \left(\Upsilon_n^{(j)} - S_n\right)^+ D \\
&\qquad - \beta \mathbb{E}_{\pi_n} \left[CVaR_{n+1}^p \left(\hat{U}_{n+1}(0, \mathcal{I}_{n+1}), \mathcal{I}_{n+1}^m \right) \right] \\
&= \beta \mathbb{E}_{\pi_n} \left[CVaR_{n+1}^p \left(\Lambda_{n+1}^{(1)}, \mathcal{I}_{n+1}^m \right) - CVaR_{n+1}^p \left(\hat{U}_{n+1}(0, \mathcal{I}_{n+1}), \mathcal{I}_{n+1}^m \right) \right] \\
&\qquad + \sum_{j=1}^{b-a} \left(\Upsilon_n^{(j)} - S_n\right)^+ D \\
&= \beta \mathbb{E}_{\pi_n} \left[CVaR_{n+1}^p \left(\nu_{n+1}(K, C), \mathcal{I}_{n+1}^m \right) \right] + \sum_{j=1}^{b-a} \left(\Upsilon_n^{(j)} - S_n\right)^+ D
\end{aligned}$$

□

Chapter 5

Commodity Operations in a Network Environment: Model, Analysis and Heuristics

5.1. Introduction

Many agricultural commodities, e.g., wheat, corn, soybean are produced in different, geographically spread locations. Energy commodities such as natural gas and crude oil are also procured and transported across multiple locations. Firms which use these commodities as inputs to their production process generally procure them from multiple locations for a variety of reasons, including price differentials across locations and capacity constraints. Similarly, firms may have processing capacities at fixed locations, with the output commodities requiring delivery to various locations. Profits for such firms are affected not only by the stochastic prices of the commodities at different locations, but also by other network characteristics such as transshipment costs, capacity constraints, transportation lead times, etc. In this essay, we study the impact of network characteristics on the integrated procurement, processing and trade decisions for a commodity processing firm operating a multi-location network.

While commodity production and distribution networks have been studied earlier (cf., Markland (1975), Markland and Newett (1976)), these papers assume deterministic commodity prices and no operational constraints. A large stream of literature in the operations management area also looks at optimal inventory and/or transshipment decisions for any given network; see for example, Karmarkar (1981), Karmarkar (1987), Federgruen and Zipkin (1984), Robinson (1990) Hu et al. (2004) etc. However, all these papers deal with known, proportional costs of procurement and do not consider multiple options for earning revenues. In contrast, we explicitly incorporate stochastic commodity prices and capacity constraints, which makes the problem non-trivial.

Papers that do incorporate stochastic commodity prices across multiple locations are usually restricted to single period models. Secomandi (2010a) considers the valuation of pipeline capacity to transport natural gas between two locations. Martınez-de Albéniz and Simón (2010) consider a capacitated commodity trading model, where a trader takes advantage of geographical spread in commodity prices by transshipping the commodity from the location with lower price and selling in the location with a higher price. They model a trader who has market power and the influence of the trader's actions on future prices of the commodities. In contrast, the firm in our model does not influence the market prices of the commodities through its actions. We also consider inventory carried across periods in our model, in contrast to Martınez-de Albéniz and Simón (2010). Finally, our model considers multiple commodities, with the ability to irreversible transform the input to an output commodity.

Somewhat related to this research is the work by Goel and Gutierrez (2008), who consider commodity procurement and distribution decisions in a supply chain. They model a two-echelon supply chain with a central warehouse supplying multiple retailers, each of who faces a stochastic demand. The central warehouse can procure the commodity from the spot and forward markets, with the two sources having different lead times. They derive optimal replenishment and distribution policies for the supply chain. In our network model, we allow transshipment between procurement and processing locations. Further, we consider processing decisions and capacity constraints, which are absent in their model.

The rest of this chapter is organized as follows. We formulate the problem for a firm operating a star network in Section 5.2. We propose various heuristic to solve the star network problem in Section 5.3 and quantify the performance of the various heuristics in Section 5.4. Section 5.5 concludes with directions for future research.

5.2. Model Formulation and Analysis

We consider the integrated problem of procurement, processing and trade over a multi-node network of M procurement nodes each with procurement capacity of K^i units per period at location $i \in \{1, 2, \ldots, M\}$. 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. 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.

The transshipment cost is $t^{(ij)}$ per unit between locations i and j, with $i \neq j$. 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[1]; i.e., $t^{(ij)} = \infty$ for $(i, j) \in \{2, 3, \ldots, M\} \times \{2, 3, \ldots, M\}$.

[1]This restriction on possible transshipment is also consistent with the actual features of the ITC

Similar to the single node problems considered in Chapter 2, the firm sells all the output using forward contracts and the procurement season for the input commodity may span multiple output forward maturities. 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 $n < N_\ell \leq N$. In addition to selling the output commodity, the firm can also trade the input itself with other processors over the horizon. To keep the exposition simple, 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^i at location i.

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 against contract ℓ for all ℓ such that $N_\ell > n$: q_n^ℓ, 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, Lemmas 2.1 and 2.2 hold for the network case also and the marginal value of output is given by equation (2.4). Further, the value function $V_n(\mathbf{e_n}, Q_n, \mathcal{I}_n)$ is separable in

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.

