OPTIMIZATION OF COVERAGE RATIOS FOR BASIS AND FREIGHT TRADING PORTFOLIOS

A Thesis
Submitted to the Graduate Faculty
of the
North Dakota State University
of Agriculture and Applied Science

By

Adam Joseph Kroll

In Partial Fulfillment of the Requirements
for the Degree of
MASTER OF SCIENCE

Major Department:
Agribusiness and Applied Economics

May 2021

Fargo, North Dakota

North Dakota State University
Graduate School

Title

Optimization of Coverage Ratios for Basis and Freight Trading Portfolios

By

Adam Joseph Kroll

The Supervisory Committee certifies that this ***disquisition*** complies with North Dakota State University's regulations and meets the accepted standards for the degree of

MASTER OF SCIENCE

SUPERVISORY COMMITTEE:

Dr. William Wilson
Chair

Dr. David Bullock

Dr. Fariz Huseynov

Dr. Frayne Olson

Approved:

June 14, 2021
Date

Dr. William Nganje
Department Chair

ABSTRACT

Changing commodity prices create opportunities for traders to profit or experience substantial losses if risk is not managed. Extensive research has been presented regarding optimal hedging strategies using futures markets to mitigate price risk of holding a physical commodity, but the literature has not introduced a method for optimizing basis or rail coverage ratios. This study introduces coverage ratios to manage basis and transportation risk. Coverage ratios are optimized for a portfolio of basis and rail positions consistent with expected utility theory using Monte Carlo simulation at various risk aversion levels. The results indicated that as risk aversion increased, optimal coverage levels increased while profit and standard deviations of profits decreased. Sensitivity analysis is used to demonstrate the effects of changing intermarket correlations, standard deviations of individual markets, and time to liquidation. This study provides insight into managing basis and rail portfolios and presents implications for traders and risk managers.

ACKNOWLEDGEMENTS

My first lesson in economics involved scarcity. While I am thankful for the many subsequent lessons in economics that followed, I may be most grateful for that first lesson as it has allowed me to see the world in a new way and more fully appreciate the significance of someone choosing to spend their scarce time with me. I will always be grateful for the people that chose to teach, mentor, and befriend me throughout my academic career.

I would first like to thank my teachers. To my elementary school teachers – thank you for instilling a passion for learning and making me feel at home in a classroom. The foundation I was given at Holy Trinity prepared me for success both academically and in life. I would also like to thank the teachers in my high school and at NSDU who went above and beyond to challenge me and help me succeed when I struggled. I owe a special thanks to my committee members who provided valuable insight and feedback during the completion of this research and to my advisor, Dr. William Wilson. His guidance and support in my research and professional endeavors has been invaluable, and I am honored to have had the opportunity to learn from him.

I would like to thank my classmates and friends for the energy, excitement, and perspective that they added to my college experience. I would especially like to thank the members of FarmHouse Fraternity and my roommates there for their friendship and support as I pursued my Master’s Degree.

Lastly, I would like to thank my parents, Duane and Linda, for their support of my academic pursuits and for encouraging me to think big. I will forever be grateful for their example and the many lessons I have learned from them.

DEDICATION

To my late grandfather, Richard Aschenbrenner, who encouraged me in my academic pursuits from early on even though he never had the opportunity to pursue education beyond the eighth grade. Of everything I learned from him, the most important lesson may have been the desire to learn.

"Learn as much as you can; a good education is the one thing no one can take away from you."

July 1993 – November 2020

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1. INTRODUCTION

1.1. Introduction

Finding the balance between risk and reward is a task that no company can escape. While it is well understood that, in general, more risk leads to higher rewards, decisions of how to effectively bear and manage different levels of risk are more difficult. Additionally, the structure of companies and the goals of shareholders, board members, and management can add another element to the equation as bearing risk adds a higher level of variability to returns.

Commodity trading businesses face a variety of risks due to risks in prices, logistics, grain quality, and other factors. Additionally, the financial risk associated with each of these increases as the size of the firm's positions increase. As a result, there is a need to manage risk in a way that allows desired returns to be achieved without subjecting the firm to unacceptable levels of risk.

Futures price risk has been discussed in great detail within academic literature (see Wilson 1982; Blank, Carter, and Schmiesing 1991; Harris, Shen, and Stoja 2010; Anderson and Danthine 1981; Chen, Lee, and Shrestha 2003; Howard and D'Antonio 1984; Cecchetti, Cumby, and Figlewski 1988). These publications cover topics such as hedging, hedging with options, cross hedging, portfolio management, optimal hedging ratios, and other topics related to managing futures risk. While discussion of futures markets is prevalent in the literature, less research has been conducted regarding basis risk and optimizing basis portfolios. Furthermore, basis coverage ratios have yet to be shown in the literature and defined for individual markets or a portfolio of markets.

1.2. Introduction to Commodity Markets

Elevators conducting trading business often operate within multiple markets. At a minimum, many are involved in the futures market and the basis market. Some elevators may not want exposure to the futures price or basis risk and may take action to eliminate them using hedging or forward contracting. This section introduces futures and basis markets as well as some of their key components and risk.

1.2.1. Futures Markets

The futures market is an organized market in which buyers and sellers exchange futures contracts. A futures contract is formally defined by the Chicago Mercantile Exchange (CME) as a financial instrument that allows a market participant to transfer or acquire the risk of price changes in an asset over time ("Definition of a Futures Contract - CME Group" 2021). Futures contracts exist for a variety of commodities including metals, currencies, lumber, animals, and grains, and in the United States, these contracts are primarily traded on the Chicago Board of Trade (CBOT). While there are many similarities between the market mechanisms and applications of the theories discussed throughout, this study focuses primarily on grains, specifically CBOT soybeans.

Futures contracts are highly standardized. Exchange rules specify the type, quantity, quality, times of trading, contract expiration, and delivery mechanisms for each contract (Kolb and Overdahl 2006). This allows for an efficient means of transferring contracts as the only term left to decide upon is the price of the contract. Throughout this study, the term "futures price" is defined as the current price of the futures market for a given contract or a price that has been set because of owning or selling a futures contract.

Futures markets are global markets for a generic contract specification regarding delivery time, location, quality, and other terms. Market prices are affected primarily by macro-level fundamental factors including changes in global supply and demand, weather (which affects the supply of grain), government subsidies, substitutes, and currency exchange rates (“Fundamentals and Agricultural Futures - CME Group” 2021). These large changes are uncontrollable by any market participant, and it is assumed that market participants have access to nearly identical market information.

1.2.2. Basis Markets

Another type of market that farmers and trading companies participate in is the basis market. The basis market is a localized market with its own set of fundamentals, but the basis market is connected to the futures market. Basis is defined as the difference between the cash and futures price. Cash price is the amount that an elevator pays or receives when grain is priced (Lorton and White). For example, if CBOT March 2021 Soybean futures were trading at $9.00/bushel and the cash price for March delivery in St. Cloud, MN was $8.50, the basis at St. Cloud would be -0.50, or “50 under” the March contract. Meanwhile, in Minneapolis, MN a barge terminal location may be bidding $8.75, so the basis in Minneapolis would be -$0.25 under the March contract. The reason the basis values differ may be a result of the opportunities each elevator has to sell the commodity, the additional cost of pulling the grain from farmers in the rural part of MN to the terminal in the large city because of transportation costs, and the competition that each elevator faces in its respective location. Changes that occur spatially and across time create opportunities for merchandisers. They can purchase grain at low basis levels, hold the grain, and sell it once the basis has appreciated or they can make advance sales of the grain and purchase the grain at a later time with hopes of a decreased basis level.

When a firm or farmer owns grain, they hold the risk of price movement in both markets. Because the dollars at risk increase as the size of the position increases, risk management become increasingly important as the firm holds larger positions. Each market has mechanisms for transferring risk from one party to another.

1.3. Rail Markets

An important element of grain merchandising involves grain shipping. The previous section suggested that the basis market works to pull grain from areas of surplus to areas of deficit. Merchants or traders recognize that in order to take advantage of the differences between origin and destination markets, they must also have a solid understanding of transportation and logistics mechanisms as well as the associated risks. A key risk that merchants face in rail risk is car placement and the prices of the transportation at various times.

Risks associated with rail car pricing, notably the secondary car market, are important as shipping costs can cause significant swings in the overall elevator margins. If grain is purchased at the origin and is priced at the destination for delivery at a later date, but the rail price is unknown, the net profit from the trade remains unknown and at risk.

The risk associated with rail pricing can be mitigated using different rail purchasing strategies. Rail cars may be purchased on either the primary market or the secondary market. The primary market allows for users of railcars to reserve cars far in advance for a random number of trips per month for a few months or a year. If a user of the railcars does not need all of the trips, the rights to use the extra trips can be sold through brokers and auctions on the secondary market. Additionally, if a merchant has not purchased freight in advance or needs to purchase additional freight, they may purchase trips on the secondary market. Using the primary and secondary markets, merchants can create a purchasing plan which can mitigate risk.

1.4. Hedging and Hedge Ratios

The futures market was briefly introduced in a previous section. Within the futures market, there are two types of traders – speculators and hedgers. A speculator is a trader who desires to acquire risk with the hopes of making a profit from the change in the asset value (Kolb and Overdahl 2006). A speculative trader may have a market bias, or thoughts on the direction of a market, and may trade accordingly without owning any of the physical commodity. In contrast, a hedger enters the marketplace in an effort to reduce or transfer the risk to other market participants (Kolb and Overdahl 2006). This type of trader either has the commodity and wants to reduce the risk of price changes that may occur while holding the commodity or plans to acquire the physical commodity at a later date and wants to reduce the fluctuations that could occur before acquiring the commodity.

While hedgers trade in a manner that allows them to reduce risk by taking opposite positions in the futures market as they hold or plan to hold in the physical market, hedging the entire position and eliminating all futures price risk may not be the profit-maximizing strategy. The proportion of the cash position that is hedged is defined as a hedge ratio. This is represented as the quantity of bushels sold in futures divided by the total cash position. Researchers have conducted research showing that an optimal hedge ratio may be found given the firm’s tolerance for risk and the correlation between cash and futures prices (see Kahl 1983; Blank, Carter, and Schmiesing 1991). A variety of methods have been used to determine optimal ratios such as minimum-variance, mean-variance, mean-semivariance, mean-Value-at-Risk, among others.

1.5. Basis Trading and Coverage Ratios

Because of the volatility and uncertainty associated with futures trading, many elevators and traders only trade basis on their physical bushels to make consistent profits (Lorton and

White 2010). Pure basis traders hedge 100% of their grain ownership or forward sales in the futures market and rely on basis trades to enhance margins. A basis trader's goal is to buy grain at a low basis and sell at a high basis. The trader may choose to create a long basis position by acquiring grain at the origin to hold with hopes of selling to a destination once basis levels are higher. Merchants may also take short basis positions by making forward sales to a destination for a deferred delivery window with the intent of buying at the origin just prior to making delivery to the destination.

A basis trader may desire to have various levels of basis or transportation price risk; this analysis uses coverage ratios to describe the amount of a purchase or sale of grain that is at risk to changing price levels. Coverage ratios can be defined as the portion of grain sold to the destination relative to the grain owned at the origin; the quantity of grain bought relative to the quantity of grain needed to be purchased; or for rail, the quantity of rail purchased relative to the amount of rail needed.

1.6. Problem Statement

If an elevator is completely and properly hedged, futures risk is eliminated leaving basis and rail price movements as the main sources of margin risk for shipper elevators. Each of these markets can exhibit varying levels of volatility and randomness which in turn creates randomness in the elevator's profits. While the randomness due to basis and rail price movements could be eliminated by using forward contracting and purchasing rail in advance, this type of strategy would eliminate many possible merchandising opportunities and significantly decrease profits. Conversely, if merchants are allowed to take very large positions without regard for the risk of the position, the company may not be able to withstand the losses.

While hedge ratios have been formally defined for managing the futures price risk, the literature lacks a discussion on mitigating basis or rail risk. This research proposes a definition for a coverage ratio and further shows that coverage ratios for basis and rail can be optimized to optimize the utility of the trader or firm.

1.7. Objectives and Procedures

The purpose of this thesis is to develop a methodology for balancing risk related to holding cash grain and rail positions in a portfolio relative to the firm's risk tolerance levels. Specific objectives are subdivided as follows:

1. Define coverage ratios for long and short basis positions as well as rail positions. This allows for models to be developed which determine profit and standard deviation of profit by taking or not taking coverage in the market.
2. Determine optimal coverage ratios for a base case scenario that illustrate the practicality and usefulness of an optimal coverage ratio.
3. Conduct sensitives to show the impacts on coverage ratios, profit, and standard deviation of profit when variables such as standard deviations of basis, standard deviation of rail, time to liquidation, or restrictions on allowable coverage are changed.

To meet the objectives, theoretical models are developed to show the expected payoff functions from a variety of merchandising scenarios. Next, the variance of the payoff function is derived which can then be substituted into an expected utility function. An analytical solution is found by maximizing the expected utility. The empirical analysis follows and uses Monte Carlo simulations in conjunction with RiskOptimizer™ within the Palisade DecisionTools Suite™ to estimate the optimal coverage ratios for each model. The models are evaluated using historical

BestFit™ distributions to illustrate a "naïve" strategy and time series forecasted distribution to illustrate a merchant that captures an "anticipatory" trade which is based on expected changes in basis and shipping costs.

This thesis contributes to the literature by discussing basis and rail risk management and offering a new procedure for managing these sources of risk. Previous literature has focused primarily on optimizing hedging ratios. This research differs in that the focus is completely on the physical grain and freight positions. Within the thesis, coverage ratios are defined and optimized for basis and rail risk. Additionally, the importance of the correlations between origin and destination basis as well as the correlations between the basis and rail markets are discussed and shown empirically using sensitivity analysis.

1.8. Organization

This thesis is organized into six chapters. This chapter gave an overview of futures, basis, and rail markets as well as an introduction to hedging and hedge ratios. Chapter 2 discusses the background and related literature which explains the details of basis and rail trading as well as types of portfolio hedging models. A variety of portfolio hedging models are discussed, but minimum-variance, mean-variance, and mean-Value-at-Risk are the focal point. Next, Chapter 3 discusses the theoretical framework, introduces the models and variables, and presents an analytical solution for the optimal coverage ratios in the single market and multi-market case. The empirical models and data are shown in Chapter 4. Chapter 4 also highlights the distributions used in the simulation models, as well as the simulation procedures and settings. All of the models are evaluated under a mean-variance and mean-semivariance framework at varying levels of risk aversion. Additionally, each model uses both historical BestFit™ and time series forecasted distributions for its random variables. Chapter 5 includes the results from each of the simulations as well as a variety of sensitivities. The sensitivities show how changing standard deviations, correlations, and timing of the positions can change the results. Lastly, Chapter 6 presents the conclusions and implications of the study as well as suggestions for further research.

CHAPTER 2. BACKGROUND AND RELATED LITERATURE

2.1. Introduction

Understanding risk exposure and risk management tools is essential in the competitive grain trading industry. Taking too little risk may leave a firm falling behind the competition while taking too much risk may result in larger losses than a firm can withstand. In grain trading, there are two primary elements of price risk: risk from futures prices and basis levels. There are a variety of ways to add or reduce risk in basis trading and several ways to model the risks.

This chapter describes relevant background information and summarizes literature relating to portfolio hedging theory. The first section introduces basis and rail trading and highlights the risks and opportunities associated with each of these types of trading. Next, arbitrage is discussed to give the reader an understanding of the unique opportunities that are available through logistical and merchandising management. Then, an overview of general risk in commodity trading follows which describes the main sources of risk for a shipper elevator. Lastly, the evolution, developments, and applications of portfolio hedging models are presented along with their application to this study.

2.1.1. Basis Trading

The forces of supply and demand can be observed both globally and locally in grain markets. The overall supply and demand for a given commodity typically affects the futures market price where the localized supply and demand affects the basis. Basis is defined as the difference between the cash price and futures price (Lorton and White 2010). By hedging, the futures portion of the risk is eliminated, and elevators become “basis traders” (Lorton and White 2010). Throughout this study, the terms “basis traders”, “traders”, and “merchants” are used similarly.

As the name implies, a basis trader looks to trade grain intending to make money from the changes in the basis. This can be accomplished by taking a "long" basis position by buying at low basis levels and selling at high basis levels or conversely by taking a "short" basis position by selling at high levels in advance and purchasing the grain later when the market values decrease. As basis increases, or becomes more positive, the basis is said to be strengthening. As basis decreases, or becomes more negative, the basis is said to be weakening. Traders can exploit these basis movements and increase elevator margins by taking positions in the basis market. Both of these strategies involve risk that the market moves against the trader's position.

An alternative strategy is a "back-to-back" transaction that involves purchasing grain at a set basis, adding in a handling margin, and immediately contracting the sale of the grain (Lorton and White 2010). This strategy reduces the risk of unfavorable market movement but eliminates the opportunity to make money from the changing basis.

While the riskless alternative may initially seem desirable, the effects of basis trading can be easily demonstrated with an example. Suppose an elevator has 2,000,000 bushels of grain that need to be bought and sold throughout a year. Assume this elevator remains completely hedged in the futures market, thus eliminating futures price risk. A riskless basis approach with a back-to-back sale at a $0.10 handling margin would generate $200,000 in revenue where even capturing an additional $0.15 from the market movement plus the $0.10 handling margin would generate revenues of $500,000. Additionally, basis trading may allow elevators to handle larger grain volumes by taking advantage of isolated buying or selling opportunities without the need to buy and sell simultaneously. With larger elevators and increased basis volatility, the importance of basis trading further magnifies.

Several factors influence basis movements. Cost of shipping to a destination market, local elevator margins, cost of storage, cost of carry, quality variations in different geographies, and localized supply and demand all can affect the basis market (Lorton and White 2010). Some of these factors are predictable, while others can change rapidly. If an elevator has a train or barge that must be filled, the merchant may temporarily increase the basis to incentivize quick movement of grain to the elevator to avoid paying demurrage, or a late fee, to the rail or barge company. Additionally, Baldwin (1986) says crop quality and conditions, transportation availability, storage availability, storage cost, and seasonality are factors that can affect the basis market. If the crop quality in an area is low, merchants may lower the basis to encourage farmers to store the grain until they have higher quality grain to blend it with. Transportation availability is an important factor for basis because if freight is expensive or unavailable, the elevator must lower basis levels to maintain their desired margins. Because basis is used as a way to control the flow of grain, basis tends to be seasonal and is related to the cost and availability of storage. When storage is relatively cheap and plentiful, the basis can be stronger than if storage is expensive and/or unavailable. This helps explain why during harvest, when there is a surplus of local supply and a shortage of storage, basis tends to be weaker than before harvest where storage is plentiful and supply is diminished.

The variations in the basis market caused by fundamental market factors provide opportunities for trading and profit creation. In fact, most merchandisers act as basis traders because of the opportunities to create higher returns with less risk than alternative grain trading methods (Lorton and White 2010). This means that the common gross margin an elevator captures is the difference between the buying and selling basis levels.

Knowing that basis is dynamic and that there are riskless and risk-bearing trading methods, the need for a method to determine a coverage ratio that gives the elevator the right amount of risk to maximize their opportunities becomes clear. The research conducted within seeks to create a theoretical framework and empirically derive an optimal coverage ratio for an elevator with a specific portfolio of basis origins and destinations as well as transportation. This allows merchants to better position themselves for maximum trading profits within given risk tolerances.

2.1.2. Rail Trading

While a large portion of the risk that an elevator faces is related to the basis market, transportation plays a critical role in merchandising activities and can be a source of risk for the firm. Because there are dynamic markets for commodities and consumer goods which require transportation, the prices of transportation can vary based on the shipping demand. Like the basis trading transactions, and thanks to the relatively recent development of car allocation methods, merchants can bear or mitigate the risk associated with rail freight prices.

Prior to 1987, there was no method for reserving railcars in advance (Wilson, Priewe, and Dahl 1998). This created uncertainty for merchandisers because of the rapid changes in prices and availability of cars (Gelston and Greene 1994). In 1988, the Certificate of Transportation (COT) program began which allowed the guaranteed the use of a quantity of shipments (Wilson and Dahl 2011). Later, long-term shipping contracts were developed that specified the number of pickups over a given time period. These rights to a shipment were tradeable and had a system of penalties to incentivize the efficient usage of the cars. This system of tradable freight is often known as the secondary market.

More recently, a number of studies analyzed the relationships between basis and secondary rail market values. Bullock and Wilson (2019) quantified how shipping costs and the secondary market impacts the soybean export basis. Lakkakula and Wilson (2021) analyzed the interrelationship between origin and destination basis and secondary car values. Wilson and Lakkakula specifically analyzed how the secondary market impacts basis values. This builds on an earlier study by Wilson and Dahl (2011) who had earlier illustrated the impacts of rail car values on basis values. Taken together, these results illustrate that the basis and secondary rail values are correlated, among fundamental factors, and that changes in rail car values impact both origin and destination basis.

2.1.2.1. Rail Freight Pricing in the Primary and Secondary Markets

Rail freight is subject to a variety of supply and demand factors that influence the prices and availability of freight. Additionally, the cost and availability of other transportation options may affect the cost of rail freight. Rail freight is traded in a "primary" and "secondary" market. Merchandisers may purchase freight on the primary market by bidding for a certain allocation of cars or trips with a rail company, paying for service at the posted tariff rate plus any fuel surcharges ("Rail Service Challenges in the Upper Midwest" 2015). The winner of the auction receives a certain number of guaranteed trips with those cars within a given time period. A situation may arise where the elevator cannot utilize all the trips they have purchased or need additional trips. The secondary car market provides the solution to this problem.

The secondary car market is comprised of buyers and sellers that work together through a third-party brokerage service (Landman 2017). The bids and asking prices are typically gathered anonymously and are quoted in dollars above or below the tariff prices. In contrast to the primary market, secondary market allocations are typically for one trip only and have guaranteed shipment windows. Typically, the tariff rates remain relatively stable compared to the rates of the secondary car market. The different pricing mechanisms of the primary and secondary markets allow for efficient allocations of cars and give merchants tools to manage, acquire, and mitigate freight risk.

2.1.2.2. Risk in Rail Trading

Rail trading has two primary sources of risk: car placement risk and price risk (Wilson, Bullock, and Lakkakula 2020). Car placement risk exists in the primary market because there is uncertainty around the exact timing and quantity of rail cars that are placed. This depends on the velocity, which is volatile, and is a key attribute of the primary market for rail cars. Car placement risk is not a problem in the secondary market because the seller is obligated to place the cars within the specified period. This thesis seeks to manage risks associated with rail car pricing rather than rail car velocity or placement risk.

The risk of rail freight pricing begins at the moment an elevator enters a position in which rail freight will be required. When an elevator is long grain, they become short freight because they will need to sell and ship the grain at a future time (Wilson, Priewe, and Dahl 1998). Thus, once the freight or grain position is established, the firm is exposed to the risk of changing rail rates. Elevators typically utilize forward rail marketing and purchase in the primary market in advance or rely on secondary market purchases of freight. The primary market and tariff rates experience few changes where the secondary market can have periods of large price fluctuations (Landman 2017). Exposure to price risk in the rail market can result in dramatic changes in the elevator's revenues.

2.1.2.3. Effects of Rail Rates on Cash and Grain Trading Activities

Across the available literature, there is a general agreement among researchers that rail performance, reliability, and pricing affect merchandising activities and cash grain prices. Studies such as Olson (2014), Usset (2014), and Ortiz (2016) show significant connections between cost and availability of transportation and cash grain prices.

Usset (2014) examines how freight delays and disruptions can impact farm revenues. He studied the effects of rail transportation disruptions on basis levels in corn, soybeans, and hard red spring (HRS) wheat in various districts in the state of Minnesota and estimated the effects of those basis changes on-farm revenues. The analysis was completed by identifying key market fundamental factors such as stocks-to-use ratios and Minnesota ending stocks numbers for each commodity and finding analogue years that mirrored those fundamental patterns. The earlier years with similar patterns served as a control against the 2013/14 crop year, the year of interest. Usset found that in the presence of freight disruptions, basis levels decreased 50-80 cents, 50-70 cents, and 30-50 cents per bushel for soybeans, corn, and wheat, respectively.

Around the same time, Olson (2014) studied the relationship between rail freight rates and farm revenue by using analogue years. He used the 2009/10 marketing year as a base year because of the similar market fundamentals to the 2013/14 marketing year and estimated the differences between the various commodity prices. While the report does not include an impact of rail disruptions on a per bushel basis, Olson does conclude that there were approximately 66.6 million dollars in lost farm revenue in North Dakota from grain that was sold between January and April 2014. This further strengthens the argument that transportation markets have significant and economically relevant effects on basis markets.

While Usset and Olson focused on generalized farm revenues across geographical areas due to generalized freight disruptions, other research conducted by the Agricultural Marketing Service (AMS) in 2015 and Ortiz (2016) study the impacts of transportation on basis. Both studies suggested that increased rail cost decreased basis. The AMS concluded that increased rail prices may have lowered corn, soybeans, and wheat local prices by 11 to 18 cents per bushel. Ortiz (2016) agrees that the link between grain basis and rail transportation is "utterly uncontroversial", and showed that for each additional dollar per barrel of oil, the basis in wheat decreases 1 to 8 cents per bushel on average. The study by Ortiz is different than the others, but it helps show that changing demand and prices for rail cars because of other commodity shipping demand impacts the basis of commodities in localized areas. Overall, the literature shows a consensus that basis and rail freight rates are related.

2.1.3. Arbitrage

Understanding the concepts of basis and rail trading opens yet another opportunity for merchandisers: arbitrage. While the previous trading strategies involved either back-to-back trades or becoming short or long the basis, arbitrage is a trading strategy that can result in large profits for traders without risk or investment (Kolb and Overdahl 2006). Kolb and Overdahl (2006) describe academic arbitrage as trades that occur simultaneously, thus eliminating any investment, and without risk.

A variety of scenarios exist which may create arbitrage opportunities. One case where arbitrage may exist is in the delivery market as CBOT contracts reach expiration. During the delivery period, the seller (who holds a short futures position) notifies the clearing firm of their intention to deliver. The clearing firm then notifies the oldest long of the seller's intention to make delivery. Deliveries are made to registered warehouses that are approved by the CBOT for a fee. During the delivery period, basis is expected to equal the delivery charge. If basis is greater than the delivery charge, an arbitrage opportunity exists to buy futures, take delivery, and sell cash grain. As a result, futures prices increase and basis decreases until the basis is equal to the delivery charge. Alternatively, if basis is less than the delivery charge, a merchant could purchase cash grain, sell futures, and make delivery.

Kolb and Overdahl state that in a “well-functioning market”, arbitrage opportunities do not exist. Knowing this, what warrants the discussion of arbitrage? Arbitrage is discussed and remains important because markets can experience imperfections, and these imperfections can present opportunities for merchandisers. The law of one price would suggest that identical goods should have identical prices; the presence of arbitrage opportunities with transaction cost is a reason that this law holds (Lamont and Thaler 2003). A common form of arbitrage is spatial arbitrage. In this form of arbitrage, grain is bought and sold at the same time for a profit of the difference less the transportation, handling, and transaction costs (Kub 2014). Skadberg et al. (2015) found that spatial arbitrage opportunities existed for certain North Dakota grain elevators trading soybeans and that the payoffs from arbitrage opportunities varied based on the elevator locations.

2.2. Risk in Commodity Trading

Before beginning the discussion of optimal coverage ratios and grain trading strategies, a formal definition of risk and the scope of the research is necessary. Risk involves a probability of an event occurring where uncertainty occurs because of a lack of information about the probability of outcomes related to a particular event (Vose 2008). This means that to effectively manage risk, the possible outcomes must be measurable and have a probability distribution. Without this information, the future is left with uncertainty which is harder to manage and model with a high level of precision. Understanding this distinction allows for a further discussion on the risks faced by merchandisers.

Arguably, some of the largest risks faced by a shipper elevator are risks related to basis values, inventory risk, and shipping costs which were discussed earlier. Basis risk results from the changes in the basis values in the origin and destination markets. One way of measuring this risk is to use the historical standard deviation of observed basis values over a particular time frame. A high standard deviation shows that there is a higher level of risk than an area or period with a low standard deviation. While individual elevators are responsible for setting their own basis values and agreeing to basis levels at a destination, changing market conditions and competition from other elevators make this a random variable.

Another source of risk for elevators is inventory risk. While there are a variety of contract solutions that elevators offer which can help regulate the flow of the grain to the elevator, many bushels are still delivered at the "spot" price. The spot price is the cash price at an elevator at any given time for which a farmer can sell their grain. Farmers typically may deliver at any point and transfer ownership to the elevator. As prices increase, economic theory suggests that the quantity supplied by the farmers would increase. The cash price is comprised of the futures plus the basis. The elevator has no control over the futures price but can adjust basis values to influence grain flows.

While price has a large impact on the grain inventory levels, other factors can affect it as well such as weather or seasonal influences which determine farmers want to haul grain. For example, during the spring, farmers may haul grain if the ground is too wet to plant but may stop hauling as soon as planting resumes. This can create risk for the elevator if they are unable to procure enough grain to meet their shipping obligations and could result in lost opportunity if they are unable to take delivery of farmer grain due to limited storage space. Because these inventory levels are uncontrollable, farmer deliveries can be considered a random variable.

Table 2.1: Risk in Commodity Trading Summary

Risk Name	Description	Risk Management/Mitigation
Basis Values	Creates risk for margin when basis values decreasing when holding long cash grain position	Purchase/sell grain in advance at known levels
Shipping Costs	Creates risk for margin when freight values increase when needing to purchase freight	Purchase freight in advance at known levels
Inventory	Elevators may purchase grain from farmers in advance using contracts or may purchase grain that farmers deliver without contracts	Require smaller delivery periods in forward contracts and incentivize forward contracting
Quality Risk	The quality of grain in storage may decay over time because of mold, moisture, insects, or other factors	Ensure elevator operations team tracks grain temperatures and quality to load out, aerate, or fumigate problematic bins

The sections within this chapter thus far have highlighted some of the key risks that merchandisers face due to changing basis levels, rail freight prices and availability, as well as inventory risks associated with farmer spot deliveries as summarized in Table 2.1. Given the clear risks associated with pure basis and rail trading, the need for risk management tools is clear. The remainder of this chapter focuses on previous work that has been done regarding portfolio hedging models to understand the potential applications to basis and rail trading risk.

2.3. Portfolio Hedging Models

The research in this thesis builds on previous work in the area of portfolio selection and optimization which lays the foundation for determining optimal coverage ratios. Portfolio theory relates risk and reward among a variety of investment alternatives. While the concepts of risk and reward could be dated to biblical times, the theory was only relatively recently formalized for investments in academic literature by Markowitz (1952) who introduced the concept of portfolio theory as a way to explain the desire for investors to achieve the largest possible return for each respective level of risk, which is represented by the variance of returns. This led to the representation of these two factors for various investments and is called the "expected returns – expected variance of returns" (E-V) rule.