$\mathbf{e_n}$ and Q_n as given by equation (2.6) and we have

$$\begin{aligned}
U_n(\mathbf{e_n}, \mathcal{I}_n) &= \max_{(\mathbf{x_n}, \mathbf{y_n}, m_n) \in \mathcal{B}_n} \Bigg\{ [\Delta_n - p] m_n - \sum_{i=1}^{M} S_n^i x_n^i \\
&\qquad - \sum_{i=2}^{M} t^{(i)} [y_n^{(1i)} + y_n^{(i1)}] \\
&\qquad - h_I \left[\sum_{i=1}^{M} (e_n^i + x_n^i) - m_n \right] \\
&\qquad + \beta \mathbb{E}_{\mathcal{I}_n} [U_{n+1}(\mathbf{e_{n+1}}, \mathcal{I}_{n+1})] \Bigg\} \quad \text{for } n < N \quad (5.1) \\
U_N(\mathbf{e_N}, \mathcal{I}_N) &= \sum_{i=1}^{M} S_N^i e_N^i \qquad (5.2)
\end{aligned}$$

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

$$\mathcal{B}_n = \left\{ \begin{array}{ll} 0 \le x_n^i \le K^i & \text{for } i = 1, 2, \ldots, M \\ 0 \le m_n \le C & \\ m_n + \sum_{i=2}^{M} y^{(1i)} \le e_n^1 + x_n^1 + \sum_{j=2}^{M} y_n^{(j1)} & \\ y^{(i1)} \le e_n^i + x_n^i & \text{for } i = 2, 3, \ldots, M \\ \mathbf{x_n} \ge 0, \mathbf{y_n} \ge 0, m_n \ge 0 & \end{array} \right\} \qquad (5.3)$$

and the state transition equations are given by

$$e_{n+1}^i = \left\{ \begin{array}{ll} e_n^i + x_n^i + \sum_{j=2}^{M} y_n^{(j1)} - \sum_{j=2}^{M} y_n^{(1j)} - m_n & \text{for } i = 1 \\ e_n^i + x_n^i + y_n^{(1i)} - y_n^{(i1)} & \text{for } i = 2, \ldots, M \end{array} \right. \qquad (5.4)$$

Notice that (5.2) 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 a result, the optimal procurement and processing policy also depend on the entire inventory vector $\mathbf{e_n}$ and it is hard to solve the network problem without additional simplifications. In the next section, we consider some simplifications and use these simplifications to develop tractable heuristics for the network problem.

5.3. Heuristics and Upper Bound for the Network Problem

The complexity in solving the DP given by (5.1)–(5.2) arises primarily from the complexity in computing the value-to-go function, $U_{n+1}(\mathbf{e_{n+1}}, \mathcal{I}_{n+1})$, for a multi-dimensional state space. We develop heuristics by considering approximations to $U_{n+1}(\mathbf{e_{n+1}}, \mathcal{I}_{n+1})$ that are easy to compute. Approximations to the value-to-go function can be achieved by reducing the number of periods considered in the remaining planning horizon or by reducing the dimensions of the state variable. We present heuristic policies based on both these approaches in this section. The Equivalent Single Node (ESN) heuristic uses the similarities between the network problem and single node problem with convex cost of procurement to reduce the dimensionality of the state space. On the other hand, the Network Full Commitment (NFC) heuristic is a myopic heuristic which approximates the value-to-go function by reducing the number of periods considered in the remaining planning horizon.

5.3.1 Equivalent Single Node (ESN) Heuristic

Solving the network problem optimally is 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. However, the network problem is tractable and is equivalent to the single node problem with piecewise linear, convex cost of procurement under some simplifying assumptions. 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. 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 (5.1)–(5.2) as

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

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

$$
\begin{aligned}
U_n(\hat{e}_n, \mathcal{I}_n) &= \max_{(\mathbf{x_n}, m_n) \in \hat{\mathcal{B}}_n} \Bigg\{ [\Delta_n - p]m_n - \sum_{i=1}^{M} S_n^i x_n^i \\
&\qquad - h_I \left[(\hat{e}_n + \sum_{i=1}^{M} x_n^i) - m_n \right] \\
&\qquad + \beta \mathbb{E}_{\mathcal{I}_n} [U_{n+1}(\hat{e}_{n+1}, \mathcal{I}_{n+1})] \Bigg\} \text{ for } n < N \qquad (5.5) \\
U_N(\hat{e}_N, \mathcal{I}_N) &= S_N \hat{e}_N \qquad (5.6)
\end{aligned}
$$

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

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

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 problem analyzed in Chapter 2, albeit with a convex cost of procurement in each period. We now re-visit the single node problem analyzed in Chapter 2, with the change that the procurement cost is a convex function of the total quantity procured, and derive expressions for the marginal value of input inventory. We then propose a tractable heuristic, the equivalent single node (ESN) heuristic, to solve the network problem.

5.3.1.1 Single Node Problem with Convex Cost of Procurement

The analysis in Chapter 2 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 will see in Section 5.3.1.2, the results obtained here will be useful in developing a heuristic for the network problem. With this motivation, we consider the situation when the

[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.

firm has a convex cost of procurement.

We assume all aspects of the operations remain the same as in Section 2.3, 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

$$\mathscr{C}(S_n, x_n) = \begin{cases} \gamma^1 S_n \times x_n & \text{if } 0 \leq x_n \leq K^1 \\ \\ \gamma^j S_n \times [x_n - K^{j-1}] & \\ \qquad + \displaystyle\sum_{i=1}^{j-1} \gamma^i S_n \times [K^i - K^{i-1}] & \text{if } K^{j-1} < x_n \leq K^j \\ & \quad \text{for } j = 2, \ldots, J \end{cases} \tag{5.7}$$

where $\gamma^j > \gamma^{j-1}$ and $K^j > K^{j-1}$ for all $j = 1, 2, \ldots, J$, with $\gamma^0 = 0$ and $K^0 = 0$.

Notice that the linear cost of procurement is a special case of this function with $J = 1$, and for which $\gamma^1 = 1$ and $K^1 = K$. One can think of $K^j - K^{j-1}$ as the procurement capacity of the j^{th} lowest cost source, from the J available sources. Further, a general convex cost of procurement can be approximated by a piecewise linear function such as the one given by equation (5.7) to any required degree of accuracy 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 2.1 holds for this case and the marginal value of output is given by equation (2.4).

We now focus on computing the marginal value of input inventory when the procurement cost is given by equation (5.7). 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 2.1, we can prove the next result.

Theorem 5.1. *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 $\mathcal{I}_n$ when the procurement cost is given by $\mathscr{C}(S_n, x_n)$, as defined in equation (5.7). 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

$$\Theta_n^{(k,j)} = \begin{cases} \Omega_n^{(k)} & j = 0 \\ \max\left\{\Omega_n^{(k+b^j)}, \min\left\{\gamma^j S_n, \Theta_n^{(k,j-1)}\right\}\right\} & \text{for } j = 1, 2, \ldots, J \end{cases} \tag{5.8}$$

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

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

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. Specifically,

Proposition 5.1. *For all $n < N$, let $\Omega_n^{(k)}$ be as defined in equation (2.11). 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 5.1 have been derived assuming the γ^j are stationary. However, equation (5.8) 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 (5.8). 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 (5.8) 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. These generalizations are useful in developing heuristics for the star network problem.

5.3.1.2 The ESN Heuristic

The SDP equations (5.5)–(5.6) 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 5.3.1.1 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 we call this the 'Equivalent Single Node' (ESN) heuristic. We develop the ESN heuristic by first replacing $\mathbf{e}_{n+1}$ with

$\hat{e}_{n+1} = \sum_i e_n^i$. We then replace the stochastic procurement cost over the network in any period by a piecewise linear, convex cost function as follows.

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

$$\gamma_n^j = \mathbb{E}_{\mathcal{I}_1}\left[\frac{S_n^{(j)}}{S_n^1}\middle| S_1^1\right] \tag{5.9}$$

$$\bar{K}_n^j = \mathbb{E}_{\mathcal{I}_1}\left[\sum_{k=1}^{j} K^{i_k}\middle| S_1^1\right] \tag{5.10}$$

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 (5.7), 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 (5.8).

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

$$\hat{U}_n(\mathbf{e_n}, \hat{\mathcal{I}}_n) = \hat{\Theta}_n^k \sum_i e_n^i + \hat{\lambda}_n^k \text{ if } (k-1)D \leq \hat{e}_n < kD \tag{5.11}$$

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 in any period $n < N$ is then given by the solution to the following optimization problem

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

5.3.2 Network Full Commitment (NFC) Heuristic

Myopic policies are examples of heuristics that approximate the value-to-go function by reducing the number of periods considered in the planning horizon. Myopic policies, as approximations to optimal policies, are well studied in the context of multi-period stochastic inventory problems; see for example, Lovejoy (1990), Lovejoy (1992), Morton and Pentico (1995), Anupindi et al. (1996), and Iida (2001). The Network Full Commitment (NFC) heuristic is a myopic heuristic and based on the full commitment policy used in practice (see Section 2.4.2 for a description of the full commitment policy in a single node context).[3]

Under the NFC heuristic, the firm only considers the value from processing and committing to sell the output immediately in the same period. Setting $t^{(11)} = 0$, let m_j be the location corresponding to the j^{th} order statistic of $(S_n^1 + t^{(11)}, S_n^2 + t^{(21)}, \ldots, S_n^j + t^{(j1)}, \ldots, S_n^M + t^{(M1)})$. We determine the procurement and processing quantities for period n in the following manner.

1. $\delta C = C$.

2. $x_n^j = 0$ for $j = 1, 2, \ldots, M$.

3. For $j = 1$ to M

   if $\max_{\ell}\{F_n^{\ell}\} - p \geq S_n^{m_j} + t^{(m_j 1)}$ where ℓ s.t. $N_\ell > n$

   - $x_n^{m_j} = \min\{K^{m_j}, \delta C\}$;
   - $y_n^{(m_j 1)} = x_n^{m_j}$;
   - $\delta C = \delta C - y_n^{(m_j 1)}$;

[3]The NFC heuristic is a modification of the full commitment policy used in practice. In practice, the value from processing and committing to sell the output immediately is compared against a weighted average cost of procurement across all locations where the procurement capacity at each location is used as the weight. Our numerical studies indicate that this heuristic performs very badly, and the NFC heuristic is a significant improvement on the full commitment policy used in practice.

4. $m_n = \sum_{j=1}^{M} y_n^{(j1)}$.

Similar to the full commitment policy in the single node case, the NFC heuristic ignores the 'option' value from postponing commitment of the output, as well as the value from holding and trading the input inventory at the end of the horizon at each location.

We can use dual penalties based on the ESN heuristic and compute an upper bound on the optimal expected profits for the network case by appropriately modifying the information relaxation procedure described in Section 2.5.2 to account for the network characteristics. This upper bound can then be used to evaluate the performance of the two heuristics for the network case.

5.4. Numerical Study

In this section, we quantify the performance of the ESN and NFC heuristics using numerical studies. We consider the soybean procurement and processing decisions as the context and use commodity market data for the soy complex for our numerical studies.

As in Section 2.4, we model the various commodity prices as single-factor mean-reverting processes. As seen in Section 2.4.2, the composite output approximation was close to optimal and for the purposes of this numerical study we model a single composite output. We investigate the performance of the heuristic for a two-node and a five-node network respectively.

5.4.1 Implementation

Price process parameters. We model the parameters of the single factor, mean-reverting price process parameters for the input and a hypothetical, composite output whose price in any period is equal to the total value of soybean meal and soybean oil produced upon processing one bushel of soybeans, where the value is calculated

Table 5.1: Price Process Parameters for Input and Output Commodities

	Input (Soybean)	Output (Composite)
Mean-Reversion Coeff κ_i	0.229	0.5348
Longrun Log level ξ_i	6.738	6.8327
Volatility σ_i	0.244	0.4360
Seasonality Factor $e^{\mu(t)}$		
Jan	0.992	0.988
Feb	0.992	0.988
Mar	0.998	0.993
Apr	0.998	0.993
May	1.000	0.995
Jun	1.000	0.995
Jul	1.017	1.037
Aug	1.010	1.013
Sep	0.991	1.000
Oct	0.991	0.987
Nov	0.989	0.987
Dec	0.989	0.984

based on the current prices of the two products. As only futures instruments are publicly traded for the different output commodities, we consider futures instruments for the composite output as well, where the futures price for a particular maturity is a combination of the futures prices of the individual output commodities , to estimate the price process parameters for the composite output. The futures price information on futures contracts traded on the Chicago Board of Trade (CBOT) for different maturities on each trading day of the month of June 2010 was used to calibrate the parameters for soybean and composite output spot price processes using the same procedure described in Section 2.4.1. The average of the estimated parameters obtained over each trading day are given in Table 5.1 and used to model the price processes.