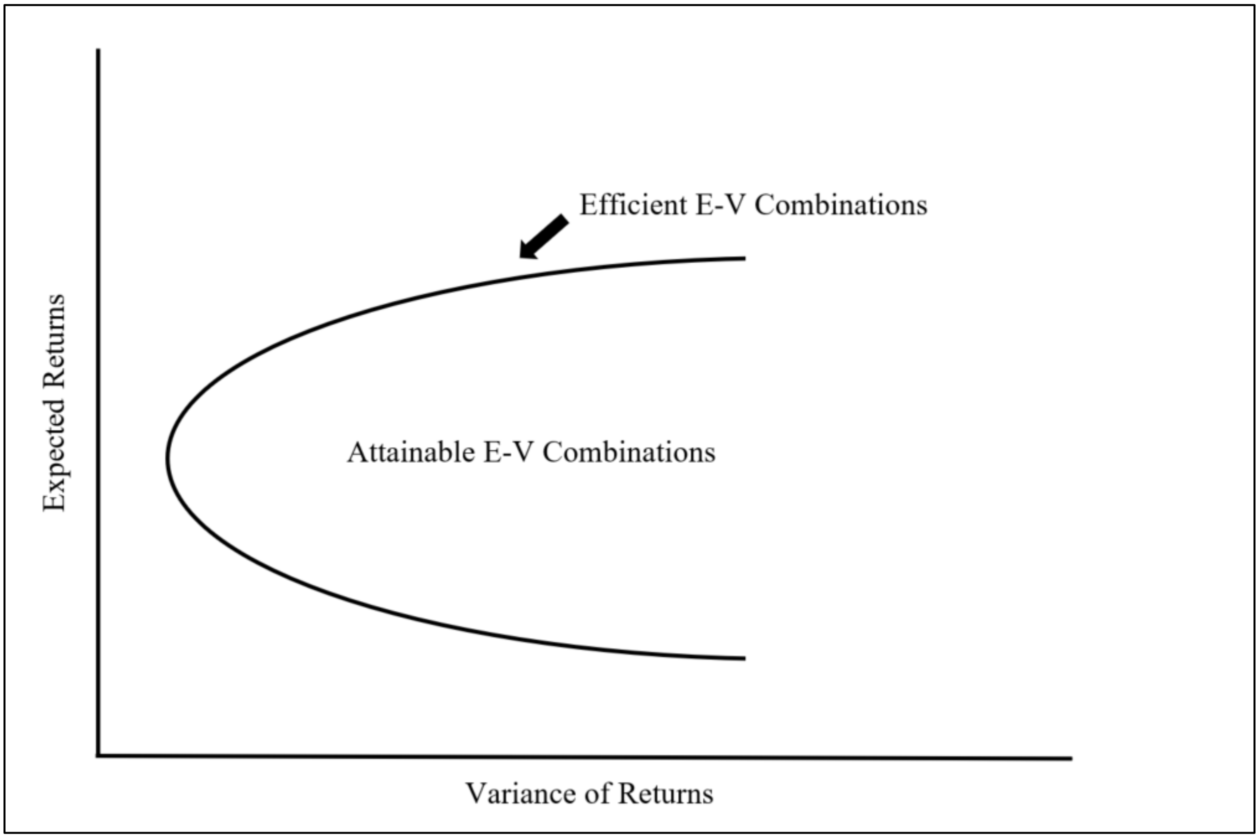

Figure 2.1: E-V Combinations

Figure 2.1 is based on Markowitz (1952) and shows a geometric illustration of this rule with an efficient frontier for combinations of expected return versus expected variance and demonstrates the alternatives for a decision-maker. An investor wanting to experience large returns would likely experience a high level of variance where an investor wanting less risk may have to sacrifice the higher levels of returns to achieve the investment objective. This showcases the tradeoffs the decision-maker encounters. Identifying optimal coverage ratios also involves portfolio optimization; the various approaches that follow guide the development of the framework used to create the objective function.

The following sections outline the evolution of portfolio hedging models starting with the minimum variance models. Mean-variance models follow along with value-at-risk (VaR) and conditional-value-at-risk (CVaR) frameworks. Lastly, we mention some other modeling strategies that have been used but have not been as widely adopted by academics or practitioners. Throughout the following sections, the pros and cons of each of the models are discussed as well as their applications or contributions to the existing literature.

2.3.1. Minimum-Variance Model

The minimum variance model is a type of model designed to reduce the variance, or risk, associated with price movement. Generally, investors or decision-makers have objectives other than to minimize or eliminate risk because investors seek to make a profit, and only a select number of opportunities exist where arbitrage, or riskless profit, is available as discussed in previous sections. This type of model fails to reflect the desired speculative returns of the trader or investor.

Johnson (1960) describes traditional hedging as a tool for risk reduction by taking opposite positions in correlated markets, and he suggests that traders tend to mix hedging activity with their speculative trading activity. This could take the form of waiting to place hedges on grain inventory because of a bullish bias or selling futures, even at times in excess of the inventory held, because of bearish market bias. As a result, Johnson (1960) calls for a more formalized hedging definition as follows: In the case of a certain quantity of physical units held in a market, a hedge is a position in another market such that the price risk encountered by holding the physical units across time is minimized. This definition of hedging is useful as the primary goal of a hedge is to minimize the price risk of the underlying physical asset. Anderson and Danthine (1981) state that traditional hedges which involve equal and opposite positions in the physical market and the futures market are generally suboptimal if it “cannot be identified as a risk-minimizing position or as an optimal speculative deviation from the minimum-risk position”.

Johnson (1960) created a minimum variance model that explained theoretically how hedgers take positions in order to reduce the variance resulting from changing prices in a given commodity. Under this framework, the pure hedger would sell futures at the time of taking ownership of a physical commodity, even with a bullish bias, because it eliminates the price risk associated with holding the physical commodity.

Ederington (1979) used a similar approach by empirically deriving optimal hedge ratios for Government National Mortgage Association (GNMA) 8% Pass-Through Certificates, Treasury-Bills, wheat, and corn with their respective futures contracts under the minimum variance approach. In both Ederington's (1979) and Johnson's (1960) research, they find the optimal hedge ratio to be covariance between the spot and futures market divided by the variance of the futures market.

When futures markets follow martingale processes, resulting in zero expected returns, the optimal hedge ratio must only minimize the portfolio variance (Harris, Shen, and Stoja 2010). In contrast, where futures have non-martingale patterns, the hedged portfolio is comprised of a minimizing risk position and a speculative position (Anderson and Danthine 1981). Even in martingale processes, an investor may have access to information unrelated to previous price data such as better weather predictions or knowledge of market fundamentals which may result in expected profits resulting from a speculative position.

When reviewing the results, one must evaluate hedge effectiveness which can be measured by the square of the correlation between the two markets of interest (Kimura 2016). A perfectly effective hedge is one in which the price movements in the two markets across times are perfectly correlated (i.e. have a correlation equal to one). Thus, any loss/gain in one market would be perfectly offset by the gain/loss in the other market when taking offsetting positions (Johnson 1960). Knowing this, traders desire to hedge in highly correlated markets. As the correlation between the markets increase, the effectiveness of the hedges also increase. In Ederington's (1979) empirical analysis, the results showed that Treasury Bills were less effective than GNMA futures for short-term hedging. The results also compare the effectiveness of hedging wheat and corn.

The minimum variance hedging model works when returns are elliptically distributed or investors have quadradic utility, but when these conditions are not met, alternative approaches are needed as the variance is no longer an appropriate measure of risk (Harris, Shen, and Stoja 2010). Additionally, as the joint distribution of cash and futures changes, the optimal hedge ratio will be calculated incorrectly if left unaccounted for (Cecchetti, Cumby, and Figlewski 1988). Other approaches that have since followed include a variety of Value-at-Risk models (VaR) such as minimum-VaR, and minimum-CVaR, as well as others that are discussed later in this chapter.

2.3.2. Expected Utility and Mean-Variance Hedging

Expected utility theory can be used to optimize a decision-maker’s expected utility in situations involving risk and reward. The Von Neumann – Morgenstern (VNM) axioms are an important part of the assumptions required for expected utility theory analysis. The axioms relate to the ordering, transitivity, continuity, and independence of the risky alternatives (Von Neumann and Morgenstern 1947). First, it is assumed that the decision maker has a preference between two options, A or B, or is indifferent between them. Next, transitivity of the options means that if the decision-maker prefers A to B and prefers B to C, then A is also preferred to C. Additionally, if A is preferred to B, and B is preferred to C, then there is a probability that exists where the decision maker is indifferent between B and an outcome yielding a combination of A and C with the probabilities of the outcome of A and C summing to 1. Lastly, if A is preferred to B, and C is a risky alternative, the outcomes of A and C are preferred to the outcomes of B and C when the probability of A is equal to the probability of C. When the axioms at satisfied, a utility function exists that relates a single utility index value to risky alternatives which can translate ranking preferences to utility ranking. E-V is consistent with VNM axioms if returns are either normally (elliptically) distributed or if the investor has a quadradic utility function.

The minimum variance model can be expanded to discuss the E-V hedging model. This model is more desirable and is more realistic as it recognizes investors' desire to maximize returns while limiting their variance, or risk, in contrast to the minimum variance hedging models which aimed to achieve maximum risk avoidance. Blank, Carter, and Schmiesing (1991) illustrated that an optimal hedge ratio can be found which incorporates both the desire for risk avoidance and returns from speculation. The mean-variance hedging model includes a risk-minimizing hedge ratio similar to the minimum variance models and was expanded to include a speculative element. These components are commonly referred to as the hedging demand for futures and the speculative demand, respectively. The mean-variance hedging method for finding optimal hedge ratios differs from the minimum variance hedging models in nearly all situations except for complete risk-averse individuals or expected futures price change is zero as in a pure martingale process (Chen, Lee, and Shrestha 2003).

The mean-variance approach is used in a number of studies relating to optimal hedging strategies (Anderson and Danthine 1981; Blank, Carter, and Schmiesing 1991; Cecchetti, Cumby, and Figlewski 1988; Howard and D'Antonio 1984; Chin-Wen, Kuo, and Cheng-Few 1994). Each of these studies incorporate the expected returns as well as the expected variance, or risk, into their models. These models maximize utility under the E-V framework. The following gives a brief overview of their findings and contributions to this area of the literature.

Anderson and Danthine (1981) recognize the application of simple portfolio theory for a hedger using both cash and futures markets due to the lack of a perfect correlation between the two markets. Additionally, the research presented is one of the first which incorporates mean and variance of returns into the study. The study also derives the utility that results from the ability to trade futures for speculation or hedging. The authors also incorporate the concept of risk aversion which is used in subsequent literature.

Howard and D'Antonio (1984) study how risk-free assets can be added to a portfolio to reduce the variance of the total portfolio under the E-V framework. However, as risk-free assets are added, the expected returns from the portfolio decrease. Because adding risk-free assets are both ways to reduce the overall variance of the portfolio, the decision lies in which instrument is more beneficial for enhancing the portfolio's risk-return characteristics (Howard and D'Antonio 1984). An important contribution offered by the article was the introduction of λ as the risk-to-excess-return relative of futures versus spot position. As a result, λ can be used to show numerically show the attractiveness of investing in cash versus futures. Howard and D'Antonio (1984) also show the importance of the relationship between λ and ρ for holding futures. The variable, ρ, or rho, represents the correlation between cash and futures. Previously, the R^2 value was used to measure hedging effectiveness, but the authors show that ρ and λ can be used together more effectively than the previous methods. They found that when λ=ρ, there was no benefit from holding any futures position, but when λ>ρ, the trader would opt to be long futures. While the goal of their work was to find a risk-return measure of hedging effectiveness, the process also created a new method for solving optimal hedge ratios between cash and futures markets.

The work by Cecchetti, Cumby, and Figlewski (1988) is consistent with the literature noting the importance of including expected returns in models rather than assuming that hedgers are pure risk minimizers, but they also add that their review of previous optimal hedging models failed to include an allowance for time variations in cash and future price change distributions. Their study uses 20-year Treasury Notes hedged with T-Bond futures to show how the problem of changing distributions can be solved. Their findings indicated that using an autoregressive conditional heteroskedasticity (ARCH) specification of joint distributions they were able to obtain a series of optimal hedge ratios for an investor with a logarithmic utility curve. The importance of this application is solidified by their empirical results showing the optimal hedge ratio varied between 52% to 91% hedged over the period of interest.

Blank, Carter, and Schmiesing (1991) present an application of portfolio hedging models graphically and mathematically as well as a discussion of limitations and practical application of the existing theory. The authors note that in a perfect hedging scenario, cash and futures are perfectly correlated, implying that the basis remains unchanged; thus, the variance of the completely hedged portfolio is zero. They also echo the arguments made earlier regarding minimum variance modeling: it should be assumed that traders maximize utility rather than minimize risk. The hedge ratio derived maximizes the utility of returns function written as a function of the expected returns, expected variance, and the risk aversion parameter, phi. The resulting optimal hedge ratio includes the risk-minimizing hedge ratio as well as the trader's speculative bias. Kahl (1983) shows that when the trader does not have a market bias, the mean-variance hedge ratio is equal to the minimum-variance hedge ratio.

Although the mean-variance strategies are an improvement over the minimum-variance strategies, they must either have quadradic utility functions or jointly normal returns to be consistent with utility maximization (Chen, Lee, and Shrestha 2003). Some researchers have worked to find alternatives to the restrictions of the mean-variance modeling. These and other hedging models are discussed in a later section.

2.3.3. Mean Value-at-Risk Framework

Value-at-Risk (VaR) has become an increasingly popular measure of risk and has been used to gauge the risk levels of investment strategies by financial institutions, insurance companies, and trading firms, to name a few. VaR is defined as the amount of losses that will not be exceeded within a certain confidence interval over a specified period of time (Vose 2008). To clarify, VaR is not a stress-testing or "worst-case scenario" tool, but rather a tool that can be deployed to give an estimation of expected losses within a statistical confidence interval. The usefulness of VaR is best shown with an example. An investment with a $50 million weekly VaR at 95% confidence would indicate that losses will not exceed $50 million 95% of the time. Alternatively, it could be said that one out of twenty weeks, a loss will occur that will be greater than or equal to $50 million. The use of VaR provides a consistent and relatively simple interpretation of risk within companies, and as a result, may be especially useful for mid and upper-level management teams to quickly understand the risk that various levels of the company are exposed to.

An early application of the VaR methodology at JPMorgan arguably helped grow the popularity of VaR as a measure of risk. They developed RiskMetrics™ as a tool to measure and manage risk internally, and a version based on the models used by JPMorgan was made available to allow market participants and financial institutions to estimate their risk exposure using VaR. The software was capable of evaluating the risks of portfolios to help investors understand the origins of their risk and the total risk within the portfolio. Additionally, firm-level risk measures were available because of the software's capabilities. For additional details on the methodologies and capabilities of the RiskMetrics™ platform see Longerstaey (1996). The creation of the software made implementing VaR metrics within a company easier than ever before, leading to its popularity in industry and research fields.

While the RiskMetrics™ software increased accessibility and use of VaR metrics, it was limited to parametric methods that were less computationally intense than simulation models (Mausser and Rosen 1999). Mausser and Rosen (1999) add to the understanding of VaR by demonstrating simulation-based methods for calculating VaR. Although these methods are less relevant today because of the significant advancements in computing technology, their work helped build the framework for future studies of simulated VaR calculations. Additionally, Mausser and Rosen (1999) decompose portfolio VaR to better understand the assets which contribute significant levels of risk to the portfolio.

One of the first academic studies available on this topic is a comparison between the mean-variance method with VaR modeling for portfolio selection by Alexander and Baptista (2002) who offer insight into how mean-variance and VaR models differ, the economic implications of these differences, the relation between mean-VaR and utility maximization, as well as an application of E-VaR under the condition of nonnormality. Their study shows that as the confidence level of the VaR model is increased to 100%, the mean-VaR efficient set converges to the mean-variance efficient set.

Chang (2011) also applies the value-at-risk approach to identify optimal hedging strategies. The findings indicated that dynamic hedging was better than static hedging strategies. Additionally, minimum variance and mean-variance strategies did not perform as well as the VaR hedging strategy. The article further elongates the trend of authors rejecting variance as a measure of risk and favoring VaR instead. Finally, the empirical results showed that the hedge ratio and hedge performance were affected by the coefficient of risk aversion and confidence level.

Finally, Awudu, Wilson, and Dahl (2016) analyze optimal hedging strategies for an ethanol processor using simulation with and without copula dependence. Three different strategies for hedging were used that included traditional hedging, linear dependence, and copula dependence. Calculations were performed for margins, utility, risk, and hedge ratios for each strategy. The study was unique to other literature in that the model was more complex in the specification of the variables, and it included both direct hedging for the corn and ethanol positions as well as cross-hedging for the products such as corn oil and dried distillers' grains (DDGS). The results showed an optimal hedging strategy using the VaR framework.

2.3.4. Conditional-Value-at-Risk

While VaR has a number of merits, the tool is not without shortfalls. VaR has been criticized for its use of portfolio management and selection because of its lack of subadditivity and convexity (Artzner et al. 1999). This means that in some scenarios the VaR of two portfolios may be greater than the VaR of the sum of returns from the two individual portfolio's. Also, VaR can be difficult to optimize as a result of showing multiple local extrema for discrete distributions (Gao and Liu 2009). Lastly, as mentioned earlier, VaR does not measure tail losses which may be important in certain scenarios. Conditional-value-at-risk, (CVaR) is used in certain applications to combat some of these issues. CVaR, also known as Mean Excess Loss, Tail VaR, or Mean Shortfall, can be applied to a variety of modeling scenarios and can be especially useful for portfolio optimization (Rockafellar and Uryasev 2000).

The approach used by Rockafellar and Uryasev (2000) focuses on portfolio optimization such that the risk of large losses is minimized under the minimum CVaR framework. In the analysis, VaR and CVaR are measured in terms of percentages, and the minimized CVaR approach is compared to the minimum variance approach. The empirical study applied to hedging was based on work done by Mausser and Rosen (1999). Rockafellar and Uryasev (2000) used similar procedures and found that the results of minimizing CVaR were similar to the results found by Mausser and Rosen (1999) which minimized VaR in one-instrument hedges. The authors show that CVaR can be used more readily in multiple-instrument hedging strategies. This application of CVaR could be useful in situations where no futures contract exists for the underlying commodity creating a need for cross hedging.

CVaR minimization was also applied to hedging equity-linked insurance contracts by Melnikov and Smirnov (2012). Partial hedging strategies were utilized as the benefit paid to the insured was related to the performance of the financial market but was unable to be perfectly hedged. Using CVaR, Melnikov and Smirnov (2012) were able to find the optimal partial hedging strategies and also estimate the financial exposure of contracts of a given age.

While Rockafellar and Uryasev (2000) and Melnikov and Smirnov (2012) minimized the CVaR, Nguyen, Nguyen, and Adegbite (2018), Gao and Liu (2009), and Giamouridis and Vrontos (2007) utilize mean-CVaR for portfolio optimization. By applying the mean-CVaR model, Gao and Liu (2009) change the problem of portfolio optimization into a linear problem. Additionally, historical simulations are used rather than an assumption of a normal distribution. Gao and Liu (2009) conclude that mean-CVaR models are "undoubtedly more advantageous" for investing in a variety of assets due to the intuitive model design and calculation superiority. Gao and Liu (2009) and Giamouridis and Vrontos (2007) compare results from mean-variance models and mean-CVaR models. Gao and Liu (2009) used the mean-variance model as a base case for the comparison between static and dynamic portfolio allocations where Giamouridis and Vrontos (2007) compare the overall performance between the two models. Giamouridis and Vrontos (2007) conclude that mean-CVaR outperformed mean-variance in the majority of investment universes with various settings. While CVaR was optimal in many applications, Giamouridis and Vrontos (2007) also found that mean-CVaR requires sophisticated inputs to account for time-variant returns distributions and loses superiority for practical applications as transaction costs increase; therefore, they suggest that investors should not rely solely on CVaR if simple inputs or high transaction costs are used.

2.3.5. Other Hedging Frameworks

Aside from the most commonly used hedging frameworks such as minimum variance, mean-variance, mean-VaR, and mean-CVaR, a variety of other techniques have been used to mitigate some of the issues that have been found with the techniques. One issue that commonly affects mean-variance strategies is the requirement of a quadradic utility function or jointly normal returns for the optimal hedge ratio to maintain consistency with expected utility maximization. Chen, Lee, and Shrestha (2003) present a thorough review of the variety of hedging and portfolio optimization techniques that may be used in situations of non-normal returns and non-quadradic utility. A summary of their conclusions is presented in this section.

To eliminate or relax some of the assumptions surrounding the issues of distributions or utility functions, frameworks such as the mean extended-Gini (MEG) coefficient, semivariance models, and lower partial moments have been developed. The MEG coefficient minimization procedure was used by Cheung, Kwan, and Yip (1990) and Shalit (1995) in the context of hedging. The mean-Gini approach was used primarily for studies relating to income inequality, but Cheung, Kwan, and Yip (1990) suggested it be used as a new framework for evaluating futures and options strategies. One of the benefits of this framework is that the assumption of normal returns is no longer needed and did not require quadradic utility functions. Shalit (1995) sought to compare the mean-variance approach with the MEG approach. The results indicated that when prices followed a normal distribution, the MEG and minimum-variance ratios converge.

Approaches focusing on the variance of the portfolio were beneficial in some cases, focusing purely on the total variance results in ignoring the preferences for upside versus downside risk. Losses are never desirable where upside potential is beneficial. The semivariance approach allows for this and is applied by Chen, Lee, and Shrestha (2001), Turvey and Nayak (2003), Hogan and Warren (1974), and Jong, Roon, and Veld (1997) in a variety of hedging and portfolio situations. Hogan and Warren (1974) introduce the benefits of using semivariance in place of variance and lay a theoretical framework for further development of semivariance in the capital market modeling. Chen, Lee, and Shrestha (2001) use the semivariance approach to determine the optimal hedge ratio and compares the effectiveness of the hedge against other commonly used hedging models. Most recently, Turvey and Nayak (2003) evaluated the semivariance-minimizing hedge ratio in Kansas City wheat and Texas steers. The study compares the minimum semivariance hedge ratio to the minimum variance ratios and finds that the two can be different. Additionally, the empirical results found that the semivariance model suggested lower optimal hedge ratios than the minimum variance ratios.

While many of the most common models used in evaluating and optimizing hedged portfolios have been discussed, other researchers have deployed an array of various methods. Chen, Lee, and Shrestha (2003) published a detailed complication of many of the methods used by researchers. Chen, Lee, and Shrestha (2003) show that some of these methods are based on simple ordinary least squares (OLS) regression analysis such as Ederington (1979) and Benet (1992) while others like Cecchetti, Cumby, and Figlewski (1988) and Baillie and Myers (1991) are based on the conditional heteroscedastic framework. Chen, Lee, and Shrestha (2003) explain that still others use the random coefficient method (Grammatikos and Saunders 1983), the cointegration method (Chou, Denis, and Lee 1996; Ghosh 1993), or the cointegration-heteroscedastic method (Kroner and Sultan 1993). The many frameworks used by researchers over time have led to an improved understanding of portfolio optimization. The theoretical and empirical discussions that follow rely heavily upon the previous work and seek to apply these principles in a new way.

2.4. Background on Optimal Coverage Ratios

The concept of optimizing hedging strategies has been of great interest in the academic community over the last 80 years, as demonstrated by the review of literature presented in the previous sections. While the research has progressed over the years in terms of objective functions, computing methods, and practicality of application to industry, nearly all of the research focused on hedging physical grain positions by taking positions in one or more future markets.

This leaves a significant gap in the literature relating to the risk that is faced by basis and rail traders. By defining a coverage ratio and a profit function for various merchandising scenarios, the utility of the trader can be maximized under various frameworks. Additionally, coverage ratios can serve as a new tool for risk management.

2.5. Summary

This chapter provides the necessary background for approaching the research goal of finding an optimal coverage ratio for a pure basis trader by introducing the concepts of basis, basis trading, mechanics of rail markets, and arbitrage. Furthermore, the chapter highlights the risks associated with commodity trading and follows with a presentation of the evolution of portfolio hedging models over time.

Throughout the review of the literature, there seems to be a vast amount of research that has been conducted related to portfolio theory. Additionally, portfolio hedging models have primarily focused on identifying optimal hedging ratios for firms holding physical grain positions and hedging in a futures market. In contrast, limited research has been conducted on basis trading and rail trading, and no research develops models to determine optimal coverage in these instruments. This study adds value to the existing literature by exploring how elevators or traders can apply portfolio models of hedging to their basis positions and optimize their coverage ratios to maximize profit subject to their risk tolerances.

CHAPTER 3. DETERMINATION OF OPTIMAL COVERAGE RATIOS

3.1. Introduction

A number of portfolio optimization models were reviewed in Chapter 2. These models included minimum-variance, expected utility maximization under mean-variance and mean-semivariance, and mean-VaR. Each of these frameworks has been used in hedge ratio optimization for traders with cash and futures positions. These models have generally involved a trader holding a position in a correlated futures market to mitigate price risk from owning the physical grain. When ownership of the physical grain was transferred, the futures positions were also liquidated. The primary idea was that gains or losses accrued by holding a physical commodity would be offset by the losses or gains accrued by holding opposite futures positions.

This chapter shifts from hedging in futures to taking coverage in a basis or freight market which can be used to manage risk when merchandising grain and managing logistics. The payoff functions for each of the models used in the empirical analysis are introduced. Additionally, the utility-maximizing objective functions are discussed for both mean-variance and mean-semivariance approaches. Lastly, the process for analytically solving the optimal coverage ratios is shown.

3.2. Conceptual Framework

This chapter relies on the concepts introduced in Chapter 2 which included basis and rail markets along with factors that affect each of these markets. Basis markets are strongly influenced by local supply and demand factors such as local processing and usage demand, demand in destination markets, local crop conditions, and local elevator margins (Lorton and White 2010). Changes in basis create trading opportunities that can be realized by making strategic purchases and sales of the commodity.

Three types of basis trading scenarios were introduced in Chapter 2: a back-to-back, a long-the-basis, and a short-the-basis. A back-to-back trade consists of purchasing grain and selling 100% of it to a destination nearly immediately after (or concurrently with) the purchase to secure a margin. Generally, this type of trade leaves small, guaranteed margins for the elevator. A merchant may take a long basis position if the basis market in a destination market is expected to strengthen. This position is created by purchasing grain at the origin market. The elevator owns the grain for a period of time and successfully liquidates the position by selling grain to the destination once the market has increased. Conversely, to take a short basis position, the elevator may make forward sales of grain at the destination for delivery at a later date and purchase the grain prior to delivering it to the destination. The differences in times are denoted throughout this thesis as Time 1 and Time 2. The exact timeframe can be specified using specific entry or exit dates in the empirical models or can be assumed as a six-month period, depending on the distributions used.

Rail freight is traded similarly. Rail cars (or shuttle trains which commonly consist of 110 cars) may be bought in advance or at the time they are needed depending on market expectations and risk tolerances. Whether a shipper enters a short-the-basis or long-the-basis position, it becomes short freight because the grain needs to be shipped at a later time. Purchasing freight in advance at known prices rather than purchasing as needed with risky prices is advantageous because it allows the origin buying basis to be adjusted for the transportation cost to secure the margin from the sale of the grain. If rail freight is at risk, the overall margin is also at risk. Thus, including rail freight in the portfolio decision model is necessary.

While a merchant may desire to take long or short basis positions or have all or no freight coverage, management may want to limit the financial risk exposure of the elevator.

Additionally, management likely does not want to lose market share or handle less grain to limit their risk. Instead, coverage may be taken in the basis or rail markets to allow for handling large quantities while managing the financial risk. Coverage can be reported as a value, such as a quantity of bushels, or as a percentage or ratio. This study determines optimal coverage ratios for various portfolios of basis and rail positions.

Two coverage ratios are developed to allow for long and short basis positions. Equation (3.1) is used in long-the-basis cases.

$$h_D = \frac{amount\ sold\ to\ the\ destination\ at\ T_1}{amount\ owned\ or\ purchased\ at\ origin\ at\ T_1} \tag{3.1}$$

Equation (3.1) shows the coverage taken at the destination for a long basis position where:

h_D = Coverage taken at destination market

T_1 = Time Period 1

For a long-the-basis scenario, Time Period 1 represents the time at which the grain is purchased at the origin. At that time, a long position is established, and the elevator becomes exposed to basis risk. In that same period, a quantity of grain may be sold to the destination, or covered, at a known price to limit risk exposure. For this study, it assumed that the remainder of the grain must be sold by Time Period 2. Variable h_D is reported in a percentage throughout the study.

Equation (3.2) is used in short-the-basis cases.

$$h_O = \frac{amount\ bought\ at\ origin\ at\ T_1}{amount\ presold\ at\ the\ destination\ at\ T_1\ for\ delivery\ at\ T_2} \tag{3.2}$$

Equation (3.2) shows the coverage taken at the origin for a short-the-basis scenario where:

h_O = Coverage taken at origin market

T_1 = Time Period 1

T_2 = Time Period 2

For a short-the-basis scenario, Time Period 1 is the time at which a sale is made to a destination market for delivery at Time Period 2. When the sale is made, the elevator becomes short-the-basis and is at risk that basis moves upward before they purchase grain at Time 2, the period in which delivery to the destination must be made. To manage the risk, some of the grain may be purchased at the origin during Time Period 1. Variable h_O is reported in a percentage throughout the study.

This section showed that basis can be traded by holding long or short basis positions. The risk of these positions can be mitigated by taking coverage at the destination for long positions and taking coverage at the origin for short positions. The coverage ratios presented in this section are used with the following payoff functions.

3.3. Payoff Function Specifications

Payoff functions are used to demonstrate a shipper that purchases grain from farmers or other elevators and makes sales to a destination market. Two sets of models are presented to show long and short basis positions. The first models in each set show a single location with no rail risk; the second shows a single location with basis and transportation risk, and finally, a special case of long-the basis is shown which involves multiple destination locations and transportation risk. The variables used in the equations are presented in Table 3.1. Any variable which exhibits randomness is denoted by an asterisk. Table 3.1 does not distinguish between random and non-random variables.

Table 3.1: Variables and Descriptions

Variable Symbol	Variable	Units	Variable Description
π	Profit	Cents per Bushel	The profit functions vary across the models, but it is always a function of earned income from the sale of grain at a destination location minus the cost of the grain at the origin less freight costs.
h_A or h_B	Grain Coverage Ratio (At Destination A or Destination B)	Percentage	Coverage ratio in the long-the-basis case is defined as the quantity sold to the destination at Time 1 divided by the quantity owned at the origin needing to be sold by Time 2.
h_O	Grain Coverage Ratio (At Origin)	Percentage	Coverage ratio in the short-the-basis cases is defined as the quantity of grain purchased at the origin at Time 1 divided by the quantity of grain sold to the destination requiring delivery by Time 2.
g	Freight Coverage Ratio	Percentage	The freight coverage ratio is the quantity of rail purchased at Time 1 for use during Time 2 divided by the quantity of rail
B_D	Destination Basis	Cents per Bushel	The destination basis is shown for various time periods in the long-the-basis models.
B_O	Origin Basis	Cents per Bushel	Origin basis is shown for various time periods in the short-the-basis models.
T	Rail Tariff Rate	Dollars per Car	Rail tariff rates are assumed to be constant in this study as they are relatively stable over time, especially when compared to daily car value and basis values.
F	Daily Car Value (DCV)	Dollars per Car	Daily car values are determined by auctions

3.4. Long-the-Basis Models

In the long-the-basis models, we assume the elevator has grain in inventory or incoming farmer deliveries that have been bought at a set basis. The elevator has the option to contract the sale of the grain to the destination at the time of the purchase, Time 1, thus securing a known margin or they may wait until a later forced liquidation date, Time 2, due to inventory constraints, cash flow constraints, or other constraints with unknown random prices. While the profit-maximizing solution in the models would be to sell at whichever time period has higher expected prices, a risk-averse and forward-looking trader expects that the future basis and rail prices are random and should contract the sale of a portion of the grain at Time 1 and the remainder of the grain at Time 2. The models that follow build a framework for optimizing the coverage ratios.

3.4.1. Model 1.1: Single Origin, Single Destination

Model 1.1 is the first model in the long-the-basis series. In this model, it is assumed that the elevator has a fixed amount of grain purchased at a given price at Time 1 that must be sold by Time 2. The merchant can either sell all or a portion of the grain at the destination now or can make the sale later. The freight is assumed to be paid by the buyer.

$$\pi = h_A B_{DA1} + (1 - h_A) B^*_{DA2} - B_{O1} - T - F \tag{3.3}$$

In this model, profit is the sum of the percentage of sales made immediately as a back-to-back sale at Time 1 and the later sale of the remaining inventory at a random price, denoted with an asterisk. While it is assumed that freight is paid by the buyer, the destination basis data used in Chapter 4 is reported for grain that is delivered to the destination. To account for this, the tariff rate and daily car values for Time 1 are used to represent the buyer's cost in acquiring freight and are subtracted from the margin that is left for the seller. This could alternatively be shown as:

$$\pi = h_A (B_{DA1} - T - F) + (1 - h_A)(B^*_{DA2} - T - F) - B_{O1} \tag{3.4}$$

where the tariff and freight are first subtracted from the basis destination price. This is shown to clarify that the tariff and freight must be subtracted from a buyer's offer if the buyer is paying the freight, but the results would be the same as in Equation (3.3). Subtracting transportation cost is necessary so that the results of Model 1.1 could easily be compared to the later models where freight is paid by the seller.