The input spot price across various locations in the network are not expected to diverge greatly, even though the realizations are not necessarily equal across locations. As a result, we set the parameters of the input spot price process at each location equal to the values given in Table 5.1.[4] However, we model the Brownian motion increments in the input spot prices across different locations to be imperfectly correlated and set the correlation coefficient, ρ_{ij}, equal to 0.9.[5] The correlation between the Brownian motion increments for the input and output commodities, ρ_{io}, was estimated as 0.883.

Evaluation of the heuristics. For computing the ESN heuristic policy, we use the re-combining binomial tree procedure described in Peterson and Stapleton (2002), which can handle mean reversion in prices, to discretize the dynamics of the price processes and approximate the joint evolution of the spot price of the input and output commodities. Each period in the discrete binomial tree corresponds to a week and we discretize the price process with δ steps between each period. In our computational studies, we set $\delta = 20$.

We generate sample paths of input spot prices across different locations and estimate γ^j and $\bar{K}^j$ using equations (5.9) and (5.10) respectively over these sample paths. Using equations (2.4) and (5.8), we can compute Δ_n and Θ_n^k for $k = 1, \ldots, (N-n)a+1$, and thereby the ESN procurement and processing policy at each node in the tree.

We evaluate the performance of the heuristics using Monte Carlo simulation. We generate sample paths of prices for each period $n = 1, 2, \ldots, N$ by sampling from the continuous time price processes. We round the realized input spot price at location 1 and output spot prices to the closest node in the binomial tree and obtain the procurement, processing and commitment quantity suggested by the ESN heuristic for the corresponding to the node and inventory level. For each sample path, the procurement and processing quantities for the NFC heuristic are determined according to the algorithm given in Section 5.3.2. Expected profits from both heuristics are computed as the average profit over 10,000 sample paths. We also compute an upper bound value for each sample path by solving the upper bound problem and average

[4] We also lacked data on spot prices at different locations, prompting us to set the parameters to the values estimated from data available for a single location.

[5] This ensures that while the individual realizations of the spot prices across locations are not identical, they are not greatly divergent either.

across sample paths to obtain the upper bound on expected profits.

Other operational parameters. For all the numerical studies, we set the variable cost of processing p to equal 72 cents / bushel, same as in the single node case of Chapter 2. The procurement capacities in the two-node network are set to 3 and 2 units respectively at the two locations, while in the five-node network the procurement capacity at each location is set to 2 units. These capacities can be considered to be in multiples of bushels, e.g., million bushels. For the base case, we set processing capacity to 60% of total procurement capacity, which is roughly the percentage of soybeans produced in the United States that were estimated to have been crushed 2008 and 2009 (Ash, 2011). We leave the exact units for the capacities unspecified as only the relative values of the procurement and processing capacities matter for computing the policies and multiplying both the capacities by a common factor will scale the expected profits also by the same factor. We set the transportation cost per unit between the central and any of the other procurement locations as 20 cents / bushel, which is roughly 20% of the expected input spot price. We assume the physical holding costs for the various commodities are negligible and normalize them to zero.

5.4.2 Numerical Results

We conduct numerical studies to compute the expected profits for the firm from its procurement and processing operations over the procurement season ranging from August to December. We initialize the prices for all the commodities to their long run average values at the beginning of the planning horizon and evaluate the performance of the heuristics for different horizon lengths. Table 5.2 gives the expected profits and upper bound for different horizon lengths when the firm uses all forward contracts available over the horizon for the output commodity.[6]

The results in Table 5.2 suggest that approximating the network effects (as is done

[6] Unless indicated, the gaps shown in all tables in this section are significant with $p < 0.05$.

Table 5.2: Performance of Heuristics for Different Horizon Lengths

(a) Two-Node Network

Horizon Length (# of Fwds. and Maturities)	Expected Profits NFC	ESN	Upper Bound	Gap (ESN-NFC) (% of ESN)	(UB-ESN) (% of UB)
5 (1, $\{5\}$)	1489.48	1485.32	1606.45	-0.28%†	7.54%
10 (2, $\{5,9\}$)	3164.31	3138.56	3434.34	-0.82%†	8.64%
20 (3, $\{5,9,18\}$)	7049.61	6999.77	8304.22	-0.71%†	15.71%

$^\dagger p - value > 0.1$

(b) Five-Node network

Horizon Length (# of Fwds. and Maturities)	Expected Profits NFC	ESN	Upper Bound	Gap (ESN-NFC) (% of ESN)	(UB-ESN) (% of UB)
5 (1, $\{5\}$)	2736.80	2621.50	3025.69	-4.40%	13.36%
10 (2, $\{5,9\}$)	5936.40	5631.90	6830.30	-5.41%	17.55%
20 (3, $\{5,9,18\}$)	13470.72	12575.80	16899.20	-7.12%	25.58%

by the ESN heuristic) is worse than approximating the effect of future period decisions (as is done by the NFC heuristic) for large networks, as is evident from the large negative gaps between the ESN and NFC heuristics for the five-node network case. The results in the table above are for a base set of parameters and the same behavior is not necessarily true under all circumstances. Specifically, under the current parameter values, a positive processing margin exists at the beginning of the planning horizon and the expected processing margin throughout the horizon is positive. Also, the expected salvage value at the end of the horizon is lower than the expected spot price throughout the horizon. With a processing capacity that is 60% of the total procurement capacity, profits from processing and selling the output are a major portion of the total profits. Taking advantage of the lowest cost of procurement, as is done by the NFC heuristic explicitly, leads to better expected profits.

To investigate how the heuristics perform under different scenarios, we consider sensitivity of the different policies to some key parameters; processing capacity, price volatilities and initial processing margin.

Impact of processing capacity. When processing capacity is limited compared to the procurement capacity, we expected the value of a forward looking heuristic such as ESN to be higher. This is because when the input spot prices are low, the ESN heuristic is likely to procure input for current period processing as well as for the future. Further, including the option to trade input inventory at the end of the horizon is more valuable when processing capacity is limited. However, the ESN heuristic takes into account the network characteristics only in an approximate manner. A myopic heuristic such as NFC does not account for future processing needs or salvage at the end of the horizon, but takes the effect of the network characteristics for the current period decision. Table 5.3 shows the expected profits when using the two heuristics as the processing capacity is varied from 20% to 100% of the procurement capacity.