3.4.2. Model 2.1: Single Origin, Single Destination with Freight Risk

Model 2.1 is similar to Model 1.1 in the fact that it shows a single destination and a single origin, but now the selling elevator pays the freight, adding freight risk to the equation. Mathematically, Model 2.1 is shown as:

$$\pi = [h_A(B_{DA1}) + (1 - h_A)(B^*_{DA2})] - [gF_1 + (1 - g)F^*_2] - B_{O1} - T \tag{3.5}$$

The first section (in the first set of brackets) of the model shows the revenue generated from immediate, back-to-back sales and from sales made at a later date with random prices, as in Model 1.1 The next section subtracts costs associated with variable shipping costs as a result of changing DCV rates. The cost for shipping is the percentage of freight bought at a known price at Time 1 plus the percentage of freight booked at an unknown random price at Time 2. Lastly, the given basis origin cost and tariff rate are subtracted to give profit.

3.4.3. Model 3.1: Single Origin, Multiple Destination with Freight Risk

Model 3.1 continues with a fixed origin price and transportation price risk but is further expanded from Model 2.1 to include multiple destinations and freight risk. The mathematical representation differs. Tariff rates differ across destinations while DCV does not. Thus, variable F, may be at risk and is not location specific where tariff rate T is known and is specific to each location. As a result, the tariff rate must be subtracted from the destination basis to determine the optimal selling location.

$$\pi = \left\{ \begin{array}{c} h_A(B_{DA1} - T_A) + h_B(B_{DB1} - T_B) + \\ (1 - h_A - h_B)MAX[(B^*_{DA2} - T_A), (B^*_{DB2} - T_B)] \\ - [gF_1 + (1 - g)F^*_2] - B_{O1} \end{array} \right\} \tag{3.6}$$

The equation shows that profit is the sum of revenues generated from coverage taken at Destination A, coverage taken at Destination B, and the maximum of the revenue earned from selling to Destination A or B during Time 2 with random prices less the cost of rail freight bought for a known price at Time 1 and at a random cost at Time 2 less the given origin basis. The logic behind this model is that the merchandiser covers a portion of the inventory at Time 1 at one or both of the destinations depending on the prices, tariff rates, and variance of each market and will sell the remaining inventory to market with the highest delivered price at Time 2.

3.5. Short-the-Basis Models

An alternative case can be presented in which the merchant has made forward sales to a destination but has not yet purchased grain at the origin. This scenario may occur when the merchant is expecting basis to decrease. The merchant is likely to sell all of their remaining inventory but may want to further expand on the merchandising opportunity. This may be accomplished by entering a short basis position in which grain is sold at the destination for delivery at a later time before all the grain is purchased at the origin. The risk-averse merchant must then decide what portion of the short sale may be left open with the hope of lower origin basis levels before the grain must be delivered to the seller. The models that follow provide the means for this type of analysis.

3.5.1. Model 1.1.1: Single Origin, Single Destination

As in the long-the-basis case, it is most intuitive to first introduce a case in which there is no freight risk. The merchant fixes the sales price at the destination for Time 2, the later delivery period, and has to decide what portion of the sale will be covered immediately versus the amount that will be left open with the expectation of decreasing prices.

$$\pi = B_{D2} - h_O(B_{O1}) - (1 - h_O)(B_{O2}^*) - T - F \qquad (3.7)$$

Profit is defined as the fixed basis received from the sale of grain during Time 2 minus the percent of grain purchased at the origin for a fixed price at Time 1 minus the percent of grain purchased for an unknown random basis at Time 2. Similar to Model 1.1, transportation costs are also subtracted to make the profit values comparable to the later models. Unlike the long-the-basis case where the randomness of the destination affects revenues, in the short-the-basis models, revenues are exogenous and held constant while randomness is introduced with regards to the cost of acquiring grain.

3.5.2. Model 2.1.1: Single Origin, Single Destination with Freight Risk

A case with one destination with a fixed price and one origin with risk was shown without rail risk in Model 1.1.1; the model below is now expanded to include rail risk as was done in Model 2.1.

$$\pi = B_{D2} - h_O(B_{O1}) - (1 - h_O)(B^*_{O2}) - [gF_1 + (1-g)F^*_2] - T \tag{3.8}$$

The profit function is mostly unchanged except for the inclusion of rail risk. Profit is represented as the margin made from the sale and purchase(s) of grain, minus the total cost of freight which includes the freight bought at Time 1 for a known price, the freight bought for an unknown price at Time 2, as well as the fixed tariff rate.

Each of the models has characteristics that allow for various scenarios to be analyzed and further understood. Long-the-basis models were presented first, starting with the margin at risk due to changing destination prices and expanding into a portfolio including rail and finally multiple destinations. The short-the-basis models allow for one origin and one destination and can include rail as well.

3.6. Analytical Derivation of Optimal Coverage Ratios

The models presented in the previous sections show the payoffs for various trading scenarios. For long basis positions, profits increase if the basis appreciates over time and if the trader has grain left to sell at Time 2. In the short basis models, profits increase if the basis declines between Time 1 and Time 2 and if the trader still has grain to purchase at Time 2. In any of the models involving rail, the total profit decreases if rail prices increase between Time 1 and Time 2 if the trader did not book freight at Time 1. With the profit functions defined, two analytical models can be developed to derive optimal coverage ratios for grain and rail.

3.6.1. Analytical Methods

The methods for analytically deriving the optimal coverage ratios follow Kahl (1983) and Blank, Carter, and Schmiesing (1991). Each of the studies defines a profit or payoff function, a variance of profit function, and a utility function which includes the profit, variance of profit, and a risk aversion parameter. The expected utility function is maximized by differentiating with respect to each variable being optimized. The setup of the analytical models in Blank, Carter, and Schmiesing (1991) differ from those used in Kahl (1983) in that Blank, Carter, and Schmiesing (1991) defines a variable h for the "proportion of inventory which is hedged or the hedging ratio" where Kahl (1983) has separate variables for the quantity of cash and future positions. A summary of the approach used by Blank, Carter, and Schmiesing (1991) is given below. They begin by defining the change in the value of grain held in inventory as:

$$\tilde{V} = (\widetilde{p_2} - p_1) + h(f_1 - \widetilde{f_2}) \tag{3.9}$$

where the random variables are denoted with tilde and variables are defined as:

$\tilde{V}$	=	Change in per unit value of inventory
h	=	Hedging ratio or portion of inventory being hedged
p_1	=	Known cash price at Time 1
$\widetilde{p_2}$	=	Random and unknown cash price at Time 2
f_1	=	Known futures price at Time 2
$\widetilde{f_2}$	=	Random and unknown futures price at Time 2

The payoff function assumes that inventory is fixed leaving the changes in cash prices to affect the entire inventory value. However, a portion of the grain may be hedged by taking an opposite position in the futures market which may be liquidated when the cash position is liquidated.

Next, the expected utility function is defined which used to determine the optimal h as

$$U(R) = E(R) + \Phi var(R) \tag{3.10}$$

where the utility of the hedger is represented by $U(R)$. The utility is a function of the expected returns, $E(R)$, plus an adjustment for risk tolerance, Φ, and the variance of the returns, $var(R)$. In this case, the risk aversion parameter is negative, so as risk aversion increases or the variance of returns increases, the total utility decreases.

The variance of returns of the portfolio, $var(R)$ is derived from the random variables and their covariances. The variance is represented as

$$var(V) = \sigma_{p_2}^2 + h^2\sigma_{f_2}^2 - 2h\sigma_{p_2 f_2} \tag{3.11}$$

where $\sigma_{p_2}^2$ is the variance of the cash price at Time 2, $\sigma_{f_2}^2$ is the variance of the futures price at Time 2, and $\sigma_{p_2 f_2}$ represents the covariance between cash and futures at Time 2.

The payoff function in Equation (3.9) and the variance function in Equation (3.11) can be substituted into Equation (3.10) to give the following expected utility function:

$$U(V) = E(\widetilde{p_2}) - p_1 + h(f_1 - E(\widetilde{f_2})) + \Phi\left[\sigma_{p_2}^2 + h^2\sigma_{f_2}^2 - 2h\sigma_{p_2 f_2}\right] \tag{3.12}$$

After making the substitution, Blank, Carter, and Schmiesing (1991) maximize the utility function with respect to h which yields:

$$\frac{\partial U(V)}{\partial h} = f_1 - E\left(\widetilde{f_2}\right) + \Phi 2\sigma_{f_2}^2 - \Phi 2\sigma_{p_2 f_2} \tag{3.13}$$

The optimal hedge ratio, h^*, is found by solving for h

$$h^* = \frac{\sigma_{p_2 f_2}}{\sigma_{f_2}^2} - \frac{f_1 - E(\widetilde{f_2})}{\Phi 2\sigma_{f_2}^2} \tag{3.14}$$

The optimal hedge as defined by Blank, Carter, and Schmiesing (1991) has two components. The first part of the equation is equal to a minimum variance model and can be interpreted as the risk-minimizing hedge, or the hedging demand for futures. The second part of the equation makes up the speculative demand which changes based on the trader bias and the risk aversion level.

The process used by Blank, Carter, and Schmiesing (1991) was applied to determine optimal coverage ratios in grain and rail portfolios. Sections 3.6.2 and 3.6.3 show analytical solutions for Model 1.1 and Model 2.1, respectively. A general mean-variance case is presented first with a general mean-variance utility maximization. At the end of each derivation, the optimal coverage ratio is derived for a CARA utility function which is used throughout the empirical analysis.

3.6.2. Model 1.1 Analytical Solution

Using the same process as Blank, Carter, and Schmiesing (1991), the optimal coverage ratio can be determined for Model 1.1. In this case, there is only one variable being optimized, h_A, which represents the portion of grain sold to the destination versus what is owned at the origin. The payoff function and variance of the payoff function are shown along with the substitution into the utility function. Lastly, the maximization of the expected utility is completed which allows an optimal h_A to be determined.

The profit function for Model 1.1 is defined as:

$$\pi = h_A B_{DA1} + (1 - h_A) B_{DA2}^* - B_{O1} - T - F \tag{3.15}$$

The random variables are denoted with an asterisk. Table 3.1 describes the variables. The function can be transformed to show the expected profit, $E(\pi)$.

$$E(\pi) = h_A B_{DA1} + (1 - h_A)E(B^*_{DA2}) - B_{O1} - T - F \tag{3.16}$$

The variance of the profit function is defined as:

$$Var(\pi) = (1 - h_A)^2 \sigma^2_{B_{D2}} \tag{3.17}$$

where $\sigma^2_{B_{D1}}$ is the variance of the random variable, B^*_{DA2}, and $(1 - h_A)$ represents the quantity of grain that must be sold during Time 2.

Next, a generalized case of expected utility is presented. The generalized utility function follows Kahl (1983) in that the risk aversion parameter is positive. Because of the positive risk aversion, the variance and risk aversion are subtracted from the expected profit, rather than added. This notation is used as it offers a more intuitive application to CARA utility which is used in the empirical analysis.

$$E(U(\pi)) = E(\pi) - \lambda var(\pi) \tag{3.18}$$

The expected profit is $E(\pi)$, the variance of profit is $var(\pi)$, and the risk aversion is captured by λ. The risk aversion coefficient is positive so that for low levels of risk aversion, utility is less dependent on the variance of profits than at high risk aversion levels. Substituting Equation (3.16) and (3.17) into Equation (3.18) gives

$$E(U(\pi)) = [h_A B_{DA1} + (1 - h_A)E(B^*_{DA2}) - B_{O1} - T - F] - \lambda(1 - h_A)^2 \sigma^2_{B_{D2}} \tag{3.19}$$

To determine the optimal h_A, the utility function is differentiated with respect to h_A.

$$\frac{\partial E(U(\pi))}{\partial h_A} = B_{DA1} - E(B^*_{DA2}) + 2\lambda(1 - h_A)\sigma^2_{B_{D2}} \tag{3.20}$$

Setting Equation (3.20) equal to zero and solving for the optimal h_A, h^*_A, yields

$$h^*_A = \frac{B_{DA1} - E(B^*_{DA2})}{2\lambda\sigma^2_{B_{DA2}}} + 1 \tag{3.21}$$

By using the second-order sufficient conditions for a maximum, it can be shown that h_A^* is a local maximum of the utility function.

The CARA utility function used in the empirical analysis is defined as

$$E(U(\pi)) = E(\pi) - 0.5\Phi var(\pi) \tag{3.22}$$

In the generalized case, λ represents any weighting on the variance from risk aversion. Setting 0.5Φ is equal to λ produces the generalized utility function presented in (3.18). Substituting 0.5Φ for λ transforms the generalized coverage ratio to the optimal coverage under CARA utility.

$$h_A^* = \frac{B_{DA1} - E(B_{DA2}^*)}{\Phi\sigma_{B_{DA2}}^2} + 1 \tag{3.23}$$

The optimal coverage ratio is the difference between the known destination basis at Time 1 and the expected destination basis at Time 2 divided by the variance and the risk aversion parameter, plus one. Depending on the inputs for each of the variables, h_A^* can exhibit a relatively large range of coverage ratio possibilities which are not always feasible. Thus, bounds are required such that $0 \leq h_A^* \leq 1$. If the calculated h_A^* exceeds those bounds, the optimal solution should be considered either zero or one.

The optimal coverage ratio is made up of the numerator, the trader's bias in the basis market, and the denominator which is the impact of risk and risk aversion. When the trader expects a basis increase from Time 1 to Time 2, the numerator is negative. For larger levels of expected basis increases, the more negative the numerator becomes which leads to lower optimal coverage ratios. When the trader expects no basis movement, the numerator becomes zero, resulting in an optimal coverage ratio of 100%. If the trader is expecting basis to decrease, the numerator is positive, creating an optimal coverage ratio greater than one. If a trader is expecting basis decrease, the decision is best described using the short-the-basis models. Risk aversion and variance are accounted for in the denominator. As risk aversion increases and the variance of the destination basis increases, the denominator becomes larger which creates larger optimal coverage ratios.

3.6.3. Model 2.1 Analytical Solution

The previous section showed an analytical derivation of a single variable optimization. This section uses a similar process to derive the optimal h_A and g, the optimal coverage ratio for grain and rail.

The profit function for Model 2.1 is defined as:

$$\pi = [h_A(B_{D1}) + (1-h_A)(B^*_{D2})] - [gF_1 + (1-g)F^*_2] - B_{O1} - T \tag{3.24}$$

The random variables are denoted with an asterisk. Table 3.1 describes the variables. Expected profit, $E(\pi)$ is shown with the following transformation.

$$E(\pi) = [h_A(B_{D1}) + (1-h_A)E(B^*_{D2})] - [gF_1 + (1-g)E(F^*_2)] - B_{O1} - T \tag{3.25}$$

The variance of the profit function is defined as:

$$Var(\pi) = (1-h_A)^2\sigma^2_{B_{D2}} + (1-g)^2\sigma^2_{F_2} + 2(1-h_A)(1-g)\sigma_{B_{D2},F_2} \tag{3.26}$$

where

$\sigma^2_{B_{D2}}$ = the variance of the random variable, B^*_{D2}

$\sigma^2_{F_2}$ = the variance of the random variable, F^*_2

σ_{B_{D2},F_2} = the covariance between random variables B^*_{D2} and F^*_2

The generalized utility function is defined as

$$E(U(\pi)) = E(\pi) - \lambda var(\pi) \tag{3.27}$$

Substituting Equations (3.25) and (3.26) into Equation (3.27) yields:

$$E\big(U(\pi)\big) = \begin{Bmatrix} [h_A(B_{D1}) + (1-h_A)E(B^*_{D2}) - gF_1 - (1-g)E(F^*_2) - B_{O1} - T] \\ -\lambda\big[(1-h_A)^2\sigma^2_{B_{D2}} + (1-g)^2\sigma^2_{F_2} + 2(1-h_A)(1-g)\sigma_{B_{D2},F_2}\big] \end{Bmatrix} \tag{3.28}$$

To determine the optimal h_A and g, the utility function must be differentiated with respect to h and g.

$$\frac{\partial E(U(\pi))}{\partial h_A} = B_{D1} - E(B_{D2}^*) + 2\lambda(1 - h_A)\sigma_{B_{D2}}^2 + 2\lambda(1 - g)\sigma_{B_{D2},F_2} \tag{3.29}$$

$$\frac{\partial E(U(\pi))}{\partial g} = -F_1 - E(F_2^*) + 2\lambda(1 - \mathrm{g})\sigma_{F_2}^2 + 2\lambda(1 - h_A)\sigma_{B_{D2},F_2} \tag{3.30}$$

This yields a system of linear equations that can be solved using elimination or other methods.

The resulting optimal coverage ratio for grain is

$$h_A^* = \left\{ \begin{array}{c} \frac{E(B_{D2}^*) - \left(B_{D1}\sigma_{F_2}^2\right)}{2\lambda} - \frac{F_1 + E(F_2^*)\sigma_{B_{D2},F_2}}{2\lambda} + \\ \frac{\sigma_{F_2}^2\left(\sigma_{B_{D2},F_2} - \sigma_{B_{D2}}^2\right) + \sigma_{B_{D2},F_2}(\sigma_{B_{D2},F_2} - \sigma_{F_2}^2)}{{\sigma_{B_{D2},F_2}}^2 - \sigma_{B_{D2}}^2\sigma_{F_2}^2} \\ \\ \end{array} \right\} \tag{3.31}$$

The resulting optimal coverage ratio for rail is

$$g^* = \left\{ \begin{array}{c} \frac{E(B_{D2}^*) - \left(B_{D1}\sigma_{B_{D2},F_2}\right)}{2\lambda} - \frac{F_1 + E(F_2^*)\sigma_{B_{D2}}^2}{2\lambda} + \\ \frac{\sigma_{B_{D2}}^2\left(\sigma_{F_2}^2 - \sigma_{B_{D2},F_2}\right) + \sigma_{B_{D2},F_2}(\sigma_{B_{D2}}^2 - \sigma_{B_{D2},F_2})}{\sigma_{F_2}^2\sigma_{B_{D2}}^2 - {\sigma_{B_{D2},F_2}}^2} \\ \\ \end{array} \right\} \tag{3.32}$$

By using the second-order sufficient conditions for a maximum, it can be shown that h_A^* is a local maximum of the utility function.

The CARA utility function used in the empirical analysis is defined as

$$E(U(\pi)) = E(\pi) - 0.5\Phi var(\pi) \tag{3.33}$$

In the generalized case, λ represents any weighting on the variance from risk aversion. Setting 0.5Φ is equal to λ produces the generalized utility function presented in (3.18). Substituting 0.5Φ for λ transforms the generalized coverage ratio to the optimal coverage under CARA utility.

The resulting optimal coverage ratio for grain is

$$h_A^* = \left\{ \begin{array}{c} \dfrac{E(B_{D2}^*) - \left(B_{D1}\sigma_{F_2}^2\right)}{\Phi} - \dfrac{F_1 + E(F_2^*)\sigma_{B_{D2},F_2}}{\Phi} + \\ \dfrac{\sigma_{F_2}^2\left(\sigma_{B_{D2},F_2} - \sigma_{B_{D2}}^2\right) + \sigma_{B_{D2},F_2}\left(\sigma_{B_{D2},F_2} - \sigma_{F_2}^2\right)}{\sigma_{B_{D2},F_2}{}^2 - \sigma_{B_{D2}}^2\sigma_{F_2}^2} \end{array} \right\} \tag{3.34}$$

The resulting optimal coverage ratio for rail is

$$g^* = \left\{ \begin{array}{c} \dfrac{E(B_{D2}^*) - \left(B_{D1}\sigma_{B_{D2},F_2}\right)}{\Phi} - \dfrac{F_1 + E(F_2^*)\sigma_{B_{D2}}^2}{\Phi} + \\ \dfrac{\sigma_{B_{D2}}^2\left(\sigma_{F_2}^2 - \sigma_{B_{D2},F_2}\right) + \sigma_{B_{D2},F_2}\left(\sigma_{B_{D2}}^2 - \sigma_{B_{D2},F_2}\right)}{\sigma_{F_2}^2\sigma_{B_{D2}}^2 - \sigma_{B_{D2},F_2}{}^2} \end{array} \right\} \tag{3.35}$$

These analytical solutions show that whether finding an optimal coverage ratio for grain or rail, similar inputs are required. This is because there is a relationship between grain and rail prices that requires the covariance term. A trader may expect basis or DCV to increase, decrease, or remain the same. The optimal coverage ratios vary based on the amount of expected market movement, the risk aversion level of the trader, and the correlation between the markets. If the trader is expecting an increase in basis, the optimal coverage ratio is expected to be lower than 100% and approach 0% as risk aversion decreases and the amount of upside in basis increases. If basis is expected to decrease, the optimal coverage ratio for a long position is expected to be 100%. If the basis is expected to remain the same, the optimal coverage ratio likely is between 0% and 100%, depending on the risk aversion level and the standard deviations of the markets. Regardless of the bias, as risk aversion increases, the optimal coverage ratio approaches 100%.

Both of the optimal coverage ratios include the trader market bias for both destination basis and daily car values as well as a risk minimizing portion. For example, for h_A^*, the first row shows the trader's bias for the basis market by subtracting the expected Time 2 basis from the known Time 1 basis, weighted for risk aversion. Next, the trader's bias for the rail market is introduced, also weighted for risk aversion. Lastly, the risk-minimizing solution is shown which involves the interaction between the variance of basis and daily covariances along with their covariances.

These equations may be interpreted similarly to Equation (3.14) which shows the hedging and speculative demand. In this case, there are no futures to be optimized, but positions are taken in related instruments, in this case, B_{D2}^* and F_2^*. In Equations (3.34) and (3.35), the first term captures the bias. In this sense, the bias is similar to the speculative demand for the instrument. If there is bias in the basis or freight market, it impacts the speculative demand for forward coverage. The bottom two rows of Equations (3.34) and (3.35) capture the demand for risk mitigation. Similar to hedging demand, this is derived using the standard deviations of each market and the intermarket correlations.

3.7. Summary

This chapter transitions from the accepted understanding of hedging ratios by defining coverage ratios for grain and rail. The optimal coverage ratio for long basis position describes what portion of grain is sold to the destination at Time 1 relative to the ownership of grain at Time 1. The coverage ratio for the short basis position describes what portion of grain is bought at the origin at Time 1 relative to the amount that would need to be delivered to the destination at Time 2. The rail coverage ratio describes the portion of rail purchased at Time 1 relative to the quantity of rail capacity required at Time 2.

The payoff functions were introduced for two sets of models for long and short basis positions along with a summary of the variables that are used throughout the analysis. The long-the-basis models have a defined buying price with an option to sell grain to the destination at

Time 1 or Time 2 which leaves revenue at risk of changing destination prices. In contrast, the short-the-basis models have a defined selling price at the destination and may purchase the grain at Time 1 or Time 2 which leaves the cost at risk of changing origin prices. A rail position is added to each set of models in which the elevator is short rail and must buy rail at Time 1 or at Time 2. Managing risk due to basis or rail is significant for overall margin management.

Lastly, a summary of how Blank, Carter, and Schmiesing (1991) derive an optimal hedge ratio is presented. Following their process, an analytical solution is derived for Model 1.1 and Model 2.1. The analytical solution for Model 1.1 shows that as the basis is expected to increase between Time 1 and Time 2, the optimal coverage ratio decreases. Additionally, the bias of the trader impacts the coverage ratios. If basis is expected to increase dramatically, the optimal coverage ratio should be lower than if the basis is only expected to increase slightly. Furthermore, the optimal coverage ratio increases as risk aversion and the variance of the destination basis increases.

CHAPTER 4. EMPIRICAL PROCEDURES

4.1. Introduction

Throughout this thesis, the logic and relevance for optimizing basis coverage ratios have been developed. Literature relating to this topic is focused primarily on hedge ratios that involve a trader holding a physical commodity position, taking an opposite position in a related futures market, and liquidating both positions at a later time in order to maximize a given objective function. Unlike futures hedging problems, basis coverage problems involve deciding what portion of a purchase or sale of grain should be made immediately versus at a deferred date under various levels of risk tolerance. The previous chapter presented a theoretical framework for optimizing coverage ratios using mean-variance. Using previous literature as a guide and the theory shown in the last chapter, we now transition to the empirical models.

This chapter presents two sets of models which are designed to evaluate common merchandising scenarios. While the models in each series are similar, they differ in that the long-the-basis series assumes the merchant has already purchased grain at the origin and now must decide how much to sell to the destination while the short-the-basis series assumes the merchant has sold grain to the destination and must decide how much coverage should be taken at the origin. Generally, it is assumed that when a merchandiser is expecting basis levels to increase they would hold a long basis position, and if they are expecting a decrease in the basis, they hold a short basis position. The first models in each series focus only on the grain portion and are later expanded to include freight and multiple locations.

The models are evaluated under the mean-variance and semivariance frameworks using historical BestFit™ distributions and time series forecasted distributions. The objective functions are discussed in detail followed by the data and historical BestFit™ and time series forecasted distributions. Lastly, the simulation procedures are presented.

4.2. Model Assumptions and Specifications

Chapter 3 introduced and defined the coverage ratios that are used throughout the empirical analysis. Separate definitions of coverage ratios are required for long and short basis positions. In a long-the-basis case, coverage is taken at the destination and is defined as the amount of grain sold to the destination at Time 1 divided by the amount of grain owned at the origin. In the short-the-basis case, coverage is taken at the origin and is defined as the amount of grain purchased at the origin at Time 1 divided by the amount of grain presold to the destination for delivery by Time 2. The coverage is taken in the rail market by making forward purchases. The rail coverage ratio is defined as the quantity of rail purchased at Time 1 divided by the amount of rail capacity required at Time 2.

The models used in the empirical analysis were discussed in Chapter 3. The models are specified to illustrate the applications for coverage ratios. Two series of models are presented: one for long basis positions and another for short basis positions. Within the series, the models increase in complexity from single market cases to multi-market optimizations.

The goal of the models that follow is to determine optimal coverage ratios for base case scenarios and to demonstrate how various levels of risk aversion, standard deviations of prices, and the relationships among the individual parts of the portfolio can have substantial effects on optimal coverage ratios. To demonstrate these concepts, some assumptions are required.

4.2.1. Model Assumptions

First, it assumed that the elevator acts as a pure basis trader. This means that the elevator is not exposed to futures market price risk. To allow for a pure basis trading portfolio, it is assumed that the elevator is a naïve hedger, holding an equal and opposite position in the futures market as in the cash market to eliminate futures market price risk. Additionally, it is further assumed that the hedges are placed in the month that the merchandiser plans to liquidate the cash position so that there is not exposure to futures market spread risk.

In reality, an elevator likely holds many types of grains with varying inventories, farmer deliveries, and forward selling opportunities. To find optimal coverage ratios and demonstrate various relationships that exist within the portfolio selection problem, the elevator's position is simplified to a quantity of grain sold relative to a given inventory or quantity of grain purchased relative to a given quantity of forward sales. If the problem was to be expanded, correlation parameters between the additional commodity's markets would be required.

Assumptions are also made regarding the timing of sales and deliveries of grain. Within this study, purchases or sales may be made in one period with the delivery occurring during another. This is a relatively bold assumption as basis is always tied to a specific delivery period. The assumption of taking coverage at Time 1 for delivery at Time 2 in the grain market allows the timing on the purchases of freight to differ from the sales of grain. Additionally, the futures and basis market may at times offer incentives to store, known as contango or a carry in the market. In this study, it is assumed that there is no cost for the elevator to carry grain into later delivery periods and that the market is not incentivizing storage.

Furthermore, it assumed that the distributions do not change. When a distribution is derived for Time 1, it is expected that the distribution may be carried forward to a later time period, Time 2. Time 2 may occur anywhere between a day in advance to multiple years in advance.

Tariff rates tend to be relatively stable, so tariff is assumed to be constant throughout the study. Freight risk is captured by changes in the daily car value (DCV) which can have high volatility across a large range of prices and is correlated to basis as shown below.

4.2.2. Long-the-Basis Models

The long-the-basis models assume that an elevator in Jamestown, ND has ownership of soybeans either in storage or through forward contracts at Time 1 and the grain must be sold to an exporter in the PNW by Time 2. The goal of models in this series is to determine the optimal quantity of grain to sell at Time 1. The models increase in complexity with the first assuming no rail risk, the second including rail risk, and the third including an alternative destination market. In the long-the-basis cases, margin risk results from changing basis levels at the destination market as well as changes in rail freight when applicable. Random variables are denoted by an asterisk.

4.2.2.1. Single Origin, Single Destination: (Model 1.1)

Model 1.1 assumes a merchant has ownership of a quantity of grain in Jamestown, ND and must decide how much of the grain should be covered now or at a forced liquidation date, Time 2. It is assumed that the freight is paid by the buyer.

$$\pi = h_A B_{DA1} + (1 - h_A) B^*_{DA2} - B_{O1} - T - F \qquad (4.1)$$

Profit is defined as the portion of revenues generated from sales at the PNW during Time 1 at known prices and the portion of revenues generated from random prices at Time 2 minus the cost of acquiring grain at Jamestown minus the tariff and DCV. Because the destination prices are reported as grain delivered to the PNW, tariff and DVC are subtracted to represent the buyer's cost in acquiring freight which is passed on to the seller of the grain. This is discussed in detail in Chapter 3.

4.2.2.2. Single Origin, Single Destination with Freight Risk: (Model 2.1)

Model 2.1 expands upon Model 1.1 by adding rail risk into the equation.

$$\pi = [h_A(B_{DA1}) + (1 - h_A)(B^*_{DA2})] - [gF_1 + (1 - g)F^*_2] - B_{O1} - T \tag{4.2}$$

The first part of the equation represents the revenues earned from the sale of soybeans to the PNW while the second part subtracts the cost of rail cars booked at Time 1 at known prices and the cost of cars booked at Time 2 at random prices. The cost of acquiring the soybeans in Jamestown and the tariff rate which are assumed constant are also subtracted.

4.2.2.3. Single Origin, Multiple Destination with Freight Risk: (Model 3.1)

The last long-the-basis model illustrates a merchant who may sell to multiple destinations at Time 1 and Time 2. Variable h_A represents the portion of grain covered at the PNW at Time 1 while variable h_B represents the portion of grain covered at St. Louis at Time 1.

$$\pi = \left\{ \begin{array}{c} h_A(B_{DA1} - T_A) + h_B(B_{DB1} - T_B) + \\ (1 - h_A - h_B)MAX[(B^*_{DA2} - T_A), (B^*_{DB2} - T_B)] \\ - [gF_1 + (1 - g)F^*_2] - B_{O1} \end{array} \right\} \tag{4.3}$$

Rail tariffs differ across locations, so it is subtracted from the destination basis. Revenue is a function of sales made to the PNW and St. Louis at Time 1 at known prices plus the sales made at Time 2 at random prices plus the sale of the remaining grain at Time 2 to the destination with the highest delivered price. The cost of acquiring rain and rail is subtracted similarly to Model 2.1.

4.2.3. Short-the-Basis Models

In the short-the-basis models, it is assumed that at Time 1 an elevator in Jamestown, ND makes forward sales of soybeans to a destination in the PNW for delivery by Time 2. The elevator may purchase the soybeans during Time 1 to secure a margin or wait until Time 2 with hopes of lower basis levels. The goal of these models is to determine the optimal coverage levels that should be taken at the origin. In this case, the margin risk is due to changing costs of acquiring the grain at the origin as well as changing rail prices when applicable. The first model only includes origin basis risk, and the second includes rail risk. Random variables are signified with an asterisk.

4.2.3.1. Single Origin, Single Destination: (Model 1.1.1)

Model 1.1.1 is similar to Model 1.1 in that freight is assumed to be paid by the buyer which eliminates the risk of price changes in the freight market.

$$\pi = B_{DA2} - h_O(B_{O1}) - (1 - h_O)(B_{O2}^*) - T - F \tag{4.4}$$

Profit is defined as the revenue from a sale of grain at Time 1 to the PNW delivered by Time 2 minus the cost of acquiring a portion of the grain in Jamestown at a known price minus the cost of acquiring the remaining grain in Jamestown at Time 2 less tariff and DCV.

4.2.3.2. Single Origin, Single Destination with Freight Risk: (Model 2.1.1)

Model 2.1.1 introduces rail risk similarly to Model 2.1 The revenues and costs related to the grain transaction remain unchanged.

$$\pi = B_{DA2} - h_O(B_{O1}) - (1 - h_O)(B_{O2}^*) - [gF_1 + (1 - g)F_2^*] - T \tag{4.5}$$

The profit function now is defined as the revenues earned from the sale of grain at the PNW at known basis levels minus the cost of purchasing grain in Jamestown at known levels at Time 1 and at unknown levels at Time 2. The costs of freight are subtracted which is made up of the cost of freight at Time 1 and the unknown cost of freight at Time 2.

4.3. Objective Function Specifications

To determine optimal coverage ratios, an objective function must be defined which corresponds to the decision maker’s utility function. The optimal coverage ratios were determined analytically in Chapter 3 using a generalized utility function with the mean-variance framework. This section presents two objective functions that are used to find empirical solutions based on mean-variance and mean-semivariance.

4.3.1. Mean-Variance

The application of mean-variance modeling for hedge ratios was demonstrated by Blank, Carter, and Schmiesing (1991) and Kahl (1983), among others. The E-V framework offers advantages in risk modeling over the minimum-variance framework as it allows for returns to be maximized while minimizing the variance of the returns – a more realistic approach than simply minimizing variance. Assuming a constant absolute risk aversion (CARA) utility function, the expected utility for the decision-maker can be shown as:

$$E[U(\pi)] = CE[U(\pi)] = E(\pi) - 0.5\Phi var(\pi) \tag{4.6}$$

where $E(\pi)$ is the expected return, Φ is the CARA risk aversion coefficient, and $var(\pi)$ is the variance of returns.

The mean-variance framework assumes that the decision-maker views upside and downside risk equally. This framework is common in the literature, likely because of the ease of interpretation, but it does not necessarily reflect the thinking of most traders or managers.

4.3.2. Mean-Semivariance

Under mean-variance, the coverage ratio would be the same for the right or left-skewed distributions given the same mean and variance of the distribution. When depicted graphically as shown in Figure 4.1 however, it is apparent that a right-skewed distribution for a profit would be preferable to the risk associated with a left-skew (“Left Skewed vs. Right Skewed Distributions” 2021).

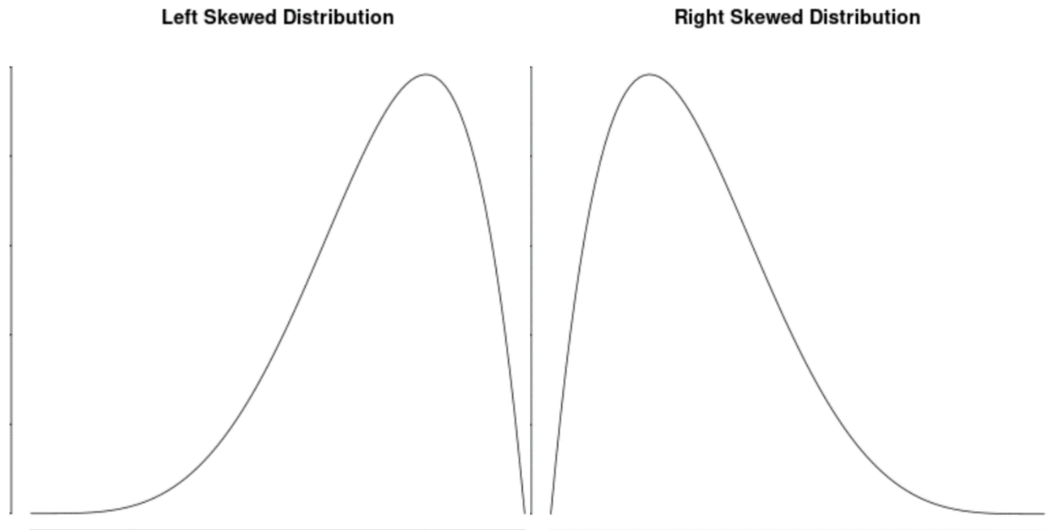

Figure 4.1: Left vs Right Skewed Distributions

While the mean-semivariance framework seems to be less prevalent in the literature, some advantages exist when evaluated using skewed distributions of data. The mean-semivariance method only tries to avoid downside deviations from profit, rather than the mean-variance which tries to minimize the variance of the profit. The CARA utility function is still assumed and semivariance of returns is substituted for the variance of returns such that:

$$E[U(\pi)] = CE[U(\pi)] = E(\pi) - 0.5\Phi \text{semi}var(\pi) \tag{4.7}$$

where $E(\pi)$ is the expected return, Φ is the CARA risk aversion coefficient, and $semivar(\pi)$ is the semivariance of returns.

4.4. Data

The data used in this study serves to illustrate the effectiveness and interpretation of the models. While specific locations and commodities were chosen, the purpose of the data is to serve as an example of conditions that a merchant may face, not necessarily to be representative of a general time period or location.

4.4.1. Data Sources

This study used weekly data collected for April 2018 through February of 2021 and was analyzed over various time periods. Base case scenarios evaluated using historical BestFit™ and time series distribution used data from July 24, 2020 to January 29, 2021 which allowed recent price history to be shown without structural breaks in the data. Data was collected primarily from TradeWest, with supplemental data from Thomson Reuters Eikon. Even with supplemental data, small gaps remained which were filled using linear interpolation.

The sources of data collected are shown in Table 4.1 below.

Table 4.1: Data Sources

Data	Source	Data Description
Basis Data	TradeWest Brokerage, 2021.	Weekly data was reported in $/bu and converted to c/bu.
Supplemental Basis Data	Thomson Reuters, 2021.	Weekly data was reported in $/bu and converted to c/bu.
PNW Rail Tariff Data	Burlington Northern Santa Fe, 2021a.	Current tariff rates were reported in $/car and were converted to c/bu.
St. Louis Tariff Data	Burlington Northern Santa Fe, 2021b.	Current tariff rates were reported in $/car and were converted to c/bu.
Daily Car Value	TradeWest Brokerage, 2021.	Daily car values were reported in $/car and converted to c/bu.

4.4.2. Data Behavior

A summary of the collected data follows which includes graphical representations of the data as well as descriptions of the data behavior.

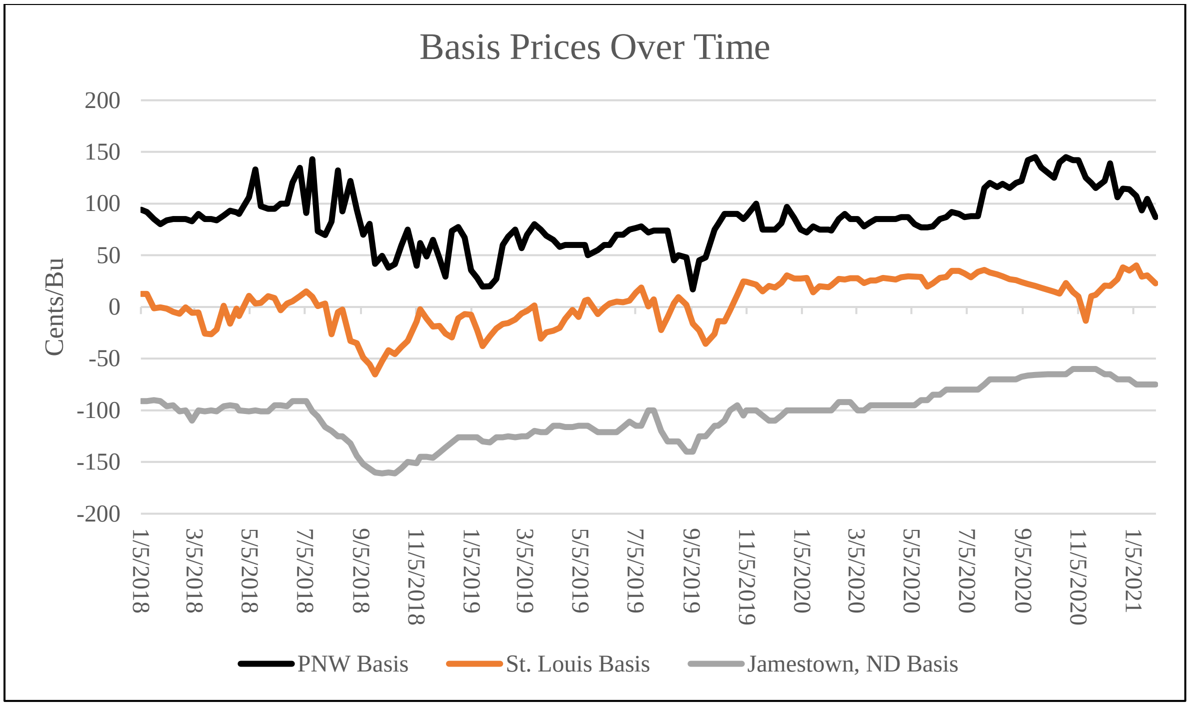

Figure 4.2: Basis Values Over Time

Basis is reported for locations in the Pacific Northwest, St. Louis, and a shipper elevator in Jamestown, ND. The x-axis shows the date, and the y-axis shows the basis values which are reported in cents per bushel. The data exhibits a degree of seasonality indicated by temporary decreases in the basis levels around the months of September, October, and November. The PNW basis shows this the most clearly, but seasonality can be observed in each of the data sets. The seasonal patterns create merchandising opportunities, but risk still exists. For example, the Jamestown basis exhibited basis decreases during the harvest seasons in 2018 and 2019, but in the Fall of 2020, the basis increased while it decreased at the PNW. While the various basis markets tend to be positively correlated, the correlation is not stable (as shown in a following table) and has significant impacts on the results.

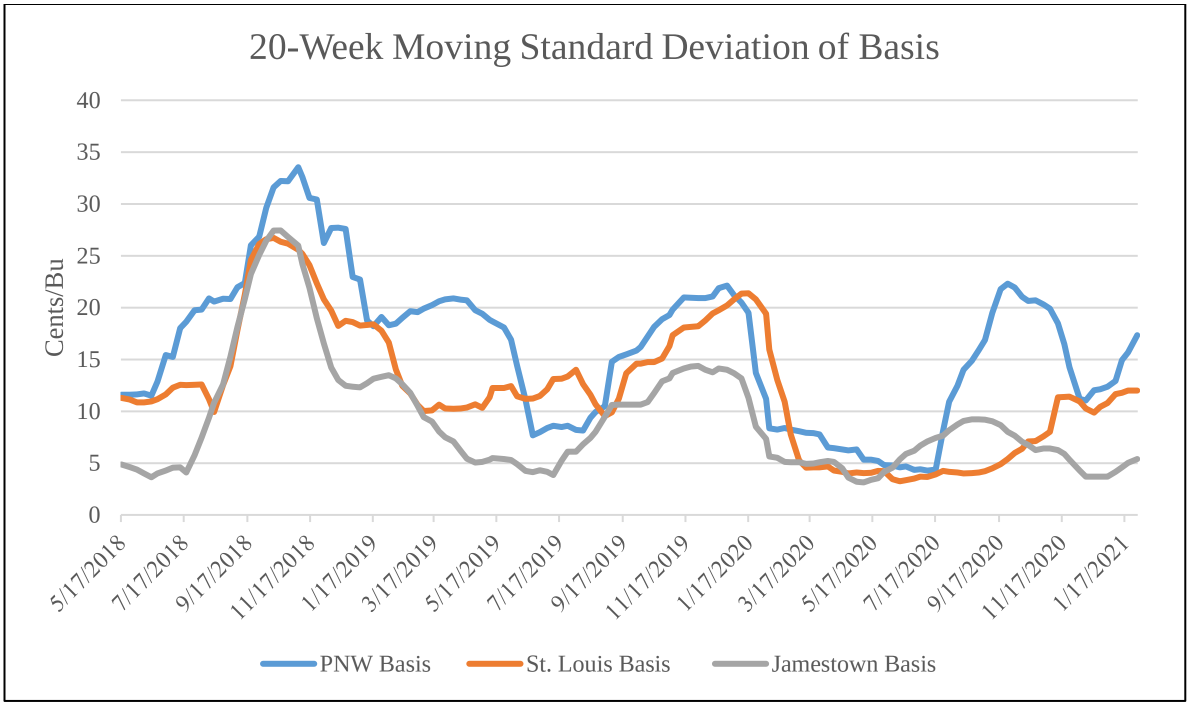

Figure 4.3: 20-Week Moving Standard Deviation of Basis

While Figure 4.2 showed the seasonality of the basis values, Figure 4.3 shows the standard deviation of basis is also seasonal. The time period of 20 weeks was chosen as it is a similar amount of time for the grain positions to be held throughout the estimation procedures. The date listed on the x-axis corresponds to the final week of the observed time period. This means that in the weeks leading up to November, standard deviation tends to increase. While the standard deviations tend to be positively correlated, they do differ at times. For example, in October of 2020, the standard deviation of the St. Louis basis was increasing while the standard deviation of the PNW basis was increasing. The implications of changes in standard deviations within a market are demonstrated using sensitivity analysis in Chapter 5.

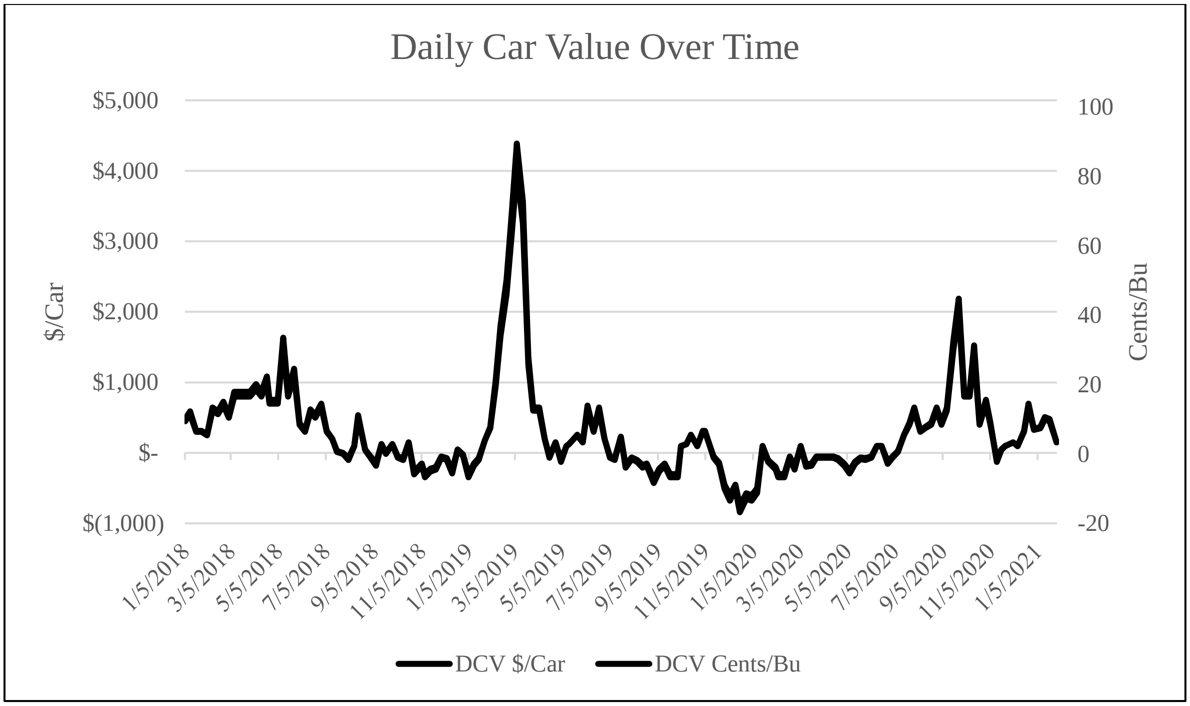

Figure 4.4: Daily Car Values Over Time

Figure 4.4 shows the changes in daily car values over time in dollars per car on the left axis and in cents per bushel on the right axis. The chart demonstrates the significance of freight to the elevator's profit margins, especially compared to the changes in basis values. The daily car values have a larger range and standard deviation, which impacts the risk that is associated with holding an open rail position.

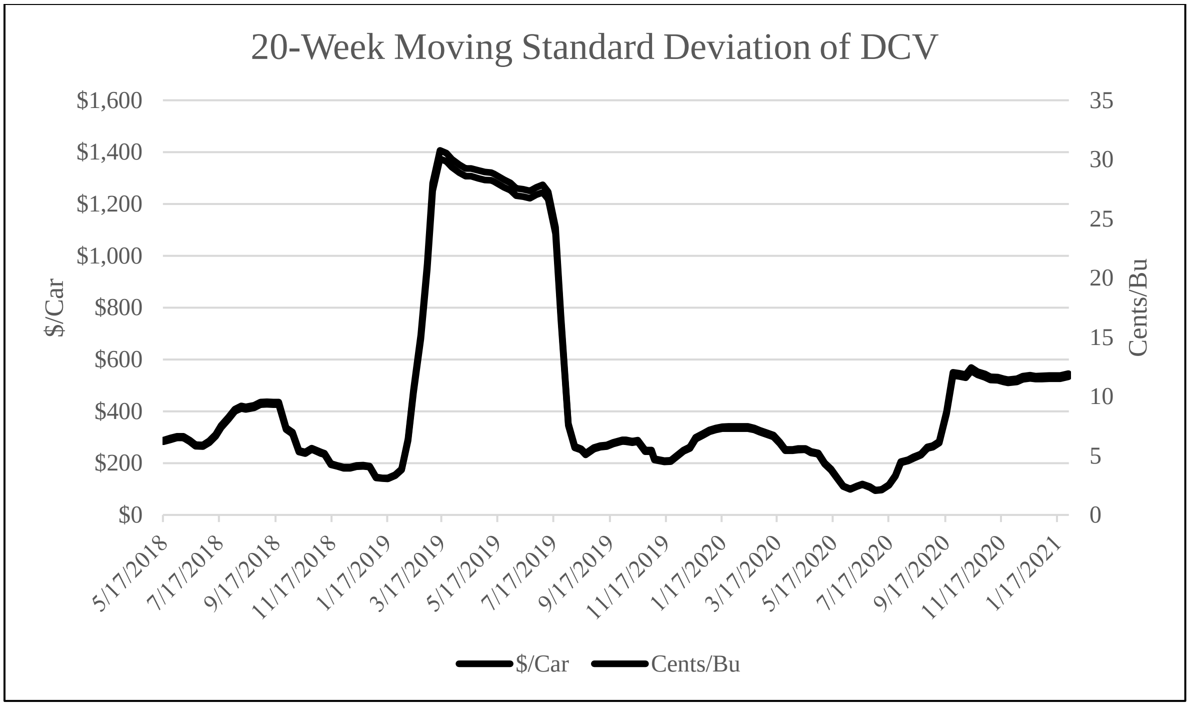

Figure 4.5: 20-Week Moving Standard Deviation of Daily Car Values

Figure 4.5 shows how the standard deviation of DCV changes throughout time. The standard deviation of the rail freight tends to be more stable than the standard deviations of the basis levels. Understanding changes in the standard deviations is important to the estimation procedures as the standard deviation may change based on the data selection process.

In addition to changes in prices and standard deviations that occur in the individual market, the relationships between the markets change over time. Figure 4.6 combines the basis and daily car values to illustrate the relationship between the raw data. Figure 4.7 below shows that in various 20-week periods, the correlations between the different variables change significantly. Thus, the selection of data can alter the correlations observed between the markets, and affect optimal coverage ratios. The effects of these changes are discussed in detail in Chapter 5 using sensitivity analysis.

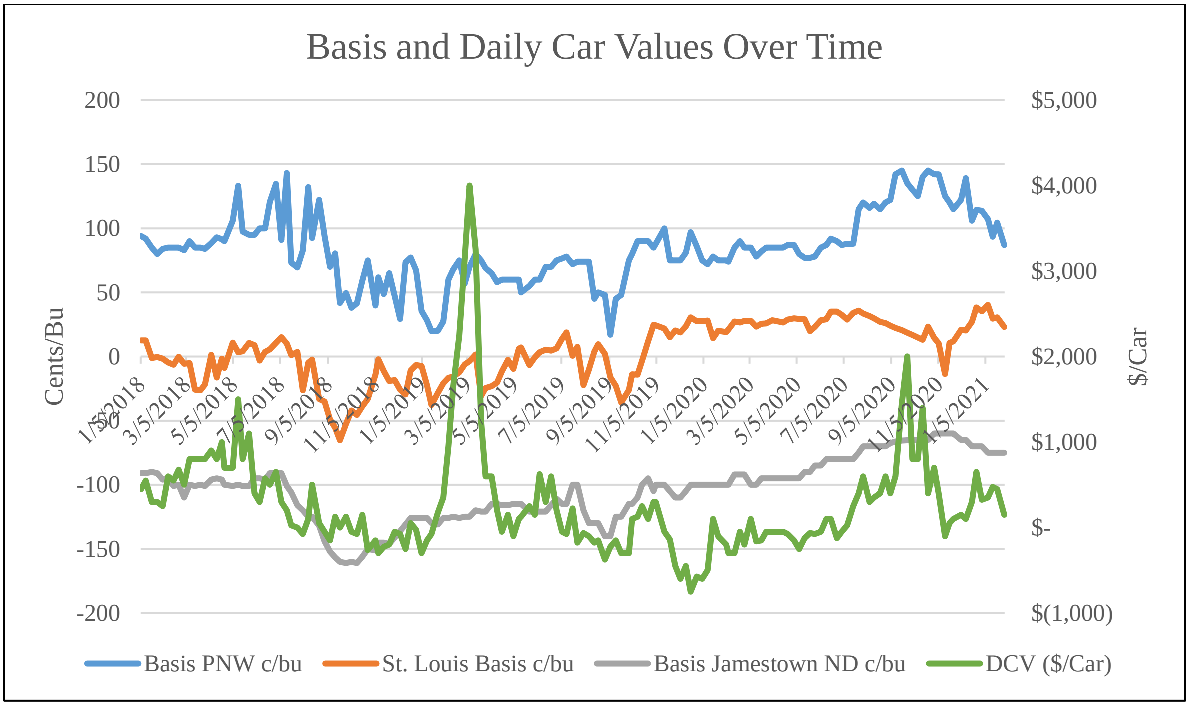

Figure 4.6: Basis and Daily Car Values Over Time

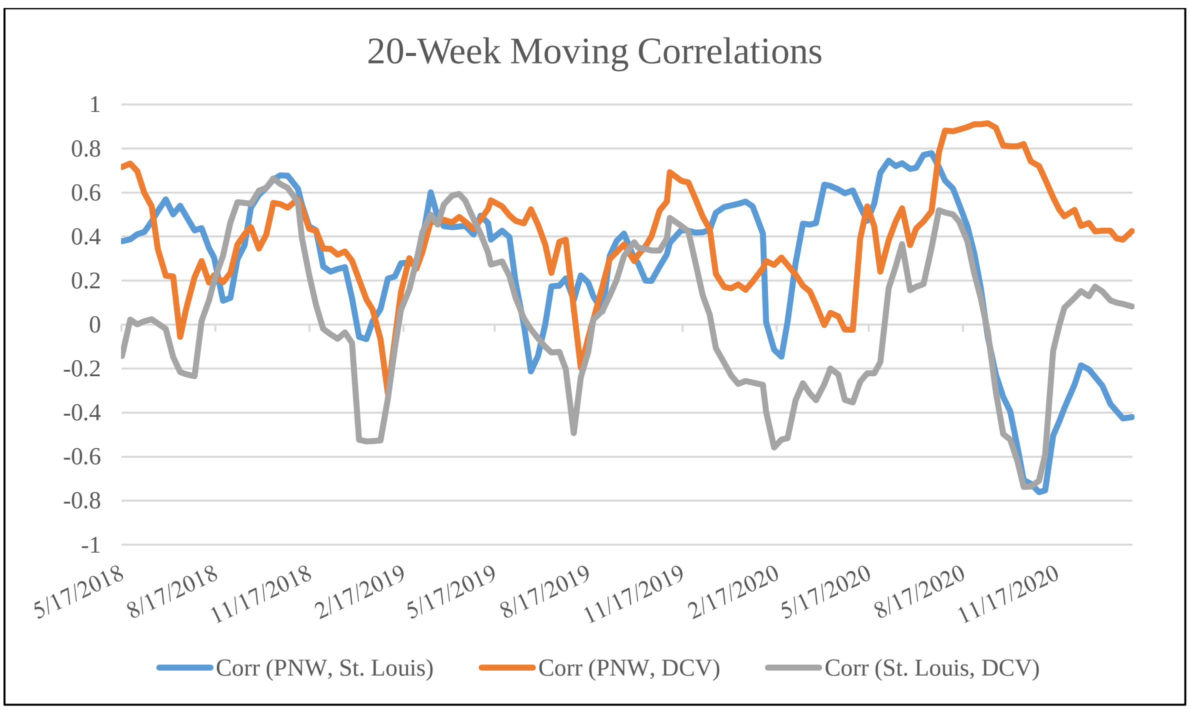

Figure 4.7: 20-Week Moving Correlations

The next set of tables shows the correlations that are used throughout the study for both long-the-basis and short-the-basis models.

Table 4.2: Historical BestFit™ Correlations Used in Long-the-Basis Models

Correlation	PNW Basis	St. Louis Basis	Daily Car Value
PNW Basis	1.000		
St. Louis Basis	-0.617	1.000	
Daily Car Value	0.372	0.020	1.000

Table 4.2 shows the correlations between the variables used in the analysis. It is important to note that depending on the time period chosen, the correlations between the variables can differ drastically as shown by Figure 4.7. The correlation between each of the markets is not surprising except for the correlation between the PNW and St. Louis basis markets. Figure 4.7 shows that this correlation is not necessarily typical, but it is the most recently observed correlation.

Table 4.3 shows the correlations between the Jamestown basis and the daily car values for the data used in the short-the-basis models.

Table 4.3: Historical BestFit™ Correlations Used in Short-the-Basis Models

Correlation	Jamestown Basis	Daily Car Value
Jamestown Basis	1.000	
Daily Car Value	-.055	1.000

4.4.3. Distributions

The five models presented are evaluated under both the mean-variance and mean-semivariance frameworks. The analysis is conducted using historical BestFit™ distributions as well as time series forecasted distributions. This section discusses the usage of these distributions and the benefits and downfalls of each method.

4.4.3.1. Overview of Historical BestFit™ (Naïve) Distributions

The models are first evaluated using historical BestFit™ distributions. These distributions should be interpreted as "naïve distributions" as it assumes the trader expects the distribution of prices and correlations between markets to remain the same across time. The primary advantage of fitting price history, a time series variable, to non-time series distributions is that it allows for the standard deviations of the variables to be easily modified. For example, a variable that has a normal distribution with a given mean and standard deviation could be easily adjusted to show the effects of an increase or decrease of standard deviation on the optimal coverage ratio. Another benefit of using a historical BestFit™ distribution is the ability to compare results against the time series forecasted distributions.

4.4.3.2. Overview of Time Series (Anticipatory) Distributions

One of the disadvantages of historical distributions is the reliance on prices that are one day old or one year old which creates an additional problem of trying to determine the proper period to use in the model. Time series forecasting solves this issue by placing more weight on the events that happened most recently as this type of distribution creates projections for future prices. With this type of strategy, the trader is assumed to anticipate the timeframe for liquidation. Additionally, the standard deviation of the forecast gets larger as the forecasting time period gets larger. Figure 4.8 demonstrates this below.

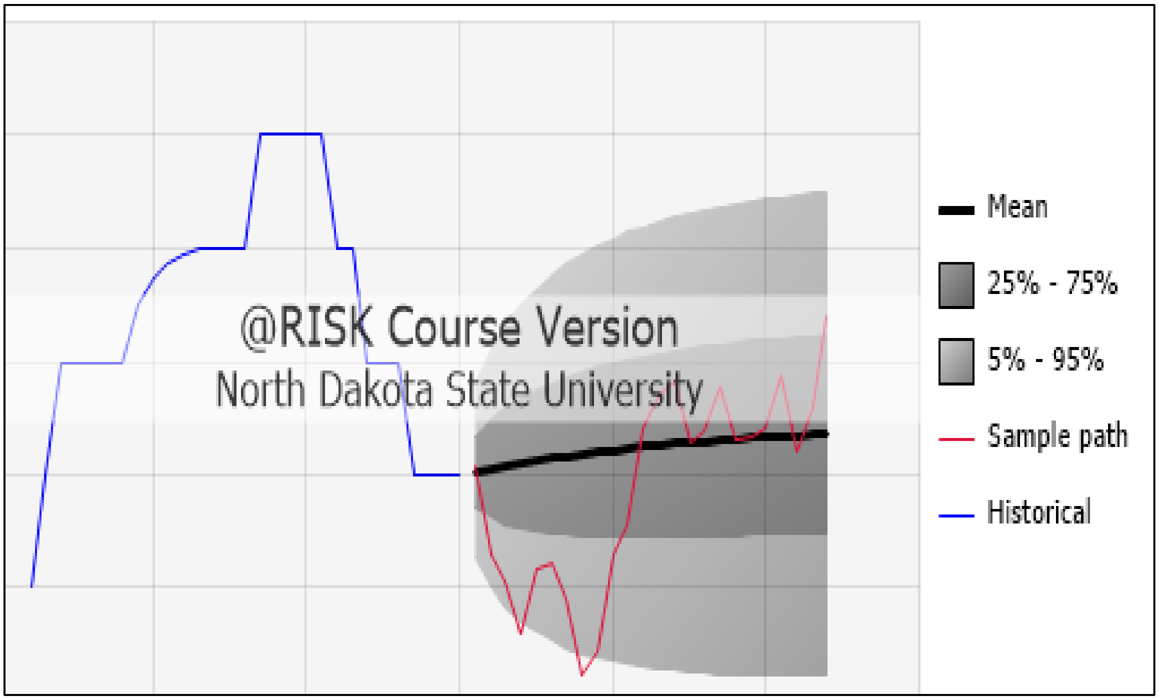

Figure 4.8: Increasing Standard Deviation of Forecast Over Time

The increasing standard deviation is shown by the shaded cone. The confidence level of the forecast is depicted by the lighter and darker shades.

4.4.3.3. Historical BestFit™ (Naïve) Distributions and Fit Parameters

The BestFit™ function of the @Risk™ software package is used to fit distributions to sets of data based on Bayesian information criterion (BIC). For purposes of consistency, the data used in the historical BestFit™ distributions is the same as what is used in the time series distributions shown in the next section. The following figures show the fitted distributions and correlations.