Table 5.3: Impact of Processing Capacity on ESN and NFC Heuristics
($N = 10$, $L = 2$, $N_\ell = \{5, 9\}$)

(a) Two-Node Network

Processing	Expected Profits		Upper	Gap	
Capacity (% of Total K)	NFC	ESN	Bound	(ESN-NFC) (% of ESN)	(UB-ESN) (% of UB)
20%	1061.41	1086.95	1474.24	2.35%	26.27%
40%	2122.83	2118.38	2457.90	-0.21%†	13.81%
60%	3164.31	3149.33	3435.23	-0.48%†	8.32%
80%	4029.82	4020.91	4248.37	-0.22%†	5.35%
100%	4895.32	4888.84	5032.34	-0.13%†	2.85%

$^\dagger p - value > 0.1$

(b) Five-Node network

Processing	Expected Profits		Upper	Gap	
Capacity (% of Total K)	NFC	ESN	Bound	(ESN-NFC) (% of ESN)	(UB-ESN) (% of UB)
20%	2180.91	2169.84	3306.81	-0.51%†	34.38%
40%	4124.09	3918.57	5041.06	-5.25%	22.27%
60%	5936.40	5660.03	6688.09	-4.88%	15.37%
80%	7624.75	7405.84	8205.72	-2.96%	9.75%
100%	9153.31	9148.37	9516.90	-0.15%†	3.87%

$^\dagger p - value > 0.1$

The benefits of using a forward looking policy such as ESN, albeit by approximating the network characteristics, are significant compared to a myopic policy such as NFC for smaller networks and highly constrained processing firms. This advantage disappears as the processing capacity increases and the gap is negative for larger networks. We see that approximating the network characteristic becomes important for large networks with moderate processing capacities, as seen from Table 5.3 for the five-node network. This is because profits from processing and selling the output are a major portion of the total profits and taking advantage of the lowest cost of procurement, as is done by the NFC heuristic explicitly, leads to better expected profits. As the expected processing margin is positive throughout the horizon, both the ESN and NFC heuristic lead to almost identical policies as the processing capacity increases and thus the gap between the two decreases. Further, because the expected salvage value at the end of the horizon is less than the expected spot price through the horizon, the gap between the heuristics and upper bound also decreases as the processing capacity increases.

Impact of price volatilities. Both heuristics, NFC and ESN, have option like features. The NFC heuristic is equivalent to exercising an European spread option between the output and input spot price for each location, where the number of options exercised is smaller of the remaining processing capacity and procurement capacity at the location. Further, the NFC heuristic incorporates the transportation cost to bring the input to the processing location while making the exercise decision. The ESN heuristic models the output commitment decision as a compound exchange option, while modeling the the procurement and processing decisions based on spread options. However, the spread option exercise decisions at each location in the network are made based on an approximation of the network characteristics and do not capture the full value.

We expect the option value, and hence expected profits under each heuristic, to

Table 5.4: Impact of Price Volatility on ESN and NFC Heuristics
($N = 10,\ L = 2,\ N_\ell = \{5, 9\}$)

(a) Two-Node Network

Price	Expected Profits		Upper	Gap	
Volatility (σ_i, σ_o)	NFC	ESN	Bound	(ESN-NFC) (% of ESN)	(UB-ESN) (% of UB)
0.25	2927.98	2897.91	3151.79	-1.04%	8.06%
0.35	3024.55	2987.42	3405.79	-1.24%	12.28%
0.45	3162.87	3114.77	3745.61	-1.54%	16.84%
0.55	3339.12	3302.25	4178.24	-1.12%†	20.97%
0.65	3549.26	3547.43	4703.01	-0.05%†	24.57%

$^\dagger p - value > 0.1$

(b) Five-Node network

Price	Expected Profits		Upper	Gap	
Volatility (σ_i, σ_o)	NFC	ESN	Bound	(ESN-NFC) (% of ESN)	(UB-ESN) (% of UB)
0.25	5467.65	5168.62	6114.04	-5.79%	15.46%
0.35	5725.01	5321.56	6888.96	-7.58%	22.75%
0.45	6058.97	5549.81	7929.61	-9.17%	30.01%
0.55	6464.19	5889.12	9252.71	-9.76%	36.35%
0.65	6933.43	6375.69	10836.98	-8.75%	41.17%

increase as the volatility of the commodity price processes increases. This can be seen in Table 5.4 which quantifies the performance of the heuristics as the price volatilities are varied. We notice that the gap between the heuristics and the upper bound increases as the volatility increases. This is because the upper bound calculation incorporates the option value in the various decisions fully, while the heuristics only capture the option value partially through approximations. Thus, the gap increases. As seen from the results for the five-node network, the exchange option value inherent in the output commitment decision as modeled by the ESN heuristic is not enough to overcome the reduction in profits because of approximating the network characteristics for large networks.

Impact of initial processing margin. As discussed earlier, the NFC heuristic determines the procurement and processing decisions in a given period based only on the margin from processing and committing to sell the output immediately. As such, we would expect the NFC heuristic to perform well when the expected processing margin is positive throughout the horizon. On the other hand, the ESN heuristic also considers future procurement and processing decisions, albeit by approximating the network as a single node, while determining the current period decisions. Thus, we expect the ESN heuristic to perform better than the NFC heuristic when the processing margins are tight and for smaller networks. Further, as the processing margin increases, we expect both the heuristics to yield similar expected profits as the share of profits obtained from the sale of output in total profits increases. Further, we also expect the gap between the heuristics and the upper bound to decrease as the processing margin increases. The results in Table 5.5 reflect this intuition (The initial margin per unit is calculated as *output spot price in period 1−processing cost−input spot price in period 1 at location 1*).

Overall, the numerical studies indicate sizeable gaps between the heuristics and the upper bound for a variety of conditions. The forward looking ESN heuristic performs well for firms operating smaller networks and facing processing capacity and/or processing margin constraints. On the other hand, the myopic NFC heuristic performs better than the ESN heuristic for large networks. While part of the gap between the upper bound and heuristics is the gap between the upper bound and optimal expected profits, the numerical study provides sufficient evidence that a heuristic policy that is forward looking and also combines the network characteristics is required to capture more of the value from operating a commodity procurement and processing network.