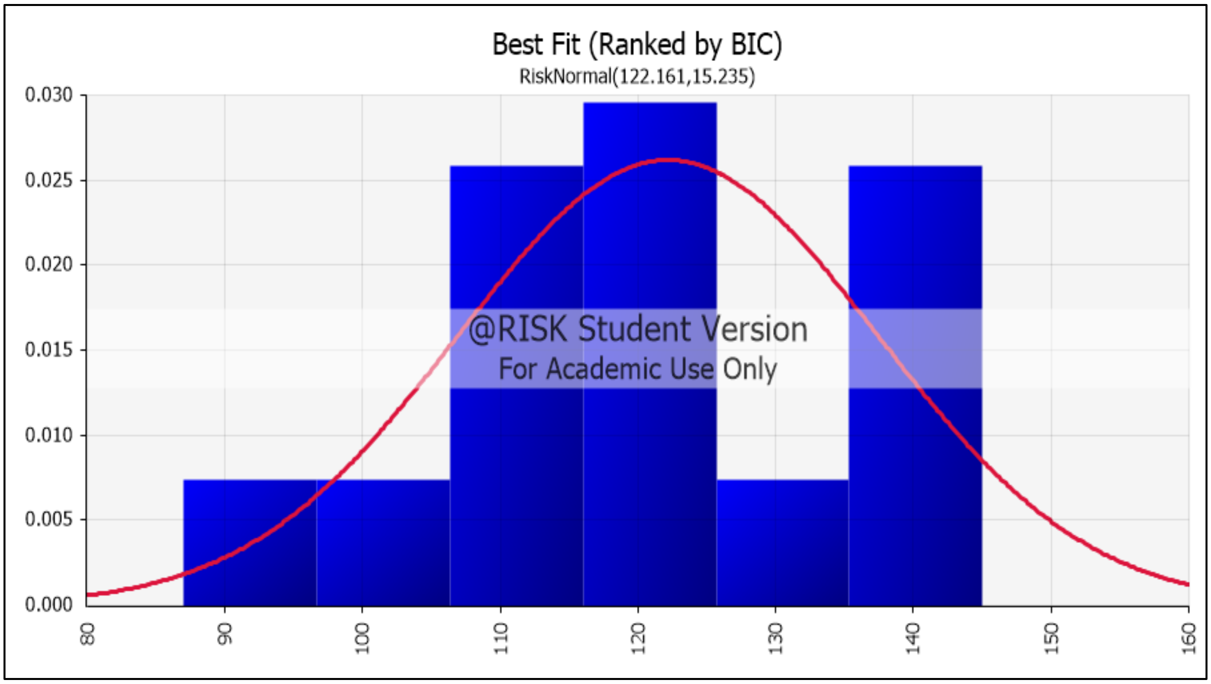

Figure 4.9: Historical BestFit™ Distribution of PNW Basis

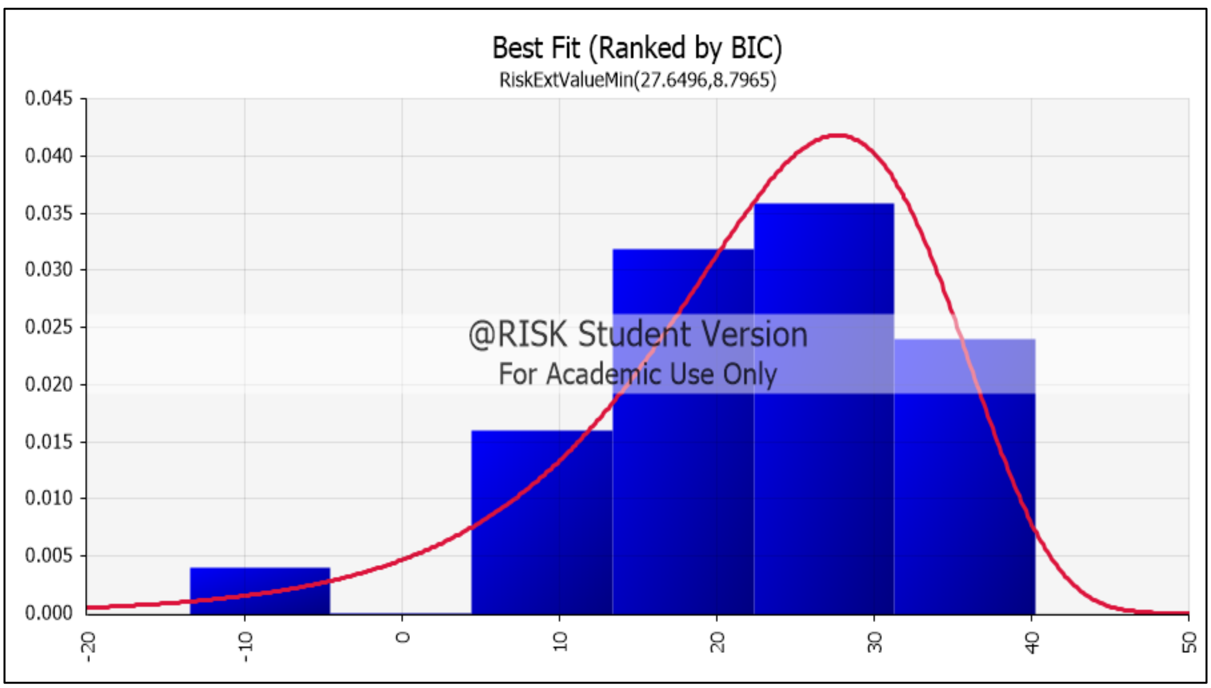

Figure 4.10: Historical BestFit™ Distribution of St. Louis Basis

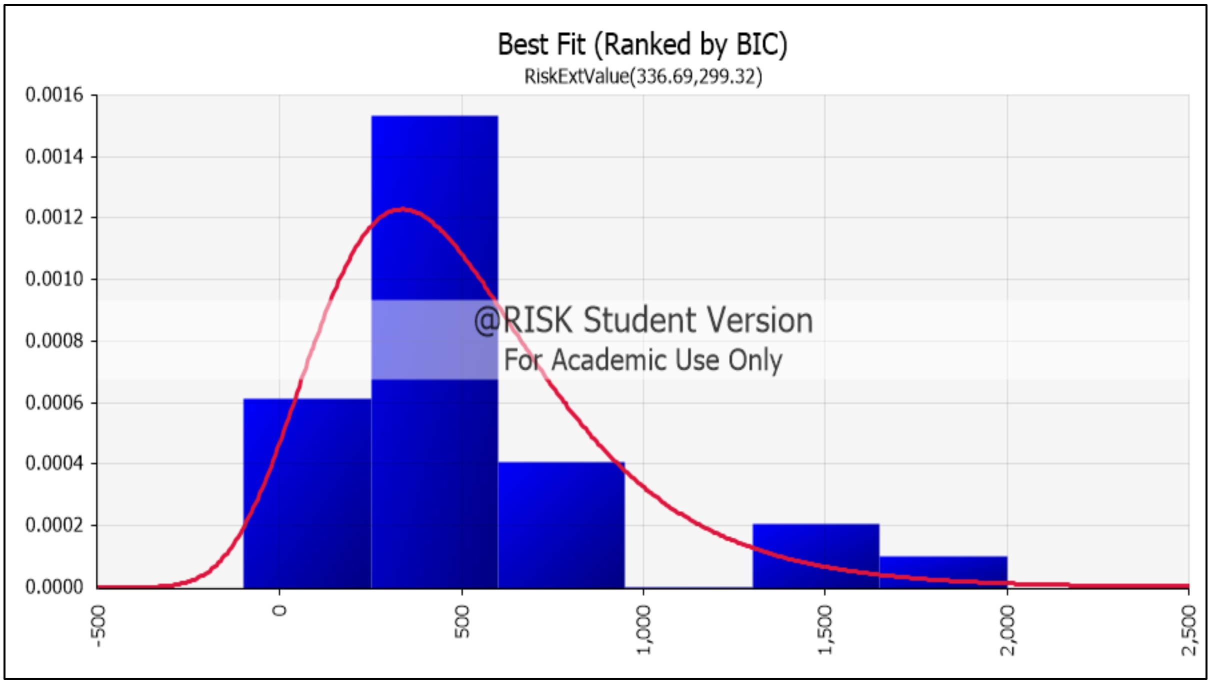

Figure 4.11: Historical BestFit™ Distribution of Daily Car Values

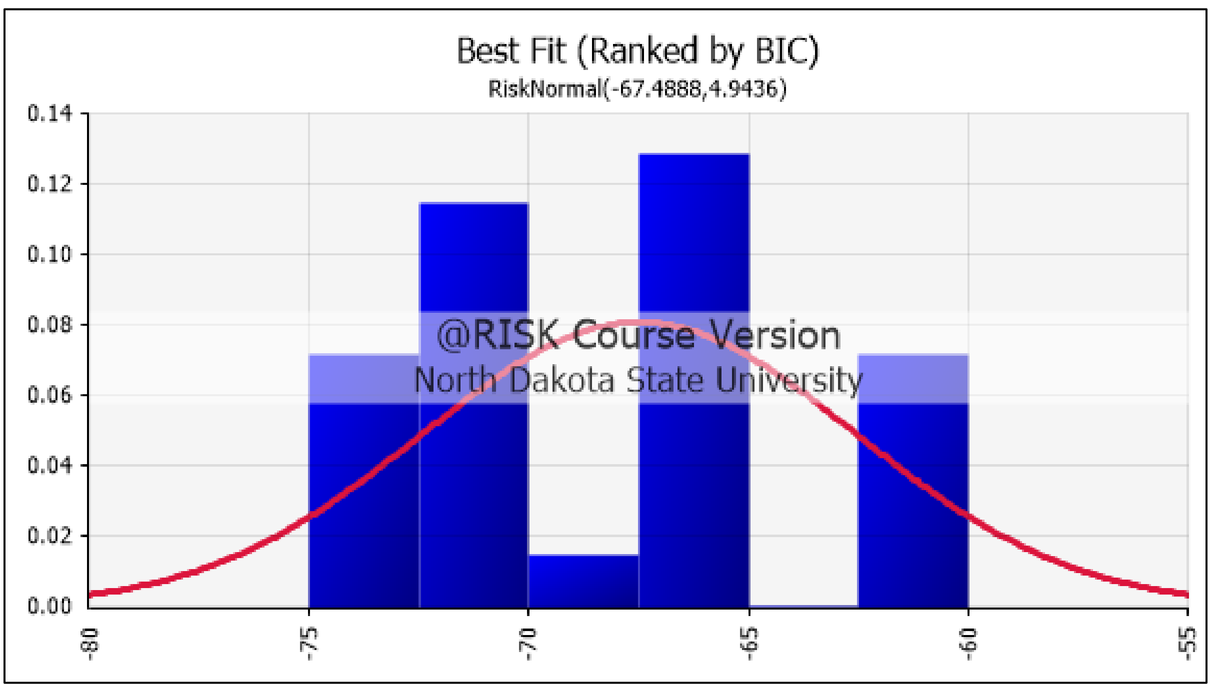

Figure 4.12: Historical BestFit™ Distribution of Jamestown Basis

Table 4.4: @Risk™ Historical BestFit™ Distribution Functions

Variable	Distribution	Function	BIC Score
PNW Basis	Normal	RiskNormal(122.161,15.235)	237.64546
St. Louis Basis	ExtValueMin	RiskExtValueMin(27.6496,8.7965)	215.997400
Daily Car Value	ExtValue	RiskExtValue(336.69,299.32)	416.5110
Jamestown Basis	Normal	RiskNormal(-67.4888,4.9436)	174.618301

Table 4.4 shows the distributions and the functions that have been assigned to each variable along with the BIC score.

4.4.3.4. Time series (Anticipatory) Distributions and Fit Parameters

An alternative, and arguably more applicable, method of approaching the coverage ratio problem is utilizing time series forecasting methods to identify a range of possible basis values for deferred time periods. @Risk™'s time series BestFit™ function allows time series data to be fitted to three different categories of distributions including autoregressive moving average (ARMA), geometric Brownian motion (GBM), and autoregressive conditional heteroscedasticity (ARCH) ("Time Series Functions" 2021). Furthermore, BestFit™ can detect and correct for trends, seasonality, or stationary before creating the distribution.

The distributions that were used for the time series analysis are presented below.

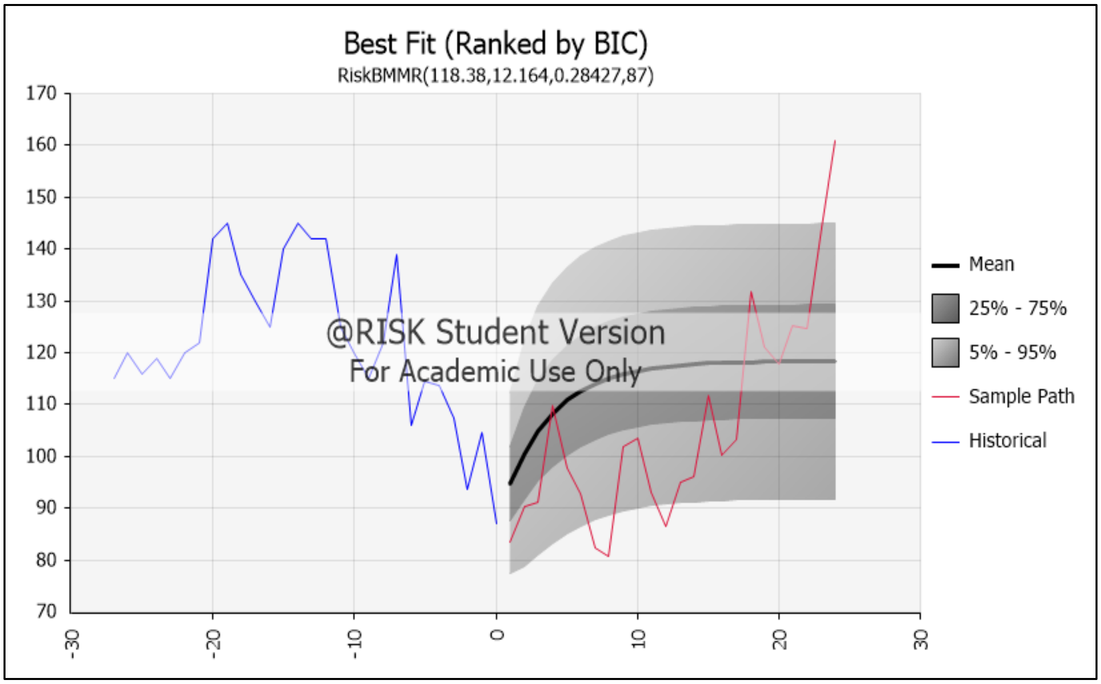

Figure 4.13: Time Series BestFit™ Distribution of PNW Basis

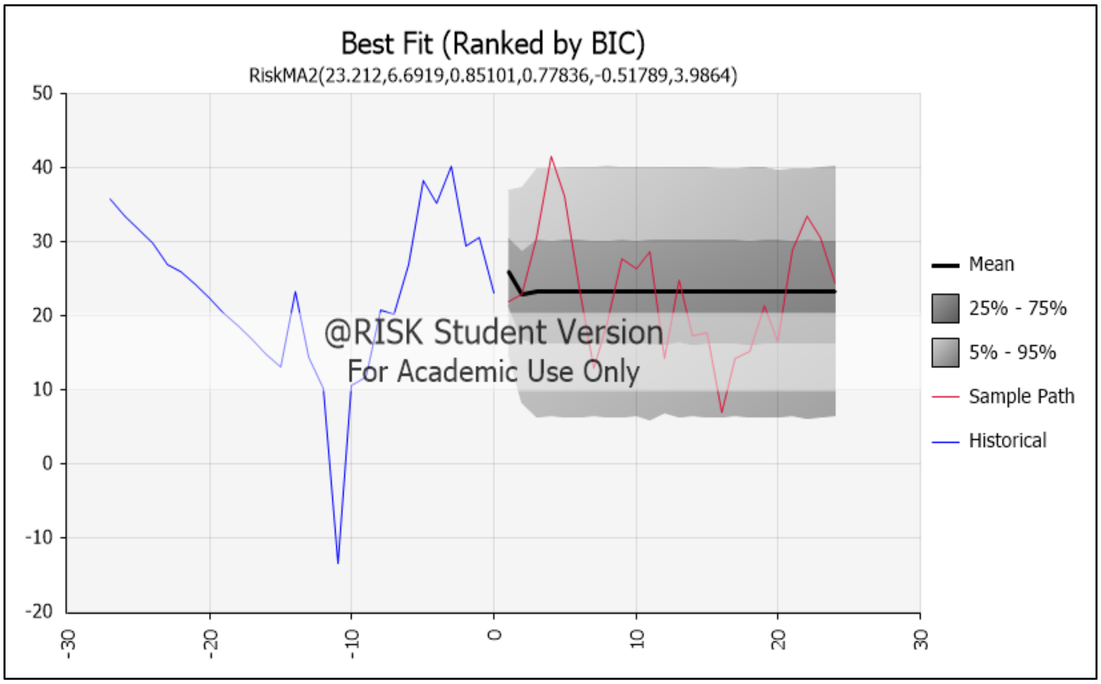

Figure 4.14: Time Series BestFit™ Distribution of St. Louis Basis

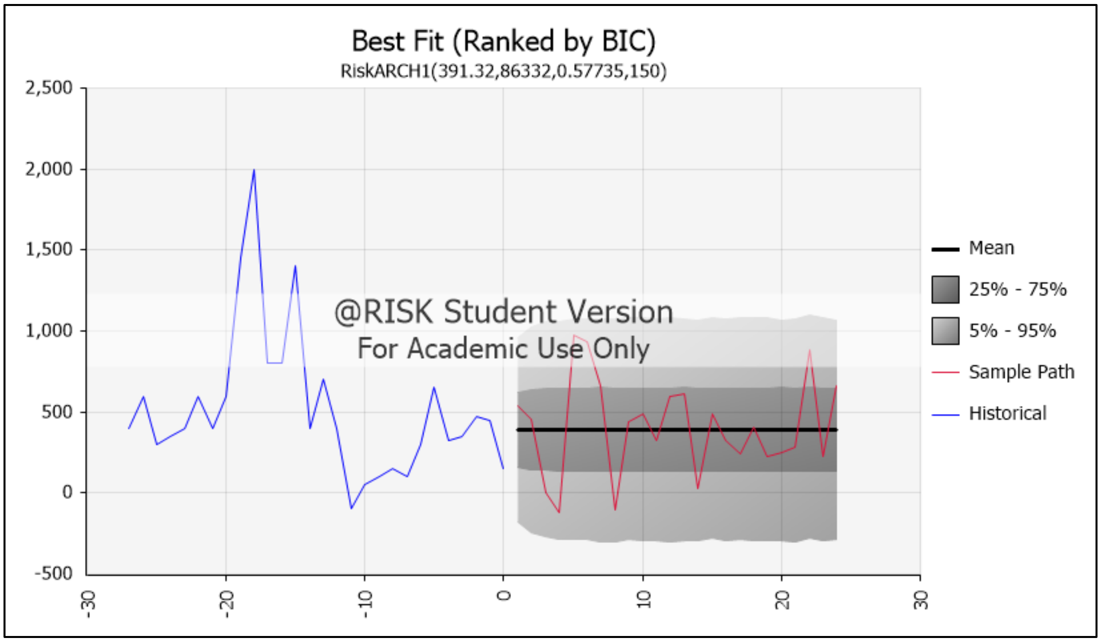

Figure 4.15: Time Series BestFit™ Distribution of Daily Car Values

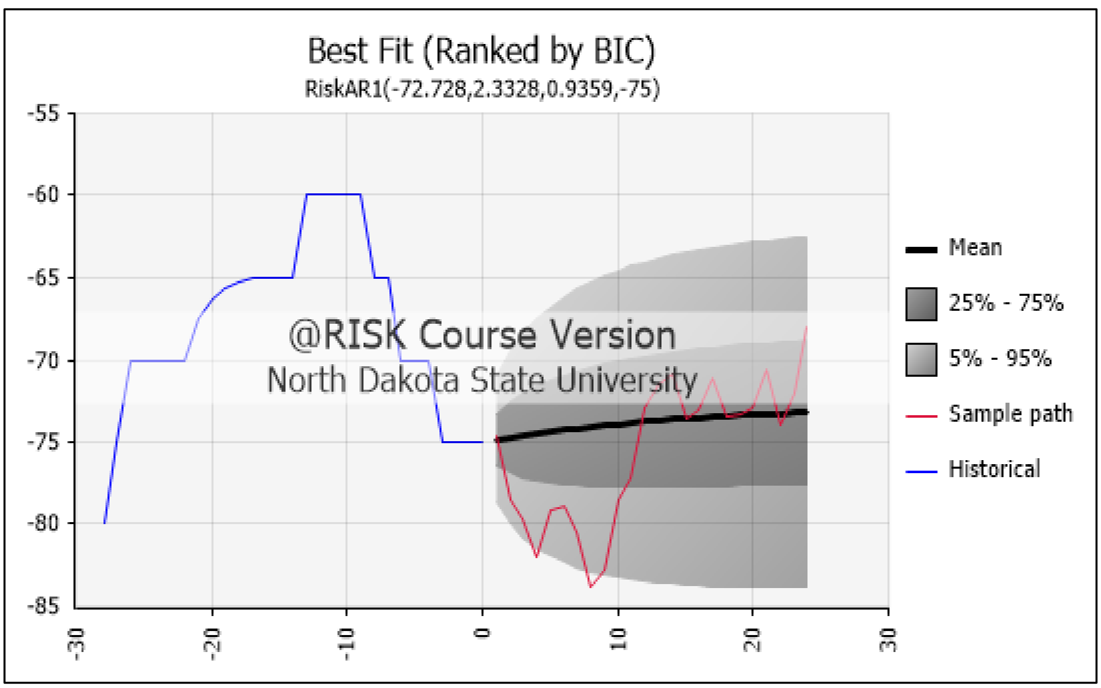

Figure 4.16: Time Series BestFit™ Distribution of Jamestown Basis

Table 4.5: @Risk™ Time Series BestFit™ Distribution Functions

Variable	Distribution	Function	BIC Score
PNW Basis	BMMR	RiskBMMR(118.38,12.164,0.28427,87)	222.6258411
St. Louis Basis	MA(2)	RiskMA2(23.212,6.6919,0.85101,0.77836,-0.51789,3.9864)	201.3605192
Daily Car Value	ARCH(1)	RiskARCH1(391.32,86332,0.57735,150)	409.1281833
Jamestown Basis	AR(1)	RiskAR1(-72.728,2.3328,0.9359,-75)	143.6165121

Table 4.5 shows the time series distributions and the functions that have been assigned to each variable along with the BIC score.

4.5. Simulation Procedures

The empirical model was evaluated using the Palisade DecisionTools Suite™. Monte Carlo simulations are used within RiskOptimizer™ and @Risk™ to estimate the optimal coverage ratios by changing the coverage ratios in each model to maximize the expected utility. @Risk™ offers a variety of settings for simulating results. Sampling methods, random value generation, and the number of iterations can affect the overall results. The @Risk settings used throughout the study are shown in Table 4.6.

Table 4.6: @Risk™ Simulation Settings

@Risk™ Specification	@Risk™ Setting
Sampling Type	Latin Hypercube
Generator	Mersenne Twister
Initial Seed Value	Fixed, 1

For each of the Monte Carlo simulations, 250 iterations were performed using Latin Hypercube sampling. The models converged quickly during the estimation procedures which allowed for a relatively small number of iterations. In base case estimations, the coverage ratios were allowed to vary between 0% and 100% with step sizes of 0.5%. The effects of varying these constraints are shown in Chapter 5.

4.6. Summary

This chapter described the models and variables that were used to find optimal coverage ratios for grain and rail positions in a portfolio under a mean-variance or mean-semivariance framework. The assumptions for the analysis followed which included the need for the elevator to hold a naive futures hedge to be considered a pure basis trader. Additionally, historical and time series BestFit™ estimation procedures were presented along with the data and distributions. Important characteristics of the data were highlighted such as varying correlations between variables and changing standard deviations. The nature of the unstable relationships among variables prompts the need for sensitivities to be conducted in later chapters. Finally, the settings used for the @Risk™ simulation procedures were presented.

CHAPTER 5. RESULTS

5.1. Introduction

Volatility and randomness of commodity prices in the markets have resulted in discussions on risk management, particularly in the realm of futures markets and futures hedging such as Blank, Carter, and Schmiesing (1991) on optimal hedge ratios. Limited research has been published relating to the changing basis, volatilities, and correlations between basis and transportation markets. Because elevators tend to be fully hedged, their primary sources of margin risk come from changes in basis or transportation cost. This study focuses on optimizing the coverage ratios for a portfolio of soybean basis and rail freight under the mean-variance and mean-semivariance frameworks.

Chapter 3 formally defined coverage ratios for basis and rail portfolios and introduced the models used in the empirical analysis. Theoretical frameworks for optimizing hedge ratios were presented with an analytical derivation of the optimal coverage ratio for a simple portfolio with only grain and more complex portfolio including grain and rail.

The data presented in Chapter 4 shows the volatility in the basis and rail markets as well as the impact that rail prices can have on the overall margin. All of the models were presented for both a long-the-basis and a short-the-basis scenario. Historical BestFit™ and time series forecasted distributions were developed for each random variable. The historical BestFit™ models assume a naïve merchant that relies solely on historical data where the time series distributions represent a merchant that is anticipating future basis and rail market behavior. Monte Carlo simulations were conducted using @Risk to evaluate the various models.

This chapter shows the simulation results beginning with the historical BestFit™ distributions in a long-the-basis scenario. The historical distributions are evaluated under the mean-variance and mean-semivariance framework. Next, sensitivity analyses are conducted which show the effects of changing parameters within the model. A similar presentation of results follows for the short-the-base case using historical distributions. The next section presents the time series forecasted models which allow for the specification of a liquidation and entry date for the basis and rail positions. Further sensitivities are conducted to illustrate the effects of varying liquidation dates and entry dates. Lastly, the short-the-basis results are presented with similar sensitivities.

Throughout this chapter, the model numbers shown correspond to those described in Chapter 4. Models that use historical BestFit™ distributions begin with an "H" followed by the model number where the models using time series forecasts begin with "TS". Throughout the analysis, expected profit is defined as E(π), and the standard deviation of profit is denoted by σ. Each of these values are reported in cents per bushel (c/bu). For the long-the-basis models, the optimal coverage ratios are defined as h_A, h_B, or g. Variable h_A represents the coverage taken at Destination A, in this case, the PNW; h_B is the coverage taken at Destination B, St. Louis; and g is the coverage taken in rail. For the short-the-basis models, the optimal coverage ratios are defined as h_O and g where variable h_O is the coverage taken at the origin, Jamestown, ND, and g is still the coverage taken in rail. Each of these ratios are listed as a percentage.

5.2. Historical BestFit™ Results

This section presents the results that used historical BestFit™ distributions of the random variables. Using historical BestFit™ distributions provides results that can be compared to the time series forecasted distributions. Additionally, the standard deviations of variables are easily modified with this set of distributions which allows for sensitivity analyses to be conducted.

When using historical BestFit™ distributions, a specific liquidation date is not defined so it is assumed that liquidation occurs some time after the position entry date.

5.2.1. Long-the-Basis Base Case Under E-V and E-SV

The model specifications for the long-the-basis base case model are presented in Table 5.1. The table shows the input name, value, and units.

Table 5.1: Long-the-Basis Base Case H Model Specifications

Specification	Input Value	Units
Time 1	12/24/2020	Month/Day/Year
Time 1 Daily Car Value	650	$/car
Time 1 Basis: Jamestown, ND	-70	c/bu
Time 1 Basis: PNW	114.5	c/bu
Time 1 Basis: St. Louis, MO	38.25	c/bu
Model Number	Varies	
Risk Aversion Level	Varies	

Table 5.2 presents different results of the three models under the mean-variance and mean-semivariance frameworks. Each model is evaluated at varying levels of phi, Φ, the coefficient of risk aversion. The cells highlighted in yellow are used as the base case throughout the remainder of Sections 5.2.1 and 5.2.2.

Table 5.2: Base Case Long-the-Basis Historical BestFit™

			E-V Models					E-SV Models				
Phi	Model Name	Description	E(π)	σ	h_A	g	h_B	E(π)	σ	h_A	g	h_B
0.05	H 1.1	Long Basis No Freight Risk	21.72	9.89	36%			24.42	15.33	0%		
	H 2.1	Long Basis Short Freight	26.67	12.71	12%	0%		27.59	14.20	0%	0%	
	H 3.1	Long Basis Multiple Destinations Short Freight	28.38	13.05	0%	0%	0%	28.38	13.05	0%	0%	0%
0.10	H 1.1	Long Basis No Freight Risk	19.29	4.98	68%			21.65	9.74	37%		
	H 2.1	Long Basis Short Freight	**23.13**	**7.95**	**50%**	**22%**		27.09	13.38	7%	0%	
	H 3.1	Long Basis Multiple Destinations Short Freight	24.96	8.96	36%	12%	0%	28.38	13.05	0%	0%	0%
0.15	H 1.1	Long Basis No Freight Risk	18.45	3.30	79%			20.05	6.52	58%		
	H 2.1	Long Basis Short Freight	21.00	5.27	67%	48%		25.25	10.69	29%	5%	
	H 3.1	Long Basis Multiple Destinations Short Freight	22.29	6.03	57%	41%	0%	27.54	11.94	10%	0%	0%

The results in Table 5.2 show how profit and standard deviation decrease as phi and coverage ratios increase as well as the differences in optimal coverage ratios between the different models and different frameworks. First, the effects of coverage ratio, profit, and standard deviation are shown in the trials of H 1.1 in which the elevator is holding inventory but has no rail risk. At the low phi, 0.05, the optimal coverage ratio was 36% under the E-V framework with a profit of 21.72 c/bu and a standard deviation of 9.89 c/bu. At a higher level of phi, 0.15, the optimal coverage ratio was 79% with a profit of 18.45 c/bu and a lower standard deviation of 3.30 c/bu. This means more risk-averse traders would sell more of their inventory to the destination during Time 1, leaving less subject to changing basis. Overall, as traders become more risk-averse, i.e. higher levels of phi, greater profits are exchanged for lower standard deviations of profit.

When evaluating the same model under the E-SV framework, the coverage ratios increase from 0% to 37% to 58% with profits decreasing from 24.42 c/bu to 20.05 c/bu. The standard deviation decreases from 15.33 c/bu to 6.52 c/bu. The differences in coverage ratios and the resulting profits and standard deviations are important because mean-semivariance is meant to only mitigate downside risk where mean-variance reduces all variance.

The same patterns exist in the results for Model H 2.1. In Model H 2.1, the elevator has a long grain position and a short freight position. The portion of freight booked at Time 1 is represented by variable g. As rail is added to the portfolio, profit increases from Model H 1.1 at all levels of risk aversion. At low levels of phi, optimal rail coverage is 0% with grain coverage at 12% and an expected profit of 26.67 c/bu, and at high levels of phi, optimal rail and grain coverage increase to 48% and 67% respectively, with profits decreasing to 21 c/bu.

In Model H 3.1, the elevator is long grain; they may sell to two possible destinations, and they need to purchase freight. The coverage ratio for h_B is always 0% due to the delivered price received at St. Louis was always lower than the available delivered price to the PNW during Time 1. While taking coverage at St. Louis may not be optimal at Time 1, the elevator may still sell to St. Louis in the spot market at Time 2 which may explain the higher expected profits observed as the second destination was added.

The distributions of profit are shown in Figure 5.1 and Figure 5.2 to illustrate how changing risk aversion changes the mean and standard deviation of profits for both the mean-variance and mean-semivariance frameworks.

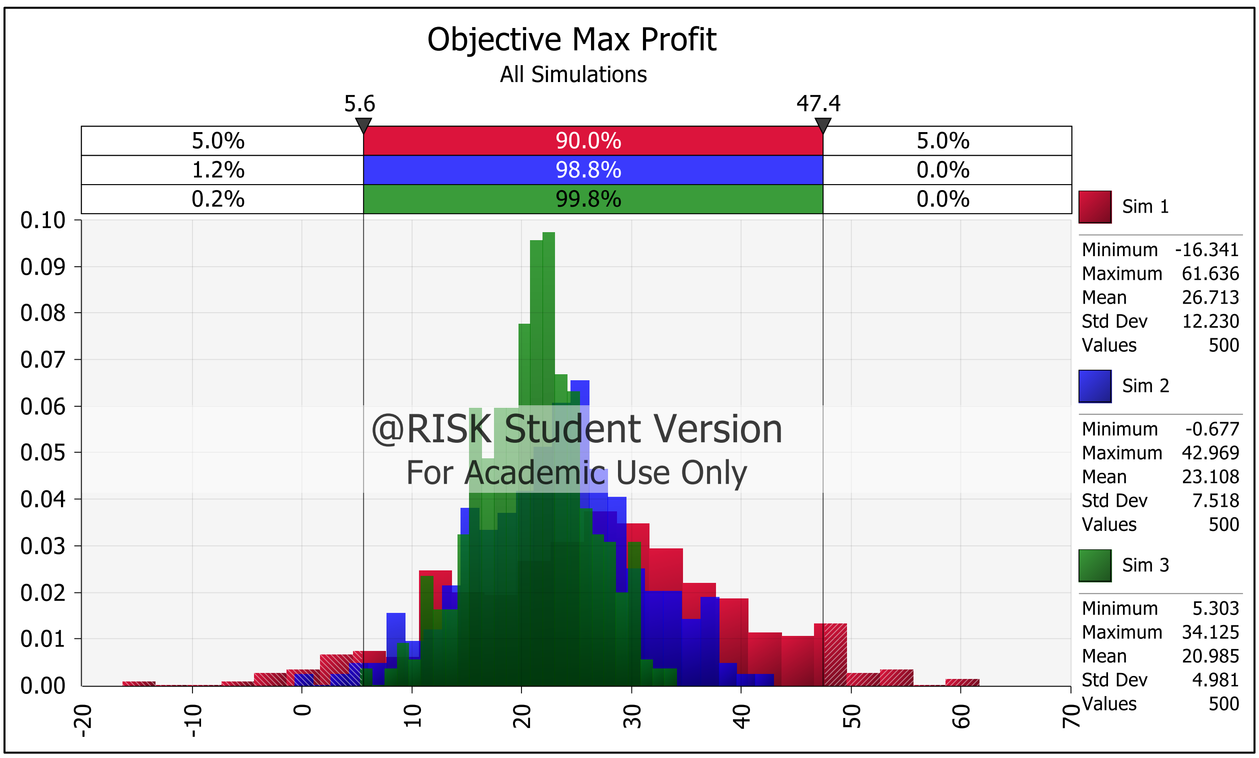

Figure 5.1: Profit Distributions from Model H 2.1 at Various Levels of Phi Under E-V

Figure 5.1 shows the profit distributions from Model H 2.1 under the E-V framework with various levels of phi. Sim 1, Sim 2, and Sim 3 correspond to levels of phi of 0.05, 0.10, and

0.15, respectively. This depiction shows that as risk aversion increases, mean profit and variation of returns decrease.

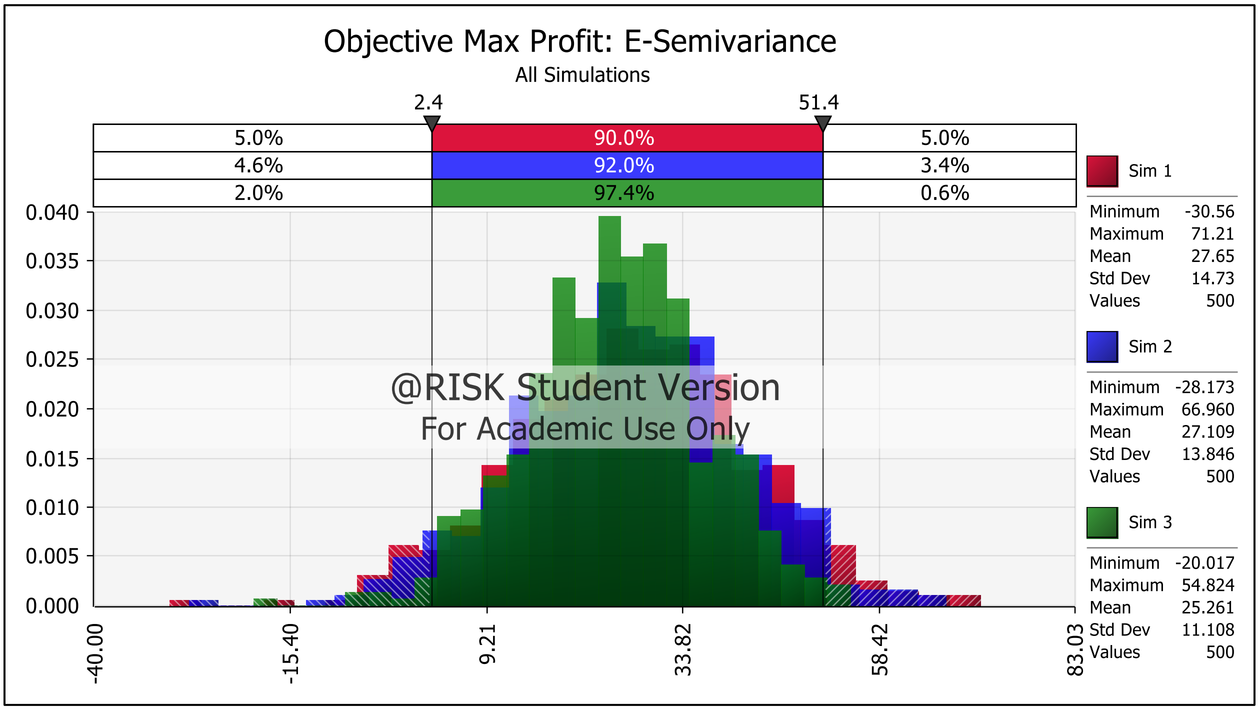

Figure 5.2: Profit Distributions from Model H 2.1 at Various Levels of Phi Under E-SV

Figure 5.2 shows the profit distributions from Model H 2.1 under the E-SV framework with various levels of phi. As in Figure 5.1, Sim 1, Sim 2, and Sim 3 correspond to levels of phi of 0.05, 0.10, and 0.15, respectively. Because mean-semivariance distinguishes between upside risk and downside risk, the levels of coverage taken under the E-SV framework are lower resulting in higher expected profit and larger standard deviations.

The relationship between profit and variance of profit was represented graphically by Markowitz (1952). Figure 5.3 below shows a variety of coverage ratios ranging between -150% and 150% evaluated on Model H 1.1 at a risk aversion level of 0.10, the base level phi. The E-V curve shows as the profit increases, the variance of the profit increases. As h_A approaches 100%, basis risk approaches zero. At this level of coverage, the expected profit is 16.8 c/bu and the variance of profit is zero. This results in lower profits and lower variation in profits. In this case, as h_A approaches zero and negative values, there is less grain sold at the destination than what is owned at the origin creating risk. To achieve a negative coverage ratio, grain would need to be owned at both the origin and the destination at Time 1. In this case, as the coverage ratios are reduced from 100% and become negative, expected profits and the variance of profits increases. Mathematically, the optimal coverage is found by finding the point where the decision maker's utility curve is tangent to the E-V curve. Representing these results graphically clarifies that the approach is consistent with utility maximization theory.

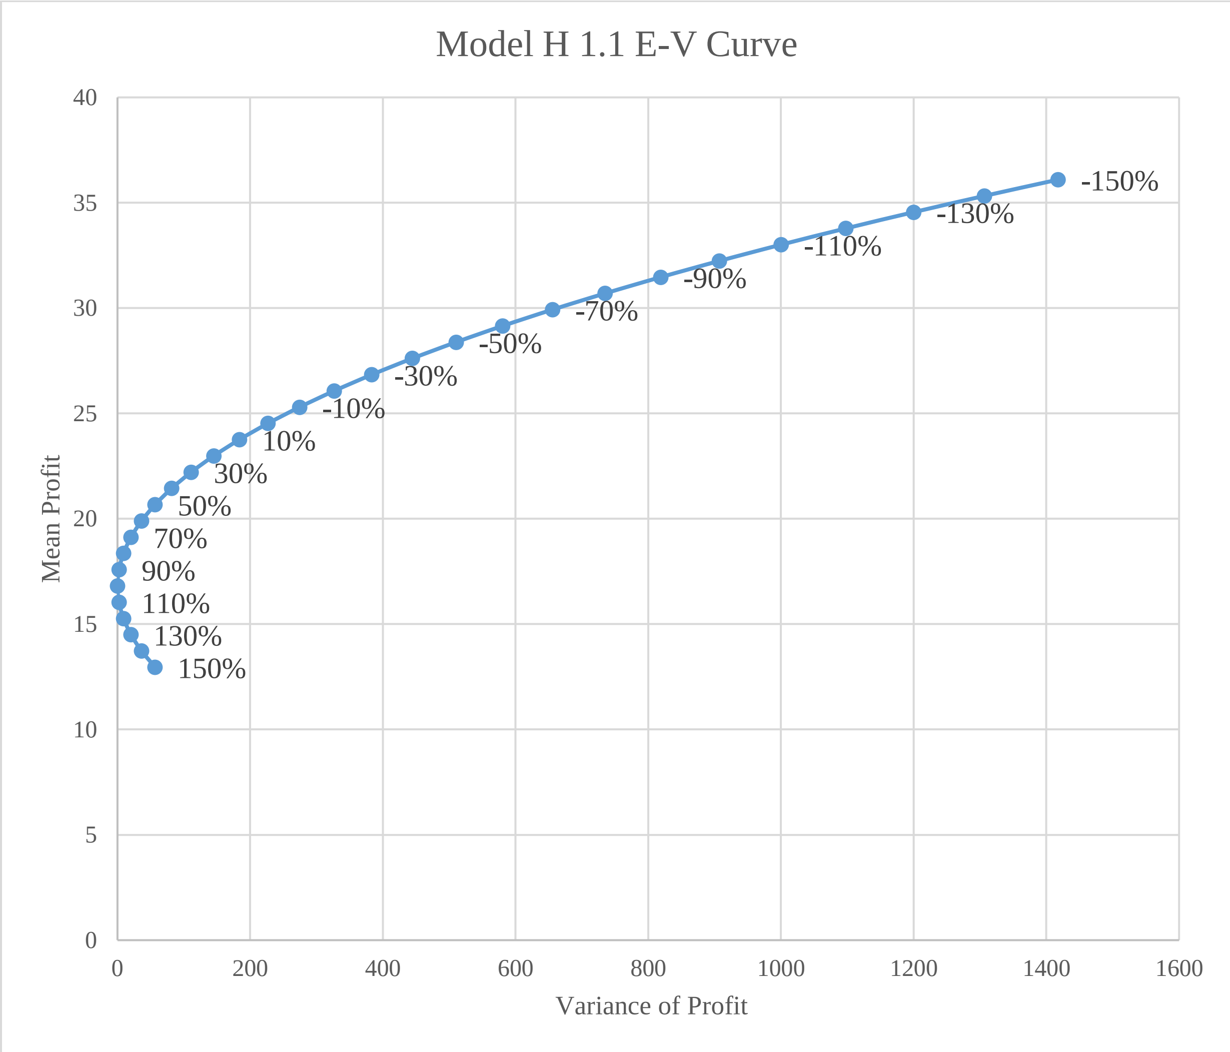

Figure 5.3: E-V Curve from Model 1.1

Figure 5.4 depicts which random variables contribute to the variance of the profit at varying levels of risk aversion when evaluated on Model H 2.1. Similar to Figure 5.1 and Figure 5.2, Sim 1, Sim 2, and Sim 3 correspond to levels of phi of 0.05, 0.10, and 0.15, respectively. This chart shows that in Simulation 1, 61.3% of the variance was due to the variance in the basis, but as the risk aversion increased, the variance of profit was due to the changes in the rail market.

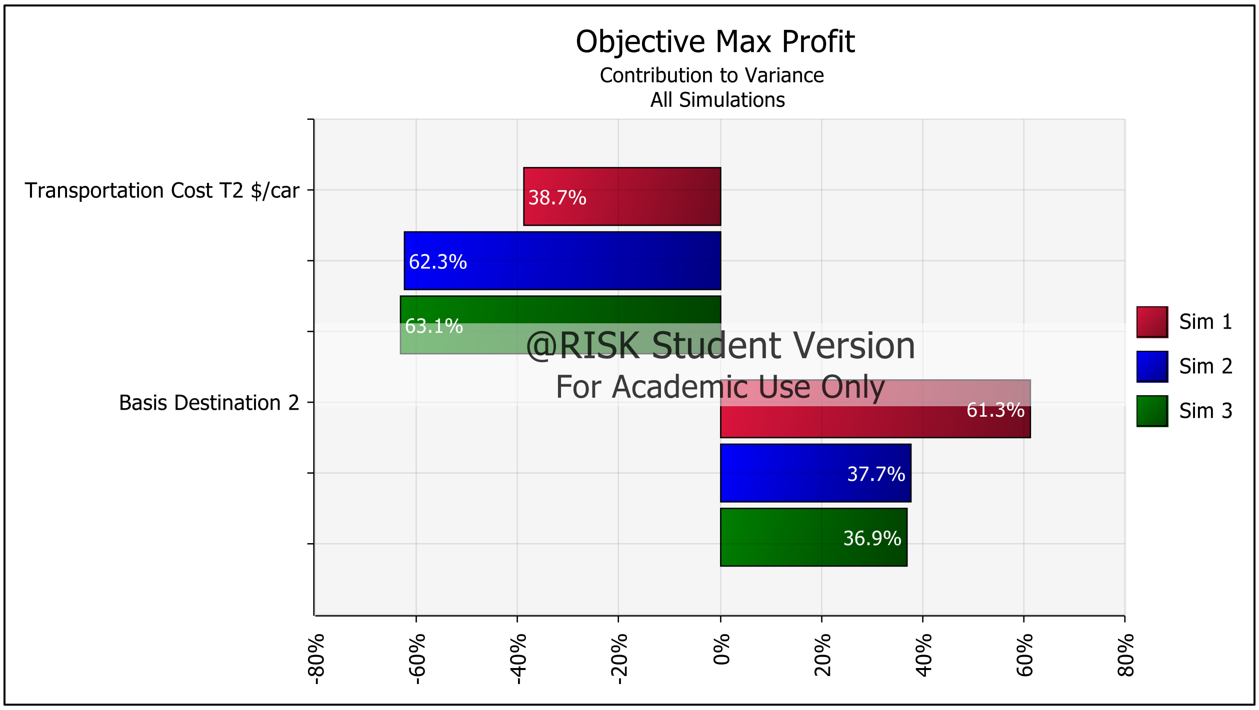

Figure 5.4: Contribution to Variance of Profit in Model H 2.1 With Varying Phi Under E-V

5.2.2. Long-the-Basis Sensitivities

A variety of factors have a significant impact on coverage ratios including the correlations between the destination basis and the daily car values or the standard deviations of basis and daily car values. The following tables show sensitivity analyses that were conducted to show the effects of these parameters on the coverage ratios. Each of the sensitivities are conducted on Model H 2.1 under the E-V framework with the parameters shown in Table 5.3 unless otherwise noted. The base case from Table 5.2 is highlighted in yellow in each table for easy comparison.

Table 5.3: Long-the-Basis Model Specifications for Sensitivities

Specification	Input Value	Units
Time 1	12/24/2020	Month/Day/Year
Time 1 Daily Car Value	650	$/car
Time 1 Basis: Jamestown, ND	-70	c/bu
Time 1 Basis: PNW	114.5	c/bu
Time 1 Basis: St. Louis, MO	38.25	c/bu
Model Number	H 2.1	
Risk Aversion Level	0.10	

Table 5.4: Sensitivity of Correlation Between PNW Basis and DCV

Correlation	Description	E(π)	σ	h_A	g
0.0	Lower Correlations	20.49	6.06	67%	63%
0.372	Base Case	**23.13**	**7.95**	**50%**	**22%**
0.9	Higher Correlation	26.26	6.45	19%	0%

Table 5.4 shows the effects of a change in correlation between the rail cost and the destination prices. The results suggest that when no correlation exists between the two markets, optimal coverage levels are higher for both destination basis and rail. When the markets are closely correlated, the optimal coverage ratios are lower for both destination basis and rail as coverage in one market also helps to cover the risk that is occurring in the other market.

Table 5.5 and Table 5.6 show the effects of a change in standard deviations of the PNW Basis and the DCV, respectively, under the parameters shown in Table 5.3.

Table 5.5: Sensitivity of Standard Deviation of PNW Basis

Standard Deviation	Description	E(π)	σ	h_A	g
-10%	Lower Standard Deviation	24.03	8.49	39%	19%
	Base Case	**23.13**	**7.95**	**50%**	**22%**
+10%	Higher Standard Deviation	22.44	7.49	58%	24%

As the standard deviation of the destination basis is increased, the optimal coverage ratio also increases. The base case correlation of 0.372 is still used in this sensitivity so as the standard deviation of destination basis increases, the coverage ratio required for the rail also increases. While optimal rail coverage increases, the increase is not as large as the increase in the coverage needed in the PNW.

Table 5.6: Sensitivity of Standard Deviation of DCV

Standard Deviation	Description	E(π)	σ	h_A	g
-10%	Lower Standard Deviation	24.48	8.61	47%	0%
	Base Case	**23.13**	**7.95**	**50%**	**22%**
+10%	Higher Standard Deviation	22.41	7.49	51%	36%

As the standard deviation of rail prices is increased, the coverage taken in both rail and grain increases, but the increase is primarily in the rail market. The optimal coverage in grain increases without the standard deviation of PNW increasing because of the correlation between the basis market and rail market. The increased coverage results in lower profit and lower standard deviations of profit.

Table 5.7: Sensitivity of Restrictions on DCV Coverage under Base Phi (0.10)

	11/13/2020				12/24/2020				1/27/2021			
Scenario	High Grain Basis, Low Priced Rail				Low Grain Basis, High Priced Rail				Low Grain Basis, Low Priced Rail			
Jamestown T1 Basis	-60				-70				-75			
PNW T1 Basis	125				114.5				87			
DCV T1 Price	-100				650				150			
	E(π)	σ	h_A	g	E(π)	σ	h_A	g	E(π)	σ	h_A	g
Freight Coverage Forced to 0%	20.15	8.10	90%	0%	24.19	9.33	45%	0%	32.59	14.20	0%	0%
Base Case: Freight Coverage Varies 0% to 100%	34.11	0.00	100%	100%	**23.13**	**7.95**	**50%**	**22%**	40.64	15.30	0%	100%
Freight Coverage Allowed up to 200%	47.76	8.62	100%	200%	23.13	7.95	50%	22%	43.59	16.83	0%	137%

Table 5.7 shows the effects of varying the restrictions placed on the freight coverage ratios under base phi with different scenarios at various times. The time periods were selected to illustrate the impacts that Time 1 prices may have on the optimal coverage ratios. The values listed correspond to the noted date which represents Time 1 in this case.

The first row shows the date that was selected to show a certain situation. The second row describes the prices of grain and rail relative to the mean of the distribution, the expected prices for Time 2. The third through fifth rows show the input prices used at Time 1. The bottom portion of the table shows the freight coverage parameters, the corresponding profit and standard deviation shown in cents per bushel, and the optimal coverage ratios.

The table shows the base case presented in Table 5.2 in yellow highlighting. In this case, the current rail price is higher than the expected rail price at Time 2. With a low level of risk aversion, the optimal strategy would be to purchase none of the freight in advance. However, when the risk aversion coefficient is 0.10 as in the base case, the optimal coverage ratios are 50% for beans and 22% for rail. Using the same base case restrictions on 11/13/2020 where the grain basis is high and rail prices are low, the optimal rail and grain coverage ratios are 100%. In contrast, on 1/27/2021 when basis and rail prices were low, the optimal grain coverage ratio was 0% while the rail coverage was 100%.

A situation may be shown where freight coverage cannot be taken, but the firm is still exposed to freight risk. When freight is forced to zero on a position entered on 12/24/2020, the coverage taken in soybean basis is reduced to 45% from 50% when freight was allowed to be covered under the base case restrictions.

Alternatively, a case is presented where freight is allowed to be covered up to 200%. This would allow the trader to overbook their freight when the price is low and sell it in the secondary spot market when the price is higher. Under the base case inputs, allowing freight coverage to be between 0% and 200% makes no impact on the optimal coverage ratio because the rail is high priced relative to expected rail prices. In a situation with low-priced rail relative to the expected, the optimal rail coverage may exceed 100% such as in column one or three where optimal freight coverage is 200% or 137%, respectively. The differences between the position sizes are in part because of the difference in the starting DCV prices. On 11/13/2020, the DCV was -$100/car where on 1/27/2021 the DCV was $150/car. This demonstrates the importance of freight as a tradable asset within a portfolio.

5.2.3. Short-the-Basis Base Case Under E-V and E-SV

The prior section focused on a long-the-basis case where grain is bought at a given price and is either sold immediately or held with the expectation of increasing prices. This section uses the same methods to evaluate a case where the trader enters a short cash position by selling cash and seeking coverage by buying basis at the origin and freight. A short cash position is typically entered when basis is high and expected to get lower. The merchant makes a sale of grain for delivery at later date and may purchase the grain when the sale is made or just before delivery is required to the destination.

Table 5.8 shows the specifications used to evaluate Models H 1.1.1 and H 2.1.1. Table 5.9 presents the results of the estimations. The interpretation of the results is fairly straightforward; for a given level of risk aversion, denoted by phi, and a given model, an optimal coverage ratio exists for the mean-variance and mean-semivariance frameworks. For example, under E-V, Model H 2.1.1 with a risk aversion level of 0.20 has a profit of 23.42 cent/bu with a standard deviation of 4.72 cents/bu when 45% of the grain is purchased at the origin and 53% of the rail is purchased during Time 1.

Table 5.8: Short-the-Basis Base Case H Model Specifications

Specification	Input Value	Units
Time 1	12/24/2020	Month/Day/Year
Time 1 Daily Car Value	650	$/car
Time 1 Basis: Jamestown, ND	-70	c/bu
Time 2 Basis: PNW	114.5	c/bu
Model Number	Varies	
Risk Aversion Level	Varies	

Table 5.9: Alternate Case Short-the-Basis Historical BestFit™

			E-V Models				E-SV Models			
Phi	Model Name	Description	E(π)	σ	h_O	g	E(π)	σ	h_O	g
0.05	H 1.1.1	Short Basis No Freight Risk	21.45	4.96	0%		21.44	4.93	0%	
	H 2.1.1	Short Basis Short Freight	28.00	9.59	0%	0%	28.00	9.59	0%	0%
0.10	H 1.1.1	Short Basis No Freight Risk	21.45	4.96	0%		21.44	4.93	0%	
	H 2.1.1	Short Basis Short Freight	27.61	9.16	0%	6%	28.00	9.59	0%	0%
0.15	H 1.1.1	Short Basis No Freight Risk	20.64	3.35	33%		21.44	4.93	0%	
	H 2.1.1	Short Basis Short Freight	24.93	6.32	26%	37%	28.00	9.59	0%	0%
0.20	H 1.1.1	Short Basis No Freight Risk	20.23	2.53	49%		21.44	4.93	0%	
	H 2.1.1	Short Basis Short Freight	23.42	4.72	45%	53%	26.24	7.79	0%	27%
0.30	H 1.1.1	Short Basis No Freight Risk	19.81	1.69	66%		20.65	3.35	32%	
	H 2.1.1	Short Basis Short Freight	21.94	3.16	63%	69%	23.94	5.34	29%	51%

While there are many similarities between the setup of the long-the-base and short-the-basis models, some key differences exist. Higher levels of risk aversion are needed to begin taking coverage in Table 5.9 as compared to Table 5.2 which showed that the trader should begin taking coverage even at low levels of risk aversion. This difference can be explained by the distributions used within the models and the specific scenario in question. The standard deviation of the PNW basis distribution is 15.2 cents/bu while the standard deviation of the Jamestown basis distribution is only 4.9 cents/bu. With a lower standard deviation, there is less risk that needs to be mitigated, thus a greater level of risk aversion is required before taking coverage to become optimal.

Another difference between Table 5.2 and Table 5.9 is the order in which coverage is taken in Models H 2.1 and H 2.1.1 under the mean-semivariance framework. In the long-the-basis case, shown in Table 5.2, coverage is first taken in the PNW basis market, and as risk aversion increases, it becomes optimal to begin taking coverage in the rail market. Conversely, in the case shown in Table 5.9, coverage is first taken in the rail market and is later taken in the basis market as risk aversion increases. These differences are a result of the standard deviation and the skews of the distributions.

The sensitivities conducted on the long-the-basis models could also be evaluated on this alternate case, but those tests would yield similar results and would provide little additional value. In essence, as standard deviations of distributions increase, the coverage needed at a given level of risk aversion also increases, and as two markets become more correlated, the amount of coverage needed in either market diminishes.

This section showed results for models evaluated using historical BestFit™ distributions for both long-the-basis and short-the-basis merchandising scenarios as well as a variety of sensitivities demonstrating how changes in distributions, correlations, and initial values change the optimal solutions. The remainder of this chapter evaluates the same sets of models using time-series forecasts.

5.3. Time Series Forecasted Results

This section uses time series forecasted distributions to show anticipated strategies. Using time series allows for specific position exit dates to be specified. Under this approach, a price and standard deviation of that period's price are projected for each time period. The standard deviation of the projection increases or remains constant as time increases. As a result, it is expected that as the time to liquidation increases, the optimal coverage levels at Time 1 should also increase. This section presents the results similarly to the previous section beginning with the base case, long-the-basis, scenario followed by sensitivities. Next, an alternative case of a short-the-basis case is shown along with relevant sensitivities.

5.3.1. Long-the-Basis Base Case Under E-V and E-SV

Table 5.10 shows the base case specifications for a long-the-basis scenario evaluated with time series distributions. The results are presented in Table 5.11.

Table 5.10: Long-the-Basis Base Case TS Model Specifications

Specification	Input Value	Units
Time 1	12/24/2020	Month/Day/Year
Time 2 (Liquidation Date)	6/2/2021	Month/Day/Year
Time 1 Daily Car Value	650	$/car
Time 1 Basis: Jamestown, ND	-70	c/bu
Time 1 Basis: PNW	114.5	c/bu
Time 1 Basis: St. Louis, MO	38.25	c/bu
Model Number	Varies	
Risk Aversion Level	Varies	

Table 5.11: Base Case Long-the-Basis Time Series Forecasted

			E-V Models					E-SV Models				
Phi	Model Name	Description	E(π)	σ	h_A	g	h_B	E(π)	σ	h_A	g	h_B
0.05	TS 1.1	Long Basis No Freight Risk	17.86	4.58	72%			18.86	8.99	44%		
	TS 2.1	Long Basis Short Freight	23.60	11.45	61%	0%		24.77	14.10	32%	0%	
	TS 3.1	Long Basis Multiple Destinations Short Freight	25.12	11.84	39%	0%	0.0%	27.06	15.84	0%	0%	0%
0.10	TS 1.1	Long Basis No Freight Risk	17.33	2.25	86%			17.84	4.50	72%		
	TS 2.1	Long Basis Short Freight	**20.31**	**5.90**	**81%**	**48%**		23.54	11.34	63%	0%	
	TS 3.1	Long Basis Multiple Destinations Short Freight	21.83	7.11	67%	36%	0.0%	25.27	12.10	36%	0%	0%
0.15	TS 1.1	Long Basis No Freight Risk	17.16	1.53	91%			17.49	2.97	82%		
	TS 2.1	Long Basis Short Freight	19.17	3.96	87%	65%		22.76	10.11	74%	7%	
	TS 3.1	Long Basis Multiple Destinations Short Freight	20.12	4.69	78%	58%	0.0%	23.34	9.77	53%	14%	0%

Table 5.11 shows the different trials of each model evaluated at various levels of phi, the risk aversion coefficient. The yellow cells represent a base case that is maintained throughout Sections 5.3.1 and 5.3.2 using consistent inputs. When sensitivities are conducted, the base case is shown next to the results of the models with changed inputs.

The results show that for Model TS 2.1, profit is 23.6 c/bu, 20.31 c/bu, and 19.17 c/bu for risk aversion levels of 0.05, 0.10, and 0.15, respectively. The standard deviations of these profits are 11.45 c/bu, 5.90 c/bu, and 3.96 c/bu. The optimal coverage ratios range from 61% to 87% for grain and 0% to 65% for rail. This means that for the most risk-averse trader shown, 87% of the grain should be sold at the time the grain is purchased and 65% of the required rail cars should be purchased in December. A trader with low levels of risk aversion should purchase all of the freight in June as it was projected to be lower than the price in December.

The results exhibit similar patterns as those shown in Table 5.2 when evaluated using historical BestFit™ distributions. Profit and standard deviation of profit increase as the models include more tradable assets. This is a logical conclusion because as more markets are added, the trader has more opportunities to profit and is exposed to additional risk. Additionally, profit and standard deviation of profit decrease as risk aversion increase. Increasing risk aversion leads to a decrease in the standard deviation of profit. The E-V curve shown in Chapter 2 shows that as the standard deviation of profit decreases, profit also tends to decrease.

Although the naïve results using historical distributions and the anticipatory results using time series distributions exhibit some similarities, there are also some key differences in the results. The optimal coverage ratios shown in Table 5.11 with time series distributions are greater than those shown in Table 5.2 with historical distributions. Additionally, the profit levels and standard deviations of profit were lower when evaluated with time series distributions.

The distributions of the profit are shown in Figure 5.5 and Figure 5.6 which show how changing the risk aversion levels affect the profit levels and the standard deviations of profit under both mean-variance and mean-semivariance frameworks. Figure 5.7 follows which shows how each random variable contributes to the variance of the profits.

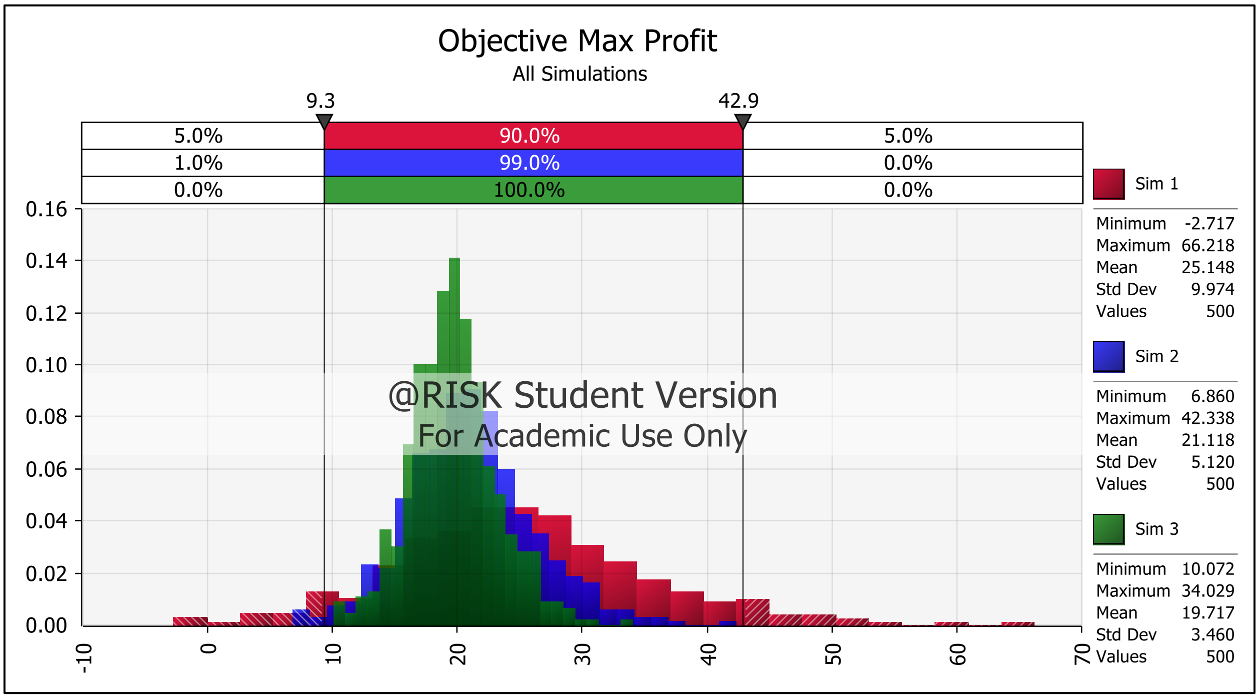

Figure 5.5: Profit Distributions from Model TS 2.1 at Various Levels of Phi Under E-V

Figure 5.5 shows the profit distributions from Model TS 2.1 under the E-V framework when evaluated at various levels of phi. Sim 1, Sim 2, and Sim 3 correspond to levels of phi of 0.05, 0.10, and 0.15, respectively. The results are consistent with those presented in Figure 5.1 which showed that as risk aversion increased, profits and standard deviations of profits also decreased.