5.5. Conclusions

Table 5.5: Impact of Initial Processing Margin on ESN and NFC Heuristics ($N = 10,\ L = 2,\ N_\ell = \{5, 9\}$)

(a) Two-Node Network

Initial Processing Mgn. (per unit)	Expected Profits NFC	Expected Profits ESN	Upper Bound	Gap (ESN-NFC) (% of ESN)	Gap (UB-ESN) (% of UB)
-15	467.31	479.11	851.50	2.46%†	43.73%
10	756.01	753.65	1126.56	-0.31%†	33.10%
35	1173.14	1153.96	1505.22	-1.66%†	23.34%
60	1646.31	1625.77	1949.47	-1.26%†	16.60%
85	2153.27	2130.99	2437.94	-1.05%†	12.59%

$^\dagger p - value > 0.1$

(b) Five-Node network

Initial Processing Mgn. (per unit)	Expected Profits NFC	Expected Profits ESN	Upper Bound	Gap (ESN-NFC) (% of ESN)	Gap (UB-ESN) (% of UB)
-15	801.72	740.19	2036.07	-8.31%	63.65%
10	1303.44	1167.96	2430.10	-11.60%	51.94%
35	2030.50	1826.84	3048.07	-11.15%	40.07%
60	2941.21	2690.09	3844.26	-9.33%	30.02%
85	3933.50	3666.94	4757.96	-7.27%	22.93%

In this chapter, we have considered the integrated optimization of commodity procurement and processing operations over a network with multiple procurement nodes and a central processing node. Our analysis shows that solving the network problem optimally is considerably more complex and computationally hard, unlike the single node problem considered in Chapter 2. We proposed two computationally efficient heuristics to solve the network problem: a) the Equivalent Single Node (ESN) heuristic approximates the network as a single node with piecewise linear cost of procurement, while b) the Network Full Commitment (NFC) heuristic is a myopic heuristic which only considers the margin from processing and committing to sell the output immediately.

We conducted extensive numerical studies to evaluate the performance of the heuristics by comparing the expected profits against an upper bound on the optimal expected profits. We find that the ESN heuristic performs better than the NFC heuristic for firms operating small networks, with tight processing capacity constraints, and when the processing margins are tight. On the other hand, using a myopic policy such as the NFC is better than approximating the network as a single node, for larger networks. We also find that the gap between the upper bound and both the heuristics is fairly high when processing capacity is tight, initial processing margins are low and commodity prices have high volatility.

Our work lays the foundation for further research in commodity processing and trading networks. Our numerical studies indicate that heuristic policies that are dynamic and also incorporate the network characteristics more explicitly are necessary to capture more of the value from operating commodity procurement and processing networks. An important extension to the present work would be to formulate improved heuristics which combine the ESN and NFC heuristics in an efficient manner to improve performance. The heuristics proposed here model the various commodity price processes as single factor, mean-reverting processes. An important area for future research would be to develop heuristics for managing commodity networks when commodity prices follow multi-factor models.

Another direction for future research is to extend the risk-averse formulation in Chapter 3 to a network setting. An especially interesting problem in this context would be the situation where decisions in the network are made in a de-centralized manner, with each procurement location making procurement decisions independently. In such a situation, the allocation of risk over the network and the impact of decentralized decision making on overall risk are important questions for the firm.

Chapter 6

Conclusion

This dissertation considers various aspects of managing price uncertainty for a commodity processing firm in the presence of operational constraints. Chapters 2, 3 and 4 explored the interdependence between procurement, processing and trade decisions for a firm managing a single location under risk-neutral and risk-averse objective functions, while Chapter 5 looked at the impact of network characteristics on the integrated decision making.

The main insight from the analysis in Chapter 2 is that operational capacity constraints affect how firms interpret price information from commodity markets. We see that a 'low' price, below which it is optimal to buy up to capacity, is dependent on the current inventory level of the input commodity. Similarly, a 'high' price, above which it is optimal to process up to capacity is also dependent on the current inventory levels. We derive the optimal procurement, processing and trade policies for a firm in the presence of operational constraints. While the optimal policies can be computed efficiently when the number of output commodities produced upon processing is small and the commodity prices follow single factor processes, we find that heuristics are needed for computing the policies in more general cases.

The second essay deals with the impact of risk aversion in managing commodity price uncertainty over a multi-period horizon. We elaborate on the notion of time-consistency in risk-averse decision making in Chapter 3. Broadly, time-consistency ensures that optimal decisions for the current period, contingent upon the state in the current period, are also optimal when evaluated in earlier periods. Surprisingly, using risk measures such as conditional value at risk (CVaR) on the total payoffs at the end of the horizon, do not necessarily lead to time-consistent decision making. We propose a dynamic risk measure, DCVaR, to model the firm's risk aversion and ensure time-consistency in decision making. Our results show that the optimal procurement and processing policy under this risk measure are characterized by 'procure up to' and 'process down to' thresholds for the input inventory. Our numerical studies indicate that using a time-consistent risk measure provides a better mean-risk tradeoff in total payoffs over the horizon, as well as better risk control by minimizing the probability of extreme losses over the entire horizon.

Chapter 4 extends the risk-averse analysis of Chapter 3 using a specific framework for modeling the commodity price uncertainty. The partially complete markets framework used in this essay distinguishes between financial market and firm specific or private factors that drive commodity price uncertainty. Extending the time-consistent risk measure introduced in Chapter 3, we characterize the optimal financial trading portfolio as a portfolio which replicates the CVaR over private states of uncertainty of the cashflows generated from operational decisions. Interestingly, we find that the optimal policy to trade the output commodity is not affected by the firm's risk aversion under this framework. Similar to the results in Chapter 3, the optimal operational policy is characterized by 'procure up to' and 'process down to' thresholds for the input inventory. Our results also show that excess processing capacity (relative to procurement capacity) does not provide any additional benefit, while excess procurement capacity provides an operational hedge to manage part of the commodity price uncertainty driven by firm specific factors.