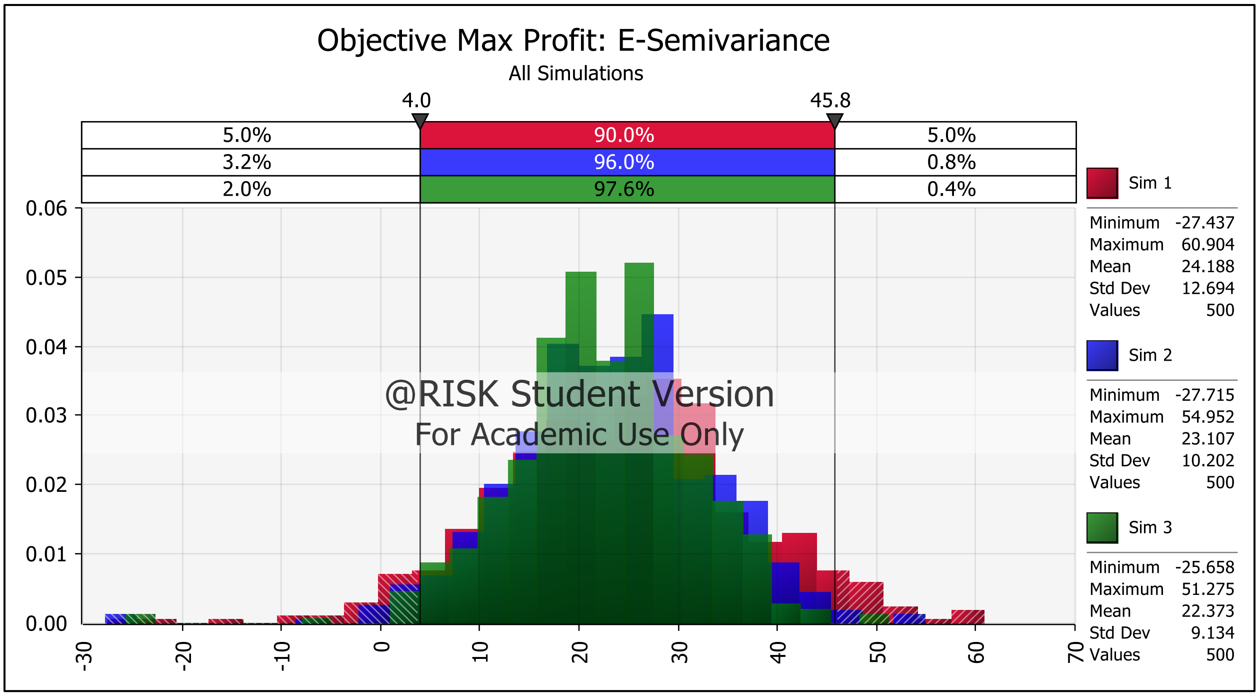

Figure 5.6: Profit Distributions from Model TS 2.1 at Various Levels of Phi Under E-SV

Figure 5.6 shows the profit distributions from Model TS 2.1 under the E-SV framework at various levels of phi. Similarly to Figure 5.5, Sim 1, Sim 2, and Sim 3 correspond to levels of phi of 0.05, 0.10, and 0.15, respectively. When mean-semivariance is used, a distinction is made between upside and downside risk. As a result, the optimal coverage ratios are lower which equates to larger profits and higher standard deviations in this scenario.

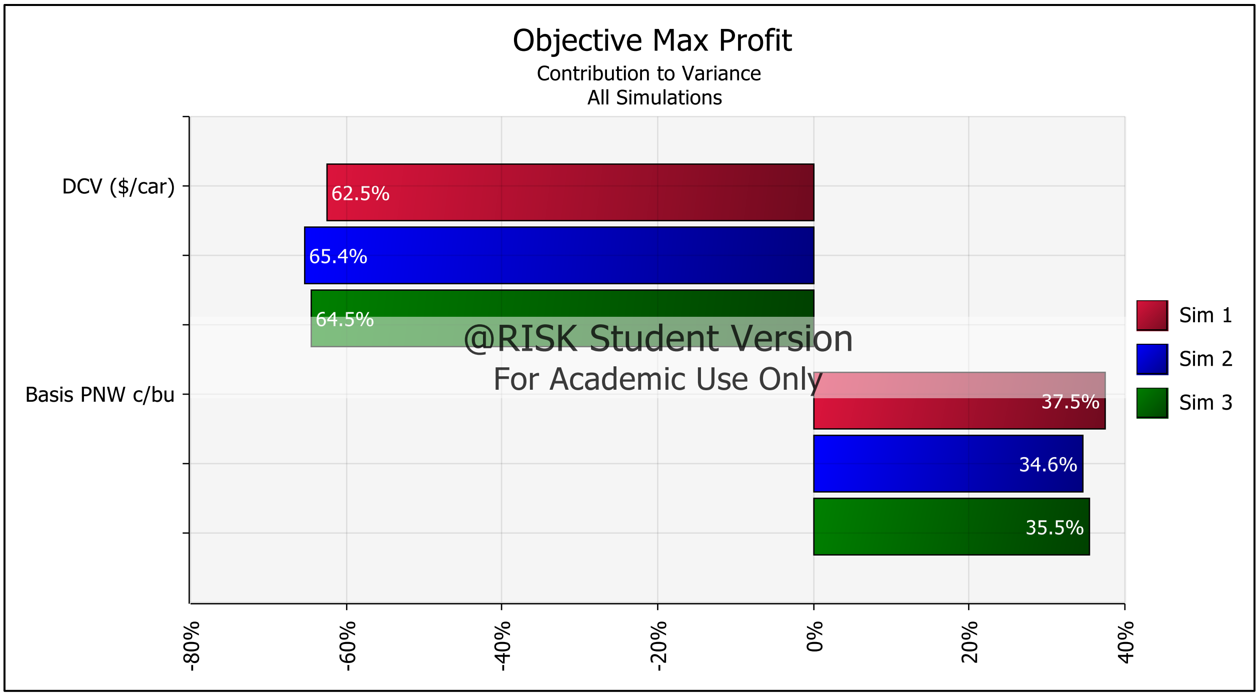

Figure 5.7: Contribution to Variance of Profit in Model TS 2.1 With Varying Phi Under E-V

The tornado graph shown in Figure 5.7 shows the contribution to the variance of the profit function at each level of risk aversion when evaluated on Model TS 2.1. Sim 1 has the least risk aversion with a phi level of 0.05 while Sim 3 has a risk aversion level of 0.15. The chart shows that when phi is 0.05, 37.5% of the variance in profit was from the PNW basis and 62.5% of the variance was from the changes in DCV.

5.3.2. Long-the-Basis Sensitivities

Several variables can impact optimal coverage ratios such as correlations between markets, time of entry for a position, or the standard deviation of a certain price. Some of these relationships between the variable and the resulting coverage ratios were shown in previous sections, but other relationships such as the timeframe for liquidation were not able to be shown because time series forecasting is required. The following tables show a variety of sensitivities that were conducted on Model TS 2.1 under the E-V framework using the specifications listed in Table 5.12 unless otherwise noted.

Table 5.12: Specifications for Long-the-Basis TS Model Sensitivities

Specification	Input Value	Units
Time 1	12/24/2020	Month/Day/Year
Time 2 (Liquidation Date)	6/2/2021	Month/Day/Year
Time 1 Daily Car Value	650	$/car
Time 1 Basis: Jamestown, ND	-70	c/bu
Time 1 Basis: PNW	114.5	c/bu
Time 1 Basis: St. Louis, MO	38.25	c/bu
Model Number	TS 2.1	
Risk Aversion Level	0.10	

Table 5.13 shows the effects of changing the correlations between the PNW basis market and the daily car values.

Table 5.13: Sensitivity of Correlation Between PNW Basis and DCV Using Time Series

Correlation	Description	E(π)	σ	h_A	g
0.0	Lower Correlation	21.14	6.57	86%	37%
0.372	**Base Case**	**20.31**	**5.90**	**81%**	**48%**
0.9	Higher Correlation	22.10	7.25	65%	26%

This sensitivity was motivated by the various correlations that were observed depending on the sampled time period (see Figure 4.7). The correlation between the PNW basis and DCV was 0.372 during the sample period but the correlation has been observed to vary between 0.0 to 0.9 over the previous two years in similar lengths of time. The results suggest that as the correlation increases, the optimal coverage ratios for both markets decrease. This is consistent with the results when the models were evaluated using historical BestFit™ distributions.

The next sensitivity shows the effects on the timeframe for liquidation; the results of this test are presented in Table 5.14. The hypothesis for this sensitivity was that as the timeframe for liquidation increased, the standard deviation of the forecasts would also increase causing the coverage ratios to increase since the model only allows for sales to be made at Time 1 or Time 2. This time series forecast does not exhibit a large increase in the standard deviation over time, which can be noted in Figure 4.13. Additionally, the projected PNW basis increased throughout time so even with an increased variance in the projection, the higher prices outweighed the increased risks. The projected rail prices were nearly constant over time with an increasing standard deviation. The results show that when evaluated under a 0.10 level of risk aversion, coverage was in grain was 95%, 81%, and 78% for grain in April, June, and August liquidation, respectively. While the optimal coverage in grain decreased, the rail coverage increased over time from 29% in April, 48% in June, to 49% in August. The results suggest that as the basis is held constant and the standard deviation of the projected basis increases, the optimal coverage ratio also increases.

Next, the impact of changing Time 1 parameters was evaluated. Table 5.15 shows how different market scenarios may affect the optimal coverage ratios. The liquidation date is held at June 2, 2021. While the market scenarios have a date listed for when the prices occurred, the standard deviations of projected prices are not affected as it was in the previous sensitivity. The results show that in scenarios when the grain is priced higher than expected and rail is lower than expected, the coverage ratio is 100% at all levels of risk aversion. When grain is priced low but rail is priced high, the coverage ratios for grain vary between 61% and 87% and for rail between 0% and 65% as risk aversion increases. In the last column, when grain and rail are priced lower than the expected price later, the optimal rail coverage ratio is 100% at all levels of risk aversion and as risk aversion increases, it becomes optimal to start taking coverage in the grain market.

These results show that the relationship between the current basis and the expected basis has a significant impact on the optimal coverage ratios.

Table 5.14: Sensitivity of Time to Liquidation Using Time Series

	4/7/2021				6/2/2021				8/4/2021			
E(PNW Basis)	116.552				118.92				118.365			
E(DCV Price)	393.07				393.14				391.08			
Phi	E(π)	σ	h_A	g	E(π)	σ	h_A	g	E(π)	σ	h_A	g
0.05	23.41	9.60	92%	0%	23.60	11.45	61%	0%	23.69	11.57	56%	0%
0.10	21.53	6.86	95%	29%	**20.31**	**5.90**	**81%**	**48%**	20.35	5.95	78%	49%
0.15	19.98	4.61	96%	52%	19.17	3.96	87%	65%	19.18	3.98	85%	66%

Table 5.15: Sensitivity of Time 1 Specification Using Time Series

	11/13/2020				12/24/2020				1/27/2021			
Scenario	High Grain Basis, Low Priced Rail				Low Grain Basis, High Priced Rail				Low Grain Basis, Low Priced Rail			
Jamestown T1 Basis	-60				-70				-75			
PNW T1 Basis	125				114.5				87			
DCV T1 Price	-100				650				150			
Phi	E(π)	σ	h_A	g	E(π)	σ	h_A	g	E(π)	σ	h_A	g
0.05	34.11	0.00	100%	100%	23.60	11.45	61%	0%	37.02	15.86	0%	100%
0.10	34.11	0.00	100%	100%	**20.31**	**5.90**	**81%**	**48%**	37.02	15.86	0%	100%
0.15	34.11	0.00	100%	100%	19.17	3.96	87%	65%	31.82	13.24	17%	100%

Table 5.16: Sensitivity of Restrictions on DCV under Base Phi (0.10) Using Time Series

	11/13/2020				12/24/2020				1/27/2021			
Scenario	High Grain Basis, Low Priced Rail				Low Grain Basis, High Priced Rail				Low Grain Basis, Low Priced Rail			
Jamestown T1 Basis	-60				-70				-75			
PNW T1 Basis	125				114.5				87			
DCV T1 Price	-100				650				150			
	E(π)	σ	h_A	g	E(π)	σ	h_A	g	E(π)	σ	h_A	g
Freight Coverage Forced to 0%	22.54	10.33	100%	0%	22.96	10.57	77%	0%	31.05	17.96	0%	0%
Base Case: Freight Coverage Varies 0% to 100%	34.11	0.00	100%	100%	**20.31**	**5.90**	**81%**	**48%**	37.02	15.86	0%	100%
Freight Coverage Allowed up to 200%	45.67	10.33	100%	200%	20.31	5.90	81%	48%	39.34	16.79	0%	139%

Finally, Table 5.16 presents the results from restricting rail car coverage when evaluated with time series distributions. The results in the first and third result columns were very similar between the historical and time series distributions while the middle column results varied slightly. The overall conclusions were consistent between the historical and time series distributions. The optimal solution is to sell all grain and purchase as much rail is allowed when rail rates are low and grain prices are high, relative to expected prices. When grain and rail are priced low, grain coverage is minimal under the base level of risk aversion, and rail coverage is high but not as high as when the grain is high and rail is low. When grain is low and rail is high relative to expectations, as in the middle column, the solution would be 0% for all of the scenarios in the absence of risk aversion.

The various sensitivities showed how changing variables such as correlations between markets, time to liquidation, and the price at the time of entry compared to expected prices may affect the overall coverage ratios, profit level, and standard deviation of profit in a long-the-basis case.

5.3.3. Short-the-Basis Base Case Under E-V and E-SV

This section builds from the short-the-basis results presented in Section 5.2.3. In the short-the-basis case, the merchant expects basis levels to fall so a sale is made at a destination before all of the grain is purchased at the destination. The freight can be purchased at the time the sale is made or at the time it is needed to ship the grain to the destination. The short-the-basis models are only evaluated for one origin and one destination.

The specifications are shown in Table 5.17, and the results are presented in Table 5.18.

Table 5.17: Short-the-Basis Base Case TS Model Specifications

Specification	Input Value	Units
Time 1	12/24/2020	Month/Day/Year
Time 2 (Liquidation Date)	6/2/2021	Month/Day/Year
Time 1 Daily Car Value	650	$/car
Time 1 Basis: Jamestown, ND	-70	c/bu
Time 2 Basis: PNW	114.5	c/bu
Model Number	Varies	
Risk Aversion Level	Varies	

Table 5.18: Alternate Case Short-the-Basis Using Time Series

			E-V Models				E-SV Models			
Phi	Model Name	Description	E(π)	σ	h_O	g	E(π)	σ	h_O	g
0.05	TS 1.1	Short Basis No Freight Risk	18.54	5.82	0%		18.55	5.63	0%	
	TS 2.1	Short Basis Short Freight	25.93	12.44	0%	0%	25.94	12.00	0%	0%
0.10	TS 1.1	Short Basis No Freight Risk	17.70	2.97	49%		18.55	5.63	0%	
	TS 2.1	Short Basis Short Freight	**22.39**	**7.48**	**22%**	**48%**	25.94	12.00	0%	0%
0.15	TS 1.1	Short Basis No Freight Risk	17.40	1.98	66%		18.07	4.08	28%	
	TS 2.1	Short Basis Short Freight	20.53	4.98	48%	66%	24.36	9.84	0%	27%

Table 5.18 shows the base case results for a short-the-basis scenario with Models 1.1.1 and 2.1.1 when using time series forecasted distributions. For each level of risk aversion, the profit and standard deviation of profit are reported. As risk aversion increases, the trader is willing to accept lower levels of profit in exchange for lower levels of risk. The base case that is used for sensitivities is Model TS 2.1.1 evaluated using the E-V framework at a risk aversion level of 0.10. The profit from that scenario is 22.39 c/bu with a standard deviation of 7.48 c/bu when the coverage ratios are 22% grain coverage at the origin and 48% of freight purchased at Time 1.

5.3.4. Short-the-Basis Sensitivities

The purpose of this section is to compare and contrast the impacts that certain variables have on coverage ratios in a short-the-basis scenario versus the long-the-basis scenario. While each of the sensitivities shown previously could be evaluated on a short-the-basis case, the results would be mostly redundant. Two sensitivities are shown that illustrate the effects of varying the timeframe for liquidation and Time 1 specification. The models were evaluated using the parameters shown in Table 5.17 unless otherwise noted.

The effects of the timeframe for liquidation on coverage ratios are shown in Table 5.19. In this scenario, the projected buying basis at the origin is increasing slightly as time for liquidation increases. The standard deviation of the projected distribution for the origin basis increases more throughout time than the standard deviation for the rail distribution. This can be noted in Figure 4.15 and Figure 4.16 in Chapter 4. In this case, as the time to liquidation is increased the optimal coverage for grain at the origin increases while rail decreases. The decreases in rail coverage are not as significant as the increases in grain coverage at the origin. The results show that when the price is held nearly constant, the coverage required increases with an increasing standard deviation. This conclusion is consistent with the long-the-basis version of this sensitivity.

Lastly, the results for the sensitivity of the Time 1 specifications in the short-the-basis case are shown in Table 5.20. While the dates that are used are the same as in Table 5.15, the results shown in Table 5.20 differ because the trader expects basis levels to be lower by the time the grain must be purchased at the origin. In the first column, grain is priced higher than expected and rail is priced lower than the expected price in June. As a result, in the short-the-basis case, the optimal coverage ratio at the origin is 0% for all tested levels of risk aversion where it was 100% at each tested risk aversion level in the long-the-basis case. In the last column, it is best to have a large amount of coverage at the origin when the prices are low relative to the expected where in the long-the-basis case, it was best to have a low amount of coverage. This example shows that the optimal solution varies significantly depending on whether the trader is in a long or short basis position.

Table 5.19: Sensitivity of Time to Liquidation Short-the-Basis Using Time Series

	4/7/2021				6/2/2021				8/4/2021			
E(Origin Basis)	-73.899				-73.418				-73.108			
E(DCV Price)	385.43				385.51				383.52			
Phi	E(π)	σ	h_O	g	E(π)	σ	h_O	g	E(π)	σ	h_O	g
0.05	24.72	10.61	0%	22%	25.93	12.44	0%	0%	25.60	11.86	0%	0%
0.10	22.57	6.95	0%	64%	**22.39**	**7.478**	**22%**	**48%**	22.32	7.43	33%	40%
0.15	21.48	5.58	11%	77%	20.527	4.979	48%	66%	20.50	4.96	55%	60%

Table 5.20: Sensitivity of Time 1 Specification Short-the-Basis Using Time Series

	10/22/2020				12/24/2020				1/27/2021			
Prices Relative to Forecasted Prices	Higher Origin Basis, Lower Priced Rail				Lower Origin Basis, Higher Priced Rail				Lower Origin Basis, Lower Priced Rail			
Jamestown T1 Basis	-65				-70				-75			
PNW T1 Basis	145				114.5				87			
DCV T1 Price	400				650				150			
Phi	E(π)	σ	h_O	g	E(π)	σ	h_O	g	E(π)	σ	h_O	g
0.05	56.09	6.25	0%	96%	25.93	12.44	0%	0%	5.51	0.00	100%	100%
0.10	56.07	6.22	0%	99%	**22.39**	**7.478**	**22%**	**48%**	5.51	0.00	100%	100%
0.15	56.07	6.22	0%	100%	20.527	4.979	48%	66%	5.51	0.00	100%	100%

5.4. Summary

This chapter showed the results from five models evaluated under mean-variance and mean-semivariance frameworks using historical BestFit™ distributions and time series forecasted distributions. Several sensitivities were conducted using each type of distribution to show the impacts that certain variables can have on the optimal coverage ratio.

The findings of the study are summarized below:

- Utility maximization theory applies to this research topic. Figure 5.3: E-V Curve from Model 1.1 shows that as coverage is reduced, profit and variance of profit increases.

- As risk aversion increases, standard deviations of profit and profits decrease. This is observed under both historical BestFit™ and time series forecasted distributions as shown in Table 5.2 and Table 5.10. Table 5.2 shows that the optimal coverage for grain and rail in Model H 2.1 increases from 12% in grain and 0% in rail to 67% in grain and 48% in rail as risk aversion increased from 0.05 to 0.15 when evaluated under mean-variance. Similarly, when evaluated with time series on Model TS 2.1, optimal coverage for grain and rail increased from 61% grain coverage and 0% rail coverage to 87% grain coverage and 65% rail coverage when risk aversion was increased from 0.05 to 0.15.

- E-V models limit upside and downside risk, while E-SV models limit downside profit-risk which results in lower optimal coverage ratios. Table 5.2 and Table 5.10 illustrate this as the optimal coverage levels are lower for every model and at each level of risk aversion when evaluated under E-SV instead of E-V. For example, Table 5.2 shows that when Model H 2.1 is evaluated at a risk aversion level of 0.10 and under both E-V and E-SV, expected profit in E-SV is 27.09 c/bu with a standard deviation of 13.05 c/bu while the expected profit in E-V is 23.13 c/bu with a standard deviation of 7.95 c/bu. Across the results, higher profits and higher standard deviations of profits are observed in E-SV models, but it is expected that those deviations will be mostly upside risk due to skews in the distributions of random variables.

- Correlations between markets vary significantly over time, and these correlations can strongly influence the optimal coverage ratios. Table 5.4 shows the correlation between the PNW basis and daily car values. The base case has a correlation of 0.372 with optimal coverage ratios of 50% and 22% for grain and rail. When correlation is increased to 0.9, optimal coverage ratios are reduced to 19% and 0% for grain and rail.

- As the standard deviations of a market increases, optimal coverage ratios increase. The impacts standard deviations have on coverage ratios are illustrated in Table 5.5

and Table 5.6. When the standard deviation of the PNW basis increased 10% the optimal coverage ratio increased 8 percentage points for grain and 2 points for rail.

- When using time series distributions that have an increasing standard deviation throughout time, more coverage is needed as the time until liquidation increases. Table 5.19 shows that as a position liquidation date moves from April to June to August, the coverage levels increase. At an April liquidation, optimal coverage under base phi is 0% and 64% for grain and rail where when deferred to August, optimal coverage increases to 33% and 40% for grain and rail.

The results show how a variety of factors contribute to the margin risk of a shipper elevator. The implications of these results are significant for academics as well as industry professionals. The findings contribute significantly to the academic literature which has limited research involving basis and freight coverage ratios. Additionally, the results demonstrate the impacts intermarket correlations and individual standard deviations can have on the elevator’s risk. In the industry, this approach could be further developed to be used in risk management groups to find optimal positions for the firm, elevators, or individual traders.

CHAPTER 6. CONCLUSIONS

6.1. Introduction

Commodity trading involves risks due to changes in futures prices, basis, and transportation costs. Risk is often quantified using the standard deviation or variance of the related price distributions. Using an E-V curve, Markowitz (1952) illustrated that with greater expected returns, there is also a larger variance of the returns. This balance between risk and reward leaves merchants, managers, and board members to decide whether they would prefer higher returns with more risk or lower returns with more certainty. They must also consider the ownership structure and the desires of the shareholders regarding risk appetite. Taking the appropriate levels of risk is essential to stay competitive while also staying in business.

Along with establishing risk limits, decision-makers must also determine which types of risks are acceptable for the company. For example, some trading companies may trade futures, basis, and rail while others may only trade basis and rail, and others may only trade futures. Lorton and White (2010) suggest that trading basis can provide more consistent profits than actively trading futures. Changes in the futures market are often due to macroeconomic supply and demand factors. These may include droughts affecting major production areas, changes in export or usage expectations, or changes in currency markets.

To mitigate the risk from changes in the futures market price, elevators can hedge their physical grain ownership conventionally by taking an opposite position in the futures market as their cash position. The elevator may eliminate their futures risk with a hedge ratio of one where a hedge ratio is defined as the $\frac{Quantity\ of\ Bushels\ Hedged\ in\ Futures}{Quantity\ of\ Physical\ Bushels\ Owned}$. By hedging all of the physical bushels, the elevator's risk is related to changes in basis and transportation prices.

Basis is defined as the difference between the cash price at the location and the futures price. The basis is tied to a delivery period, location, and futures contract. For example, soybeans delivered in December to an elevator in Jamestown may be purchased with January futures at $13.00/bu and the basis at $1.00 under the January futures. The net purchase price is then $12.00/bu. Every location has a unique basis market due to localized supply and demand factors but is also related to other origins and destination markets. Additionally, the basis is expected to be the highest in the period before harvest when local stocks are depleted while during harvest basis typically declines as local stocks are plentiful. In a way, basis is a tool that controls the flow of grain across geographies and time. Furthermore, the temporal variability of shipping costs impact the inter-market basis relationships over time.

A merchant that consistently hedges all of the elevator's physical grain is often referred to as a basis trader. In many situations, traders are purely basis traders in the sense that they make purchases in reference to the basis and make sales in reference to a basis. Hence, a complete transaction may be consummated purely in reference to the buying and selling basis.

As a basis trader, merchants can make strategic purchases and sales to take advantage of basis changes. If the basis is expected to increase, the merchant may enter a long basis position by purchasing grain at the origin and sell the grain to a destination after the basis has appreciated. Conversely, if the basis is expected to decrease, an advance sale may be made to a destination market for deferred delivery, and the grain may be purchased at the origin just prior to making the delivery. An alternative is a back-to-back sale in which grain is purchased at the origin and sold to the destination at the same time (i.e., within the same day). This secures a guaranteed basis margin for the merchant but does not allow for the opportunity to make money from changing basis.

Because basis trading often requires the elevator to ship grain from its location to a destination market, freight prices become a source of risk to the elevator’s margin. When grain is acquired at the origin or sold at the destination for a certain delivery window, the elevator becomes short freight. This means that the elevator margins benefit from decreasing freight prices and suffer from increasing prices. The rights to use rail cars can be purchased in the primary market. This reserves the use of the rail rights over a set amount of time at a given price. An alternative for purchasing rail car usage is the secondary market. This market uses auctions to determine the rail car value for one trip at a given time. Merchandisers can use this for purchasing additional cars or selling excess cars. The primary market pricing is relatively stable where the daily car value can exhibit large variations in prices. Thus, if the merchant is highly risk-averse, purchasing freight in advance may be an optimal option for them.

6.2. Problem Statement

Many merchants working at/for shippers or elevators, are pure basis traders. This may be a result of centralized futures trading at a corporate office or because of a policy to hedge all futures risk. With futures risk eliminated, the attention is shifted to risks resulting from basis movements and shipping prices. While low-risk trading approaches such as back-to-back trades with pre-purchased freight can allow merchants to secure known margins, these approaches can limit merchandising opportunities and fail to recognize shareholders' desires to earn profits. Taking positions in the basis or freight market can yield higher returns but also increases the firm’s risk exposure. This introduces the problem of finding a trading strategy with a balance of risk-bearing and risk-free trades.

Coverage ratios are introduced as a step to solving this problem. A coverage ratio can be defined as the portion of grain sold to the destination relative to the grain owned at the origin; the quantity of grain bought relative to the quantity of grain needed to be purchased; or for rail, the quantity of rail purchased relative to the amount of rail needed. For an elevator holding grain at a time, Time 1, which must be sold later, Time 2, this can be represented as the $\frac{amount\ sold\ to\ the\ destination\ at\ T_1}{amount\ owned\ or\ purchased\ at\ origin\ at\ T_1}$ where T_1 represents Time 1. In this example, the coverage ratio indicates the portion of the portfolio exposed to basis risk. If all of the grain which is owned at Time 1 is also sold at Time 1, as in a back-to-back trade, the coverage ratio is 1.0. If the elevator owns grain at the origin and the coverage ratio is 0.0, the elevator profit margins decrease as if the basis decreases. Each of the coverage ratios used throughout the analysis are presented in Chapter 3.

An optimal coverage can be found which incorporates the expected returns, the variance of the asset, and the risk aversion level of the firm. Additionally, a coverage ratio can be optimized for just one location holding grain or in a case with a grain position, a freight position, and multiple destination possibilities. Determining the optimal coverage ratios could be useful for merchants, risk managers, and future academic research topics.

6.3. Theoretical Conclusions

Chapter 3 derived each of the coverage ratios used throughout the thesis as well as the payoff functions used to represent a variety of common merchandising scenarios. Two sets of payoff functions are defined: one for long-basis positions where an elevator has physical ownership at the origin and one for short-basis positions where the elevator has made forward sales to the destination for delivery at a later time. In the long-the-basis models, profit is a function of the grain sold at a known price at Time 1, the remaining grain sold at a random price at Time 2, and the costs of freight and acquiring grain at the origin. In models that included freight risk, freight could be purchased at a known price at Time 1 or at a random price at Time

2. The short-the-basis models defined profit as the revenues from the sale of grain to a destination at Time 1 for delivery at Time 2 minus the cost of acquiring grain at a known price at Time 1 and the cost of acquiring the remaining grain at Time 2 as well as transportation costs. In the long-the-basis models, randomness in profit is due to randomness in the destination basis where in the short-the-basis models, randomness in profit is due to randomness in the origin basis.

An analytical solution is presented which is adapted from Blank, Carter, and Schmiesing (1991). Their model used the mean-variance framework to optimize hedging ratios given levels of risk aversion as well as a market bias. This type of derivation uses calculus to maximize the utility function rather than using simulations and stochastic optimization. The Blank, Carter, and Schmiesing (1991) model was the foundation for optimizing coverage ratios in this thesis.

Two models were used to illustrate the process for analytically deriving the optimal coverage ratios. Model 1.1 represented a basis trader with a long basis position and no freight risk. Model 2.1 represented a basis trader with a long basis position and a short freight position. Using the respective payoff functions, the expected profit is defined as well as the variance of the expected profit. Next, the expected profit and variance of profit are substituted into the utility function. Taking partial derivatives yields the utility-maximizing coverage ratios for grain in Model 1.1 and grain and rail in Model 2.1.

The analytical solutions indicate that as risk aversion and the variance of returns increase, the optimal coverage ratio also increases. Additionally, as the expected gains from holding grain from Time 1 to Time 2 increase, the optimal coverage ratio will decrease, all else equal. The single variable case is used to illustrate the analytical process and is advantageous because of the ease of interpretations.

In the multiple variable case, the optimal coverage ratio for both grain and rail is dependent on the variance of both grain and rail as well as the correlation between the two markets. For example, the variance of freight and the relationship between the current freight price and the expected freight price influences the optimal grain coverage ratio, h. A similar pattern is observed when optimizing the rail coverage ratio.

6.4. Empirical Model and Results

Chapter 5 presented the results of the empirical analysis which optimized the coverage ratios using stochastic optimization with Monte Carlos simulations within RiskOptimizer™. Two types of distributions were used to evaluate the models. Historical BestFit™ distributions were used to represent a "naïve" merchandising strategy in which the merchant believed that the historical basis behavior was indicative of the average and standard deviation of future basis behavior. Timeseries distributions represented an "anticipatory" strategy in which the merchant would anticipate the basis and standard deviation for a specified number of periods forward. Each of the models and distributions were evaluated under mean-variance and mean-semivariance frameworks as well. This allowed for comparison between the two frameworks.

6.4.1. Empirical Results: Historical BestFit™ Models

The long-the-basis base case (Model H 2.1 evaluated under the mean-variance framework at phi of 0.10) has an optimal coverage ratio of 50% in grain and 22% in rail which results in a 23.1 c/bu profit with a 7.9 c/bu standard deviation of profit. When evaluating the same model under a lower level of risk aversion, 0.05 the optimal coverage ratios are 12% and 0% for grain and rail, respectively. This strategy yields a profit of 26.7 c/bu and a standard deviation of profit of 12.7 c/bu. At a higher level of risk aversion of 0.15, the optimal coverage is 67% for grain and 48% for rail.

When comparing Models H 1.1, H 2.1, and H 3.1, the results show that as more tradable assets are added to the portfolio, the profits are greater. Model H 1.1 only includes a grain position, Model H 2.1 includes a grain and freight position, and Model H 3.1 includes a freight position with two possible destination markets. At base phi under mean-variance, profit is 19.3 c/bu, 23.1 c/bu, and 25 c/bu for each respective model. Additionally, as more tradable assets are added, the optimal coverage ratios for each category of coverage decrease. As the coverage ratios decrease, the standard deviations of profits increase.

The relationship between risk and return is demonstrated by the E-V curve shown in Figure 6.1 which is formed when Model 1.1 is evaluated at varying coverage ratios. The figure shows that expected profits increase with the expected variance of profits. As the coverage at Destination A, h_A, approaches 100%, the basis risk approaches 0. The figure shows that at the lowest level of risk, the resulting profit is also lowest.