The research in Chapter 5 extends the single node problem considered in Chapter 2 to a network setting and considers the integrated optimization of commodity procurement and processing operations over a network with multiple procurement nodes and a central processing node. Our analysis shows that solving the network problem optimally is considerably more complex and computationally hard, unlike the single node problem. We proposed two computationally efficient heuristics to solve the network problem: a) the Equivalent Single Node (ESN) heuristic, which approximates the network as a single node and b) a myopic heuristic, the Network Full Commitment (NFC) heuristic, which only considers the margin from processing and committing to sell the output immediately. Our numerical studies show that the ESN heuristic performs better than the NFC heuristic for firms operating small networks, with tight processing capacity constraints, and when the processing margins are tight. On the other hand, using a myopic policy such as the NFC is better than approximating the network as a single node, for larger networks.

The work in this dissertation addresses some of the real world issues involved in managing commodity operations. However, there are many other problems that we have not considered here which could potentially be addressed using the framework developed in this research. For instance, firms may have a choice regarding what output commodity to produce from the same input; e.g., oil refineries can refine crude oil to yield different proportions of various gasoline products. The framework considered in this research has potential to be extended to include the firm's choice of what output to produce, given current commodity prices and various operational constraints including lead times to switch production from one output to another. Another aspect of commodity operations that is not considered in this research is that of stochastic demand. While a vast literature exists on inventory management in the presence of stochastic demand, there is not much work that explores the effect of both demand and commodity price uncertainty, especially in the presence of capacity constraints. It is an interesting research topic to extend the decision making framework considered here to include demand uncertainty.

The research in this dissertation also indicates that heuristic policies that are forward looking while also incorporating the network characteristics more explicitly, are necessary to capture more of the value from operating commodity procurement and processing networks. An important extension to the present work would be to formulate improved heuristics which combine the ESN and NFC heuristics in an efficient manner to improve performance. Another direction for future research is to extend the risk-averse formulation to a network setting. An especially interesting problem in this context would be the situation where decisions in the network are made in a de-centralized manner, with each procurement location making procurement decisions independently. In such a situation, the allocation of risk over the network and the impact of decentralized decision making on overall risk are important questions for the firm.

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., T. E. Morton, D. W. Pentico. 1996. The nonstationary stochastic lead-time inventory problem: Near-myopic bounds, heuristics, and testing. *Management Science* **42**(1) 124–129.

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.

Artzner, P., F. Delbaen, J.M. Eber, D. Heath. 1999. Coherent measures of risk. *Mathematical Finance* **9**(3) 203–228.

Artzner, P., F. Delbaen, J.M. Eber, D. Heath, H. Ku. 2007. Coherent multiperiod risk adjusted values and bellman's principle. *Annals of Operations Research* **152**(1) 5–22.

Ash, M. 2011. Oil Crops Outlook. Electronic Report from the Economic Research Service, Report OCS-11b, United States Department of Agriculture.

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.

Bickel, J.E. 2006. Some determinants of corporate risk aversion. *Decision Analysis* **3**(4) 233–251.

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

Borovkova, S., F.J. Permana, H.V.D. Weide. 2007. A closed form approach to the valuation and hedging of basket and spread option. *The Journal of Derivatives* **14**(4) 8–24.

Bouakiz, M., M.J. Sobel. 1992. Inventory control with an expected utility criterion. *Operations Research* **40**(3) 603–608.

Brealey, R.A., S.C. Myers. 2003. *Principles of corporate finance*. McGraw-Hill, New York.

Brown, D.B., J.E. Smith, P. Sun. 2010. Information relaxations and duality in stochastic dynamic programs. *Operations Research* **58**(4) 785–801.

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, F., F. Gao, X.C. Chao. 2008. Joint optimal ordering and weather hedging decisions: A newsvendor model. Working Paper.

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

Choi, S., A. Ruszczynski, Y. Zhao. 2009. A multi-product risk-averse newsvendor with law invariant coherent measures of risk. Working Paper, Rutgers University.

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.

Detlefsen, K., G. Scandolo. 2005. Conditional and dynamic convex risk measures. *Finance and Stochastics* **9**(4) 539–561.

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

Duffie, D., J. Pan. 1997. An overview of value at risk. *The Journal of derivatives* **4**(3) 7–49.

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

Federgruen, A., P. Zipkin. 1984. Approximations of dynamic, multilocation production and inventory problems. *Management Science* **30**(1) 69–84.

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.

Geman, H., S. Ohana. 2008. Time-consistency in managing a commodity portfolio: A dynamic risk measure approach. *Journal of Banking and Finance* .

Goel, A., G Gutierrez. 2008. Procurement and distribution policies in a distributive supply chain in the presence of commodity markets. Working Paper, Management Department, University of Texas at Austin.

Gotoh, J., Y. Takano. 2007. Newsvendor solutions via conditional value-at-risk minimization. *European Journal of Operational Research* **179**(1) 80–96.

Heyman, D.P., M.J. Sobel. 1984. *Stochastic models in operations research, Vol. II: Stochastic optimization*. McGraw-Hill, New York.

Hu, X., I. Duenyas, R. Kapuscinski. 2004. Optimal joint inventory and transshipment control under uncertain capacity. Working Paper, Ross School of Business, University of Michigan.

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

Iida, T. 2001. The infinite horizon non-stationary stochastic multi-echelon inventory problem and near-myopic policies. *European Journal of Operational Research* **134**(3) 525–539.

Jaillet, P., E.I. Ronn, S. Tompaidis. 2004. Valuation of commodity-based swing options. *Management science* **50**(7) 909–921.

Karmarkar, U. S. 1981. The multiperiod multilocation inventory problem. *Operations Research* **29**(2) 215–228.

Karmarkar, U. S. 1987. The multilocation multiperiod inventory problem: Bounds and approximations. *Management Science* **33**(1) 86–94.

Keppo, J. 2004. Pricing of electricity swing options. *The Journal of Derivatives* **11**(3) 26–43.

Kleindorfer, P.R. 2008. Integrating physical and financial risk management in supply management. H. Geman, ed., *Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy*. 33–50.