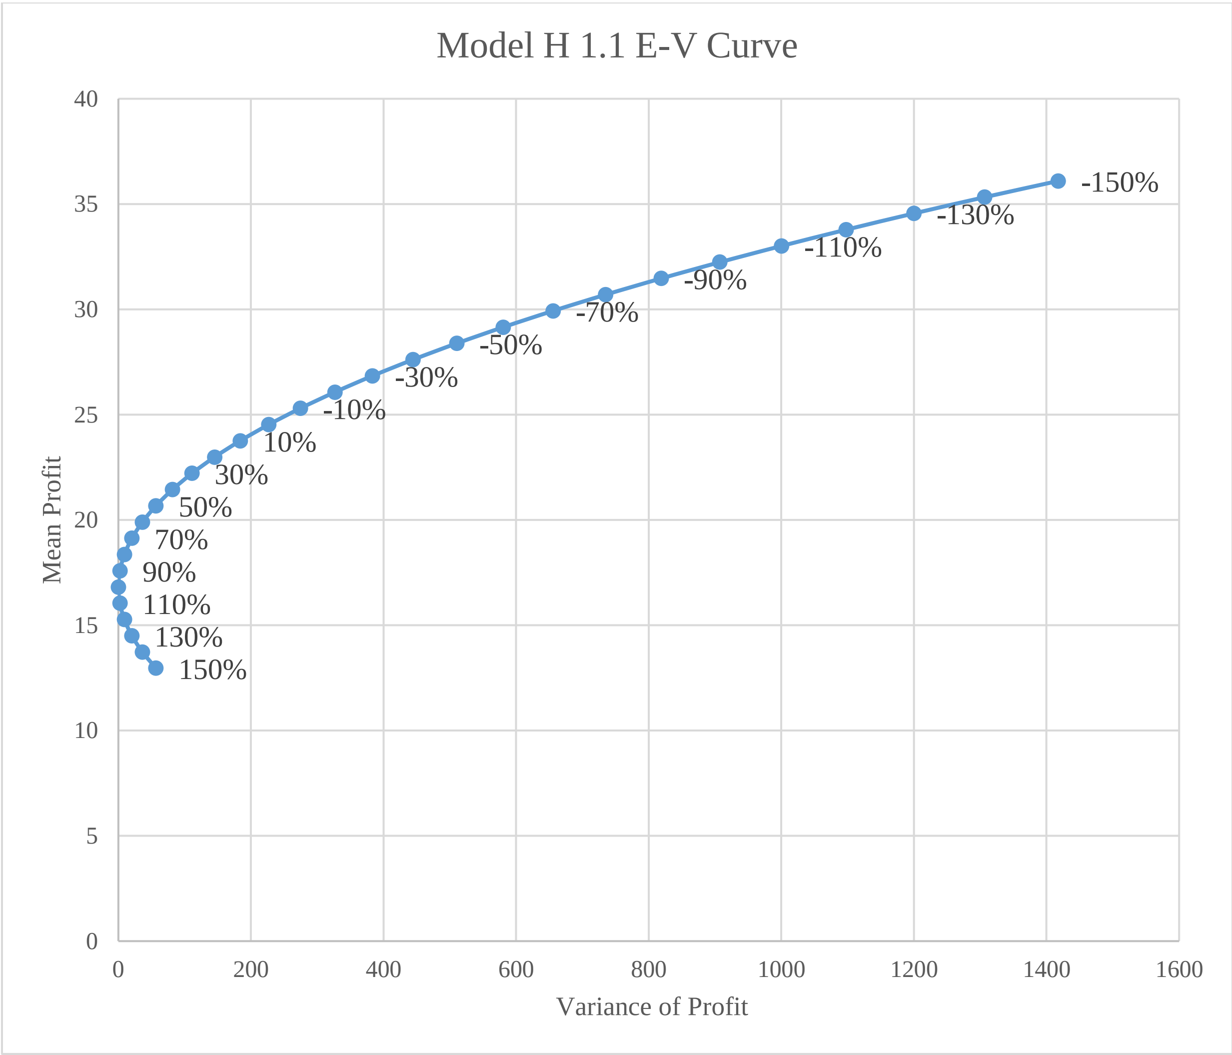

Figure 6.1: E-V Curve from Model 1.1

When the models were evaluated under the mean-semivariance framework, at all levels of risk aversion, the optimal coverage ratios were lower due to the skews of the data distributions and the distinction between upside and downside risk. Only taking coverage to avoid downside risk increases profit. Under mean-variance, Model H 2.1 had a profit of 23.1 c/bu with a standard deviation of 8.0 where the profit was 27.1 c/bu and 13.4 c/bu under mean-semivariance. When comparing all of the base case results, the models evaluated under mean-semivariance consistently had lower coverage, higher standard deviations, and higher profits.

Sensitivities were conducted which illustrated impacts of changing correlations and standard deviations. The base case has a correlation of .372 between the PNW basis and the DCV price. When the two variables varied independently, the expected profit was reduced by 3 c/bu from the base case with coverage in grain increasing 17% and rail increasing 41% from the base case. At higher correlations, profits are higher with lower coverage ratios in both grain and rail. The sensitives relating to standard deviation of basis and rail showed that as the standard deviation of the market increases, the coverage taken in both markets increases.

The short-the-basis models exhibited similar patterns as those shown in the long-the-basis models. In this case, partially due to a lower standard deviation in the origin basis, higher levels of risk aversion were required for any coverage to be optimal. When evaluating Model 2.1.1 with a risk aversion of 0.10, only 6% coverage was taken in rail. For the optimal coverage ratio to be over 50% for either model, a risk aversion level of 0.30 was required. A difference between the short basis model and the long basis model is the order in which coverage is taken. In the long-the-basis models, coverage is first taken in the PNW and then in the rail market where in the short-the-basis models, coverage is taken in the rail market first.

6.4.2. Empirical Results: Time Series Models

The long-the-basis base case (Model TS 2.1 evaluated under the mean-variance framework at phi of 0.10) has an optimal coverage ratio of 81% in grain and 48% in rail. The corresponding profit is 20.3 c/bu with a standard deviation of 5.9 c/bu. At a lower risk aversion level of 0.05, the optimal coverage ratio in grain is reduced to 61% and 0% in rail with a profit of 23.6 c/bu. When evaluated at a higher level of risk aversion, 0.15, the optimal coverage ratios increase to 87% in grain and 65% in rail. The profit decreases to 19.2 c/bu with a standard deviation of 4.0 c/bu. Similar to the historical BestFit™ distributions, as risk aversion increases, the optimal coverage levels increase. This reduces the standard deviations of profits but also reduces the expected profit. The time series models differed from the historical BestFit™ models in that the optimal coverage ratios were higher in the time series models which resulted in lower profits and standard deviations of profit. The results are sensitive to the distributions, so it is not surprising that using different distributions caused the results to differ.

Comparing Models TS 1.1, TS 2.1, and TS 3.1 shows that as the merchandiser has more trading opportunities available, the higher the expected profit. At each level of risk aversion and under both the mean-variance and mean-semivariance frameworks, the expected profit increases as additional merchandising opportunities are incorporated into the model.

The results show that when mean-semivariance is used instead of mean-variance, the optimal coverage ratios are consistently lower with higher profits at each level of risk aversion and for each model specification. Although this leads to higher standard deviations in profit, the deviations are positive deviations which are more desirable for a company than negative deviations from expected profit. This conclusion is consistent with the results that used historical BestFit™ distributions.

Sensitivities were performed to evaluate the impact of varying correlations, time to liquidation, Time 1 specification, and restrictions on DCV coverage on the optimal coverage ratios and profit levels. Similar to the sensitivity relating to correlations between the PNW basis and DCV which used historical BestFit™ distributions, the time series results showed that the optimal coverage levels were higher when the market had a low correlation and lower when the markets were closely correlated. When evaluating the impact of the timeframe for liquidation, the results contradicted the original hypothesis that the coverage would increase over time. Because the PNW basis was increasing throughout time more than the standard deviation, the results showed that it was optimal to take lower levels of coverage as the timeframe for liquidation was increased. Impacts of changing the Time 1 specification were also evaluated. A variety of market scenarios were tested which showed that when grain was priced relatively high at Time 1 compared to the expectation of Time 2 and rail was priced relatively low at Time 1 compared to the expected Time 2 price, the optimal coverage levels were 100% in both markets. When grain was low and rail was high, coverage ratios ranged from 61-87% in grain and 0-65% in rail. Lastly, a scenario was presented in which freight coverage was restricted to 0% or allowed up to 200%. Under base inputs, allowing freight to be increased to 200% did not affect the optimal coverage ratios. However, when evaluated under the 11/13/2020 conditions with high basis and low rail prices, the optimal freight coverage was 200% with grain coverage at 100%.

For the short-the-basis base case, coverage was optimal at lower levels of risk aversion than the historical BestFit™ distributions. Similar to the long-the-basis case, the results show that as risk aversion increases, the coverage levels increase. In the base case scenario (Model TS 2.1 evaluated under base phi of 0.10 and mean-variance framework), the profit is 22.4 c/bu and the standard deviation is 7.5 c/bu with 22% coverage in grain at the origin and 48% coverage in rail. The optimal coverage levels are still lower and with higher corresponding profits when evaluated under the mean-semivariance framework than the mean-variance framework.

Two sensitivities are conducted with time series in the short-the-basis series which evaluate again the time to liquidation and the Time 1 specification. The time to liquidation is interesting because the results differ from those observed in the long-the-basis time series sensitivity. In this case, the optimal coverage at the origin increases as the time to liquidation increases while the optimal coverage in rail decreases as the time to liquidation increases. This pattern can be explained by the projected movement in prices and the standard deviations of the prices. The origin basis distribution exhibits an increase in standard deviation and expected basis levels throughout time where rail exhibits a relatively stable standard deviation over time. The increased risk and expectations of increasing prices throughout time cause the optimal coverage ratios to increase. Lastly, the sensitivity relating to the Time 1 specification shows that when there are high origin basis levels at Time 1 relative to the expected origin basis at Time 2, the optimal coverage ratio is 0% in grain. In contrast, when the origin basis is low relative to the expected along with lower than expected rail, the optimal coverage ratios for rail and grain are 100%.

6.5. Implications of Results

The results of the empirical models have important implications for industry practitioners as well as researchers. First, the results show that the selection of the objective function is significant in determining optimal coverage ratios and has significant impacts on profit levels. Mean-variance models limit upside and downside risk, while mean-semivariance models limit downside profit-risk which causes them to generally find lower optimal coverage ratios. This yields higher profits and higher standard deviations of profits, but it is expected that those deviations will be mostly upside risk due to skews in the distributions of random variables.

Next, by evaluating the profit and corresponding variance of Model H 1.1 at various coverage ratios, an E-V curve was created which illustrates the applicability of E-V modeling and utility maximization to basis and rail trading portfolios. Additionally, the analytical derivations of optimal coverage ratios in Chapter 3 reinforce that adjusting the coverage ratios can lead to a point of maximum utility.

The theoretical and empirical results show that the relationship between the basis at Time 1 and the expected basis at Time 2 is critical to determining the optimal coverage ratio. If a large basis change is expected, the optimal coverage levels are lower than when a small basis change is expected. Additionally, the relationship between the basis, destination, and rail markets is important as taking coverage in one may reduce the risk of the other market.

Intermarket correlations vary significantly across time, and these correlations have significant impacts on the optimal strategy. In practice, many market participants and analysts assume that the correlations between basis and destination markets is equal to 1.0. Regardless of the distribution used to evaluate the effects of intermarket correlations, significant changes in coverage ratios were noted with changing correlations. As the correlations increased, the optimal coverage ratios decreased.

The sensitivities also show that when intermarket correlations exist, the coverage taken in one market helps to reduce the risk. This is demonstrated in the sensitivities relating to market correlations and the standard deviations of individual markets. First, when the correlations are reduced from the base case, the coverage level is higher for both rail and freight, but there is a greater increase in coverage in rail than in grain. Next, when the standard deviation of the PNW basis is adjusted with the DCV correlated at the base level of 0.372, a 10% change in the standard deviation results in about a 10-percentage point change in the optimal grain coverage ratio and 2-3 percentage point change in the rail coverage ratios. A similar theme is shown when adjustments are made to the standard deviation of rail car prices and the standard deviation of the PNW basis remains unchanged. This shows that when intermarket correlations exist, optimal coverage strategies in one market are affected by changing conditions of another market.

Overall, the analysis shows that an important element of a basis trading strategy is the portion of grain and rail coverage that is taken in a market once a market bias has been determined. Risk preferences of the decision-makers have a significant impact on the optimal strategies.

6.6. Limitations

This analysis is limited by many of the simplifying assumptions as well as the usage of data. First, many elevators handle multiple commodities and transport the grain using a combination of rail, truck, or barge. While these features could have been added, additional data, correlation matrixes, model specifications, and coverage ratios would be required. The value that would be gained by these additional features would be minimal relative to the additional complexity they would have added to the interpretations.

An additional limitation is the assumption regarding prices at Time 1 and Time 2, especially with the basis data. Basis is always related to a specific delivery period. Thus, equating the basis at Time 1 to the basis available for a forward contract during delivery at Time 2 is a bold assumption. While the method is not necessarily representative of reality, data showing forward basis quotes are not available in the form that would have been required for this analysis. Similarly, daily car value prices at Time 1 are not necessarily representative of the price of rail freight at Time 2. Additionally, purchasing rail cars far in advance does not usually happen using the secondary rail market.

Furthermore, the study is limited by assuming that there is only carry in the futures market and that cost of carry is equal to zero. Because basis may be used to control flows of grain across space and time, there may be times where the basis market pays holders of grain to store. In this type of scenario, the Time 2 basis may be higher than the Time 1 price. When the cost of carrying the grain to Time 2 is subtracted from the Time 2 basis, there may be little or no difference between selling at Time 1 or Time 2 which would affect the optimal coverage ratios.

The study is also limited in the stability of the distributions. This study used historical BestFit™ and time series forecasted distributions to optimize coverage ratios using simulations. However, these distributions are clearly not stable and as illustrated, have important impacts on the optimal coverage ratios. A limited length of history was able to be used to structural breaks in the data. Furthermore, for the historical BestFit™ distributions, in order to demonstrate the applications of a coverage ratio, the position entry dates needed to have a lower Time 1 basis at the destination than the average of the historical distribution for long-the-basis models.

Furthermore, the applicability of the utility functions and risk aversion levels to firms and decision-makers is another limitation. While a CARA utility function was assumed along with a variety of risk aversion levels, determining an individual's or firm's utility function and risk aversion level outside of theory is difficult.

6.7. Contribution to Literature

The conclusions from this research add to a plethora of literature relating to risk management and portfolio optimization. First, this work defines coverage ratios for grain and freight which can be optimized by maximizing the expected utility objective function. While there has been much research on basis and on optimal hedge ratios (i.e., hedging in futures), there have been few studies relating to risk management in basis and rail markets, and no study has sought to define and optimize coverage for these types of trading strategies.

Next, the data collected for this study revealed that the origin and destination are not always perfectly correlated as it seems to be assumed throughout the literature. Rather, this study showed that the correlations may vary significantly, and at times can be negatively correlated.

Additionally, the importance of the correlation between the secondary rail car market and the destination basis market was illustrated. The results showed that changing the correlations between the markets impacts the optimal coverage ratios.

The standard deviation of each market also has significant impacts on the coverage ratio and coverage ratios of correlated assets. For example, if for a long-basis position the standard deviation of the destination basis market is increased, the optimal coverage ratio for grain will increase along with the optimal coverage ratio for rail. Additionally, the data showed that the standard deviation of the origin basis is not necessarily the same as the standard deviation of the destination basis, and the standard deviations change throughout time.

6.8. Suggestions for Further Research

This research could be further developed by adding additional time periods to liquidate. The models assumed that a merchant could make sales or purchases at only Time 1 or Time 2 when in reality, there are far more than two times in which the merchant could make a purchase or sale. For example, if a merchant is long cash grain in December, Time 1, and must liquidate by June, Time 2, sales to the destination could happen at any point between Time 1 and Time 2. Expanding this could be useful for short-basis positions especially. Management may approve a short basis position at the origin six months prior to delivering to a destination but may want to see increasing coverage as the delivery period nears so that they are sure they have adequate stocks to fill their contract.

Additionally, the models could be expanded to include a portfolio of multiple origins, destinations, and modes of transportation. This type of analysis would be more representative of industry practice as it is common for a merchant to manage multiple elevator positions while typically having access to some combination of rail, barge, truck, or ship. These models could be used to represent a regional or company-wide basis coverage optimization.

These methods could also be applied to determine optimal coverage ratios for processors such as flour millers or soybean crushers. Instead of purchasing at an origin and selling to a destination market, these types of traders purchase a commodity from one market, transform it, and sell it in a separate but related market. In the case of a flour mill, the firm may decide to be short flour and need to purchase their wheat later or may choose to hold a long wheat position and sell the flour at a later time.

Further analysis regarding the stability of distributions and distribution selection is needed. The distributions are significant to the results, but structural breaks in data can limit the amount of data that is used. Additionally, a further analysis related to the correlations between the distributions could also prove valuable.

Lastly, this research could be enhanced by optimizing the hedge ratio with the coverage ratios for grain and transportation risk. While many elevators operate as pure basis traders, a better strategy may be to determine optimal strategies for hedging and coverage. Because of the correlations between futures, basis, and transportation, analyzing hedge ratios and coverage ratios together is likely a better strategy than analyzing them separately.

6.9. Summary

Risk management in commodity trading is important for long-term business success. While it is important to manage risk, many times risk management is not synonymous with risk elimination. To produce attractive returns, some amount of risk tolerance is necessary, but risk aversion varies across different people. This thesis focused on the risk associated with pure basis and rail trading portfolios. The concepts of the futures, basis, and rail markets were introduced along with the key market drivers. Understanding that futures hedge ratios can be optimized, coverage ratios were defined for long basis, short basis, and short freight positions. Using mean-variance an analytical solution for an optimal coverage ratio was derived.

The empirical models used daily car value prices and basis data for Jamestown, St. Louis, and the PNW and to illustrate a variety of merchandising scenarios. Distributions were fit to the data to create historical and timeseries distributions. The historical distributions represent a naïve strategy in which the merchant assumes future basis market activity to be similar to previous market activity where timeseries distributions represent an anticipatory strategy where the merchant anticipates future basis levels and standard deviations. Each of the models were evaluated under mean-variance and mean-semivariance frameworks and using both types of distributions at various risk aversion levels.

The results of this analysis showed the viability of optimizing coverage ratios for a grain and rail portfolio to improve risk management. The importance of the standard deviations of the markets, intermarket-correlations, and risk aversion levels of the decision-makers are demonstrated using sensitivity analysis.

REFERENCES

Alexander, Gordon J., and Alexandre M. Baptista. 2002. "Economic Implications of Using a Mean-VaR Model for Portfolio Selection: A Comparison with Mean-Variance Analysis." *Journal of Economic Dynamics and Control*, Finance, 26 (7): 1159–93. https://doi.org/10.1016/S0165-1889(01)00041-0.

Anderson, Ronald W., and Jean-Pierre Danthine. 1981. "Cross Hedging." *Journal of Political Economy* 89 (6): 1182–96.

Artzner, Philippe, Freddy Delbaen, Jean-Marc Eber, and David Heath. 1999. "Coherent Measures of Risk." *Mathematical Finance* 9 (3): 203–28. https://doi.org/10.1111/1467-9965.00068.

Awudu, Iddrisu, William Wilson, and Bruce Dahl. 2016. "Hedging Strategy for Ethanol Processing with Copula Distributions." *Energy Economics* 57 (June): 59–65. https://doi.org/10.1016/j.eneco.2016.04.011.

Baillie, Richard T., and Robert J. Myers. 1991. "Bivariate Garch Estimation of the Optimal Commodity Futures Hedge." *Journal of Applied Econometrics* 6 (2): 109–24. https://doi.org/10.1002/jae.3950060202.

Baldwin, E Dean. 1986. Review of *Understanding and Using Basis - Grains*, by John Ferris and David Holder. Edited by Edited Duane Griffith and Stephen Koontz. *Producer Marketing Management*, 10.

Benet, Bruce A. 1992. "Hedge Period Length and Ex-Ante Futures Hedging Effectiveness: The Case of Foreign-Exchange Risk Cross Hedges." *The Journal of Futures Markets (1986-1998)* 43 (1): 163.

Blank, Steven C., Colin Andre Carter, and Brian H. Schmiesing. 1991. *Futures and Options Markets: Trading in Commodities and Financials*. Vol. 1. Prentice Hall.

Bullock, David W., and William W. Wilson. 2019. "Factors Influencing the Gulf and Pacific Northwest (PNW) Soybean Export Basis: An Exploratory Statistical Analysis." AgEcon Search. May 16, 2019. https://doi.org/10.22004/ag.econ.288512.

Burlington Northern Santa Fe. 2021a. *Soybeans – PNW Soybeans Rates from BNSF Origins.* BNSF-4022-61115-M. Web.

———. 2021b. *Soybeans to St. Louis Destination Group*. BNSF-4022-69902-M. Web.

Cecchetti, Stephen G., Robert E. Cumby, and Stephen Figlewski. 1988. "Estimation of the Optimal Futures Hedge." *The Review of Economics and Statistics* 70 (4): 623–30. https://doi.org/10.2307/1935825.

Chang, Kuang-Liang. 2011. "The Optimal Value-at-Risk Hedging Strategy under Bivariate Regime Switching ARCH Framework." *Applied Economics* 43 (21): 2627–40. https://doi.org/10.1080/00036840903299771.

Chen, Sheng-Syan, Cheng-Few Lee, and Keshab Shrestha. 2001. “On a Mean—Generalized Semivariance Approach to Determining the Hedge Ratio.” *Journal of Futures Markets* 21 (6): 581–98. https://doi.org/10.1002/fut.1604.

Chen, Sheng-Syan, Cheng-few Lee, and Keshab Shrestha. 2003. “Futures Hedge Ratios: A Review.” *The Quarterly Review of Economics and Finance* 43 (3): 433–65. https://doi.org/10.1016/S1062-9769(02)00191-6.

Cheung, C. Sherman, Clarence C. Y. Kwan, and Patrick C. Y. Yip. 1990. “The Hedging Effectiveness of Options and Futures: A Mean-Gini Approach.” *Journal of Futures Markets* 10 (1): 61–73. https://doi.org/10.1002/fut.3990100106.

Chin-Wen, Hsin, Jerry Kuo, and Lee Cheng-Few. 1994. “A New Measure to Compare the Hedging Effectiveness of Foreign Currency Futures versus Options.” *The Journal of Futures Markets (1986-1998)* 14 (6): 685.

Chou, W.L., K.K.Fan Denis, and Cheng F. Lee. 1996. “Hedging with the Nikkei Index Futures: The Convential Model versus the Error Correction Model.” *The Quarterly Review of Economics and Finance* 36 (4): 495–505. https://doi.org/10.1016/S1062-9769(96)90048-4.

“Definition of a Futures Contract - CME Group.” 2021. Introduction to Futures. 2021. https://www.cmegroup.com/content/cmegroup/en/education/courses/introduction-to-futures/definition-of-a-futures-contract.html.

Ederington, Louis H. 1979. “The Hedging Performance of the New Futures Markets.” *The Journal of Finance* 34 (1): 157–70. https://doi.org/10.2307/2327150.

“Fundamentals and Agricultural Futures - CME Group.” 2021. Using Fundamental Analysis When Evaluating Trades. 2021. https://www.cmegroup.com/content/cmegroup/en/education/courses/using-fundamental-analysis-when-evaluating-trades/fundamentals-and-agricultural-futures.html.

Gao, J., and L. Liu. 2009. “Mean Conditional Value-at-Risk Model for Portfolio Optimization.” In *2009 International Conference on Business Intelligence and Financial Engineering*, 246–50. https://doi.org/10.1109/BIFE.2009.64.

Gelston, William, and Scott Greene. 1994. “Assessing the Potential for Improved Functioning of the Grain Merchandising/Transportation System.” *USDOT Pub. No. GRA-RRP-92-01. Final Report Prepared by Apogee Research, Inc., Submitted to the US Department of Transportation, Federal Railroad Administration, Washington DC*.

Ghosh, Asim. 1993. “Hedging with Stock Index Futures: Estimation and Forecasting with Error Correction Model.” *The Journal of Futures Markets (1986-1998)* 13 (7): 743.

Giamouridis, Daniel, and Ioannis D Vrontos. 2007. “Hedge Fund Portfolio Construction: A Comparison of Static and Dynamic Approaches,” 19.

Grammatikos, Theoharry, and Anthony Saunders. 1983. “Stability and the Hedging Performance of Foreign Currency Futures.” *The Journal of Futures Markets (Pre-1986)* 3 (3): 295.

Harris, Richard D. F., Jian Shen, and Evarist Stoja. 2010. “The Limits to Minimum-Variance Hedging.” *Journal of Business Finance & Accounting* 37 (5–6): 737–61. https://doi.org/10.1111/j.1468-5957.2009.02170.x.

Hogan, William W., and James M. Warren. 1974. “Toward the Development of an Equilibrium Capital-Market Model Based on Semivariance.” *The Journal of Financial and Quantitative Analysis* 9 (1): 1–11. https://doi.org/10.2307/2329964.

Howard, Charles T., and Louis J. D’Antonio. 1984. “A Risk-Return Measure of Hedging Effectiveness.” *The Journal of Financial and Quantitative Analysis* 19 (1): 101–12. https://doi.org/10.2307/2331004.

Johnson, L. L. 1960. “The Theory of Hedging and Speculation in Commodity Futures.” *The Review of Economic Studies* 27 (3): 139–51.

Jong, Abe De, Frans De Roon, and Chris Veld. 1997. “Out-of-Sample Hedging Effectiveness of Currency Futures for Alternative Models and Hedging Strategies.” *Journal of Futures Markets* 17 (7): 817–37. https://doi.org/10.1002/(SICI)1096-9934(199710)17:7<817::AID-FUT5>3.0.CO;2-Q.

Kahl, Kandice H. 1983. “Determination of the Recommended Hedging Ratio.” *American Journal of Agricultural Economics* 65 (3): 603–5. https://doi.org/10.2307/1240514.

Kimura, Norifumi. 2016. “Hedging Default and Price Risks in Commodity Trading.” Fargo, ND: North Dakota State University. https://library.ndsu.edu/ir/handle/10365/28055.

Kolb, Robert W., and James A. Overdahl. 2006. *Understanding Futures Markets*. 6th ed. Malden, MA: Blackwell Publishing.

Kroner, Kenneth F., and Jahangir Sultan. 1993. “Time-Varying Distributions and Dynamic Hedging with Foreign Currency Futures.” *The Journal of Financial and Quantitative Analysis* 28 (4): 535. https://doi.org/10.2307/2331164.

Kub, Elaine. 2014. *Mastering the Grain Markets: How Profits Are Really Made*. Omaha, Nebraska: Kub Asset Advisory, Inc.

Lakkakula, Prithviraj, and William Wilson. 2021. “Origin and Export Basis Interdependencies in Soybeans: A Panel Data Analysis.” *Journal of Agricultural and Resource Economics* 46 (1): 69–84. https://doi.org/10.22004/ag.econ.302464.

Lamont, Owen A, and Richard H Thaler. 2003. “Anomalies: The Law of One Price in Financial Markets.” *Journal of Economic Perspectives* 17 (4): 191–202. https://doi.org/10.1257/089533003772034952.

Landman, Daniel Jacob. 2017. "Real Option Analysis of Primary Rail Contracts in Grain Shipping." Fargo, ND: North Dakota State University.

"Left Skewed vs. Right Skewed Distributions." 2021. *Statology* (blog). January 13, 2021. https://www.statology.org/left-skewed-vs-right-skewed/.

Longerstaey, Jacques. 1996. *RiskMetrics Technical Document*. 4th ed. New York: J.P.Morgan/Reuters. https://www.bancoinvest.pt/docs/default-source/fundamentais-value-at-risk/riskmetrics.pdf.

Lorton, Sherry, and Don White. 2010. *The Art of Grain Merchandising*. Champaign, IL: Stipes Publishing Co.

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

Mausser, H., and D. Rosen. 1999. "Beyond VaR: From Measuring Risk to Managing Risk." In , 163–78. IEEE. https://doi.org/10.1109/CIFER.1999.771115.

Melnikov, Alexander, and Ivan Smirnov. 2012. "Dynamic Hedging of Conditional Value-at-Risk." *Insurance: Mathematics and Economics* 51 (1): 182–90. https://doi.org/10.1016/j.insmatheco.2012.03.011.

Nguyen, Linh, Linh Xuan Diep Nguyen, and Emmanuel Adegbite. 2018. "Does Mean-CVaR Outperform Mean-Variance? Theoretical and Practical Perspectives." SSRN Scholarly Paper ID 3143827. Rochester, NY: Social Science Research Network. https://doi.org/10.2139/ssrn.3143827.

Olson, Frayne. 2014. "Effects of 2013/14 Rail Transportation Problems on North Dakota Farm Income." Fargo, ND: North Dakota State University.

Ortiz, Laura. 2016. "No Train No Grain: The Impact of Increased Demand for Rail Services by The Energy Sector on Wheat Prices—A Preliminary Analysis." *International Journal of Food and Agricultural Economics* 4 (3): 103–25.

"Rail Service Challenges in the Upper Midwest: Implications for Agricultural Sectors - Preliminary Analysis of the 2013–2014 Situation." 2015. United States Department of Agriculture Office of the Chief Economist and the Agricultural Marketing Service. https://www.ams.usda.gov/sites/default/files/media/USDA%20Analysis%20of%202013-14%20Rail%20Service%20Challenges%20for%20Senators%20ThuneKlobuchar.pdf.

Rockafellar, R. Tyrrell, and Stanislav Uryasev. 2000. "Optimization of Conditional Value-at-Risk." *The Journal of Risk* 2 (3): 21–41. https://doi.org/10.21314/JOR.2000.038.

Shalit, Haim. 1995. "Mean-Gini Hedging in Futures Markets." *Journal of Futures Markets* 15 (6): 617–36.

Skadberg, Kristopher, William W. Wilson, Ryan Larsen, and Bruce Dahl. 2015. "Spatial Competition, Arbitrage, and Risk in U.S. Soybeans." *Journal of Agricultural and Resource Economics* 40 (3): 442–56.

Thomson Reuters. 2021. *PNW Soybean Prices from April 2018 to February 2021.* Retrieved from Thomson Reuters Eikon.

"Time Series Functions." 2021. Palisade Help Resources. April 1, 2021. https://help.palisade.com/v8_1/en/@RISK/Function/Time-Series-Functions.htm?Highlight=time%20series%20distributions%20.

TradeWest Brokerage. 2021. "Daily Market Report". *PNW Soybean Prices and Secondary Rail Car Market Values from April 2018 to February 2021.*

Turvey, Calum G., and Govindaray Nayak. 2003. "The Semivariance-Minimizing Hedge Ratio." *Journal of Agricultural and Resource Economics*, 100–115.

Usset, Edward. 2014. "Minnesota Basis Analysis: Final Report for the Minnesota Department of Agriculture." University of Minnesota.

Von Neumann, J, and O Morgenstern. 1947. *Theory of Games and Economic Behavior*. 2nd Revised Edition. Princeton University Press.

Vose, David. 2008. *Risk Analysis: A Quantitative Guide*. 3rd ed. West Sussex, England: John Wiley & Sons, Ltd.

Wilson, William W. 1982. "Hedging Effectiveness of U.S. Wheat Futures Markets." Agricultural Economics Report ; No. 165. Fargo, ND: Deptof Agricultural Economics, North Dakota Experiment Station, North Dakota State University.

Wilson, William W., David Bullock, and Prithviraj Lakkakula. 2020. "Dynamic Changes in Rail Shipping Mechanisms for Grain." AgEcon Search. February 28, 2020. https://doi.org/10.22004/ag.econ.302409.

Wilson, William W., and Bruce Dahl. 2011. "Grain Pricing and Transportation: Dynamics and Changes in Markets." *Agribusiness* 27 (4): 420–34. https://doi.org/10.1002/agr.20277.

Wilson, William W., Steven R. Priewe, and Bruce Dahl. 1998. "Forward Shipping Options for Grain by Rail: A Strategic Risk Analysis." *Journal of Agricultural and Resource Economics* 23 (2): 526–44.