Kleindorfer, P.R., L. Li. 2005. Multi-period VaR-constrained portfolio optimization with applications to the electric power sector. *Energy Journal-Cleveland* **26**(1) 1–26.

Kleindorfer, P.R., DJ Wu. 2003. Integrating long-and short-term contracting via business-to-business exchanges for capital-intensive industries. *Management Science* **49**(11) 1597–1615.

Kouvelis, Panos, Rong Li, Qing Ding. 2009. Inventory management and financial hedging of storable commodities. Working Paper, Lee Kong Chian School of Business, Singapore Management University.

Lai, G., F. Margot, N. Secomandi. 2010a. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations Research* **58**(3) 564–582.

Lai, G., M.X. Wang, S. Kekre, A. Scheller-Wolf, N. Secomandi. 2010b. Valuation of the real option to store liquefied natural gas at a regasification terminal. *Operations Research* (forthcoming).

Lapan, H., G. Moschini. 1994. Futures hedging under price, basis, and production risk. *American Journal of Agricultural Economics* **76**(3) 465–477.

Lovejoy, W. 1990. Myopic policies for some inventory models with uncertain demand distributions. *Management Science* **36**(6) 724–738.

Lovejoy, W. 1992. Stopped myopic policies in some inventory models with generalized demand processes. *Management Science* **38**(5) 688–707.

Mandaro, Laura. 2008. Credit, corn costs push VeraSun into Ch. 11. http://www.marketwatch.com/story/verasun-files-for-bankruptcy-after-corn-spiked-credit-tightened, 1 Nov 2008.

March, J.G., Z. Shapira. 1987. Managerial perspectives on risk and risk taking. *Management Science* **33**(11) 1404–1418.

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.

Markowitz, H. 1952. Portfolio Selection. *The Journal of Finance* **7**(1) 77–91.

Martínez-de Albéniz, V., D. Simchi-Levi. 2006. Mean-variance trade-offs in supply contracts. *Naval Research Logistics* **53**(7) 603–616.

Martınez-de Albéniz, V., J.M.V. Simón. 2010. A capacitated commodity trading model with market power. Working Paper, IESE Business School, University of Navarra, Barcelona, Spain.

Morton, T. E., D. W. Pentico. 1995. The finite horizon nonstationary stochastic inventory problem: Near-myopic bounds, heuristics, testing. *Management Science* **41**(2) 334–343.

Moschini, G., H. Lapan. 1995. The hedging role of options and futures under joint price, basis, and production risk. *International Economic Review* **36**(4) 1025–1049.

Nascimento, J.M., W.B. Powell. 2009. An optimal approximate dynamic programming algorithm for the lagged asset acquisition problem. *Mathematics of Operations Research* **34**(1) 210–237.

Paroush, J., A. Wolf. 1989. Production and hedging decisions in the presence of basis risk. *Journal of Futures Markets* **9**(6) 547–563.

Peterson, S.J., R.C. Stapleton. 2002. The pricing of Bermudan-style options on correlated assets. *Review of Derivatives Research* **5**(2) 127–151.

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.

PR Newswire. 2008. Yanglin Soybean Inc. reports Second Quarter Fy08 Results. 8 September, 2008.

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

Robinson, L. W. 1990. Optimal and approximate policies in multiperiod, multilocation inventory models with transshipment. *Operations Research* **38**(2) 278–295.

Rockafellar, R.T., S. Uryasev. 2000. Optimization of conditional value-at-risk. *Journal of Risk* **2** 21–42.

Roorda, B., JM Schumacher. 2007. Time consistency conditions for acceptability measures, with an application to tail value at risk. *Insurance Mathematics and Economics* **40**(2) 209–230.

Roorda, B., JM Schumacher, J.C. Engwerda. 2005. Coherent acceptability measures in multiperiod models. *Mathematical Finance* **15**(4) 589–612.

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.

Ruszczynski, A., A. Shapiro. 2006. Conditional risk mappings. *Mathematics of Operations Research* **31**(3) 544–561.

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. 2010a. On the Pricing of Natural Gas Pipeline Capacity. *Manufacturing & Service Operations Management* **12**(3) 393–408.

Secomandi, N. 2010b. Optimal commodity trading with a capacitated storage asset. *Management Science* **56**(3) 449–467.

Serres, C. 2008. Bumper crop of debt. *Star Tribune (Minneapolis)*. 20 July 2008.

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

Smith, J.E. 1998. Evaluating income streams: A decision analysis approach. *Management Science* **44**(12) 1690–1708.

Smith, J.E. 2004. Risk Sharing, Fiduciary Duty, and Corporate Risk Attitudes. *Decision Analysis* **1**(2) 114.

Smith, J.E., R.F. Nau. 1995. Valuing risky projects: option pricing theory and decision analysis. *Management Science* **41**(5) 795–816.

Soyatech. 2008. How the Global Oilseed and Grain Trade Works. Report prepared for U.S. Soybean Export Council and United Soybean Board by HighQuest Partners and Soyatech.

Tomlin, B. 2006. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science* **52**(5) 639.

Tomlin, B., Y. Wang. 2005. On the value of mix flexibility and dual sourcing in unreliable newsvendor networks. *Manufacturing & Service Operations Management* **7**(1) 37–57.

Tseng, C.L., G. Barz. 2002. Short-term generation asset valuation: a real options approach. *Operations Research* 297–310.

Tseng, C.L., K.Y. Lin. 2007. A framework using two-factor price lattices for generation asset valuation. *Operations Research* **55**(2) 234–251.

Walls, M.R., J.S. Dyer. 1996. Risk propensity and firm performance: a study of the petroleum exploration industry. *Management Science* **42**(7) 1004–1021.

Wang, T. 1999. A class of dynamic risk measures. Working Paper, University of British Columbia.

Wu, DJ, P.R. Kleindorfer. 2005. Competitive options, supply contracting, and electronic markets. *Management Science* **51**(3) 452–466.

Wu, O.Q., H. Chen. 2010. Optimal control and equilibrium behavior of production-inventory systems. *Management Science* **56**(8) 1362–1379.

Zhu, W., R. Kapuscinski. 2006. Optimal operational versus financial hedging for a risk-averse firm. Working Paper, Ross School of Business, University of Michigan.