Financial Innovation in Derivatives: Understanding the Use and Properties of Volatility and Dairy Derivatives

Ilnara Midhatovna Gafiatullina

A thesis submitted to Auckland University of Technology in fulfilment of the requirements for the degree of Doctor of Philosophy

Department of Finance

Faculty of Business, Economics and Law

November 2018

Abstract

Derivative products are undoubtedly among the most important financial innovations in modern history. Since the beginning of the 1970s, derivative markets have seen an expansion in the variety and complexity of products. There is a degree of consensus that derivatives bring significant benefits to the financial system. However, if not properly understood and applied, they can bear significant risks. Therefore, it is important to understand the use and properties of new derivative products. Among different uses, a large body of theoretical and empirical literature recognises derivatives for their informational role, contribution to price discovery, and risk management. This thesis aims to address these important functions with a focus on two recently introduced derivative products on volatility and dairy.

The first product class we investigate are volatility exchange-traded notes (ETNs), introduced in 2009. The most two successful volatility ETNs are the VXX and XIV. Broadly, the VXX allows taking a long position in the US stock market volatility, and the XIV a short position. Historically, the VXX delivers negative long-term returns and, despite that, attracts a tremendous trading volume. This observation might lead to the conclusion that trading in the VXX is dominated by noisy traders, and thus the VXX prices poorly reflect information about volatility. The XIV has a similar dynamic to the VXX but inverted. At times when the XIV increases, the VXX decreases. Informed traders can either switch between these two products, or stay in one market and take either a short or long position, depending on the information they have about future volatility. Understanding which product is more informative and at what times is an important question.

In the first essay, we investigate the informational leadership between the VXX and XIV using high-frequency data and attempt to answer what the key determinants of informational leadership are. Results show that the contribution to price discovery between these two competing markets is time-varying. This time variation is explained by three factors: trading costs, market liquidity, and market conditions. These findings contribute to a better understanding of these relatively new products. Time-variation in price leadership means that both ETNs can quickly react to new information about the future value of volatility and reflect it in their prices.

The second product class we investigate are dairy futures and options introduced by the New Zealand Stock Exchange (NZX) in 2010. Exports of dairy commodities play an important role for the New Zealand economy, dominated by exports in Whole Milk Powder (WMP). Currently, there are eight different dairy derivatives traded on the NZX. Futures and options contracts on WMP are amongst the most popular ones. WMP is one of the most volatile commodities globally. This high volatility affects the decision-making of all supply chain participants, including dairy processors and farmers. The high volatility of dairy commodities translates into a high milk price risk for dairy farmers, which adversely affects the financial strength of farmers and potentially stability of the New Zealand banking sector. Given the importance of the dairy sector to the New Zealand economy, in the next two essays we investigate the benefits these derivatives offer.

In the second essay, we investigate the information content of the dairy derivatives market. We develop a dairy-implied volatility index, termed the DVIX, from option prices on WMP futures. As for the properties of the DVIX, we document the asymmetric return-implied volatility effect. Additionally, we find that the DVIX contains significant information about future volatility, and outperforms the volatility forecast based on historical averages or the GARCH-type forecast. Overall, we find that the relatively new dairy derivatives market contains important information that can be used by market participants.

In the third essay, we evaluate WMP futures from a risk management perspective. We develop a profit margin hedging strategy which aims to protect New Zealand dairy farms from the downside risk of low liquid milk prices. We conduct the analysis both for a representative farm and for a sample of New Zealand farms. The

results show that the profit margin hedging strategy decreases the risk, reduces the likelihood of financial distress and improves the returns, even after controlling for commissions and different levels of basis risk. This study demonstrates that WMP futures are useful for a cross-hedging of milk price risk, despite the presence of the basis risk. Overall, our finding means that WMP futures can benefit New Zealand farms and improve the stability of the New Zealand banking sector, as high indebtedness of the dairy farm sector makes it vulnerable to low dairy prices.

All in all, our findings shed light on some properties and uses of relatively new derivative products, which were introduced in the last ten years - volatility and dairy derivatives. We document that both volatility ETNs can efficiently incorporate new information about the future value of volatility. We find that dairy options gained some level of informational efficiency, and information contained in their prices can be used to obtain insights about dynamics of the physical market. Additionally, we document that dairy futures are suitable for one of the most traditional purpose of futures, that is protecting against unfavourable moves in the physical market.

Table of Contents

List of Figures

List of Tables

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Signed:

Date:

Acknowledgements

During the last three years I have become a believer in persistence and help being the two most important components in completing the most challenging things. I would like to compare my experience of being a PhD student with climbing a very twisty, tall mountain with several resting spots. Definitely, it is not easy, but the moment you climb up the top and see magnificent scenery makes you realise that it all has its purpose. In this section, I want to express my sincere gratitude to those people whose encouragement and care gave me support not to slide off this twisty path, and whose involvement allowed me to complete this thesis.

First and foremost, I would like to thank my primary supervisor, Professor Bart Frijns. I am very grateful that in May 2015 he helped me to prepare my first research proposal, which was later accepted, and I moved to Auckland to undertake my PhD. I am very happy it happened this way. To me, Bart is a man of many talents. First, he is an encouraging, bright academic supervisor. Bart's constant confidence in my abilities, patience and great ideas often empowered me to overcome the most challenging times. He taught me how to be a better writer and researcher. He highlighted the importance of thinking about the big picture first and then digging into details. Second, he is proactive, and always looking into seizing interesting opportunities. His connections with the industry made my PhD not only academically focused, but allowed me to build several meaningful industry connections. For all those reasons, I am grateful.

I am also very thankful to my second supervisor, Professor Alireza Tourani-Rad. He has played an important role in the overall process and constantly provided useful feedback and insights regarding my empirical chapters. He also introduced me to the world of teaching, as he offered my first teaching assistant position to me. This experience was very valuable, and I am grateful for that.

My third supervisor, Dr. Adrian Fernandez-Perez, also positively influenced the development of this study from the first draft to the thesis. His friendly attitude made me feel comfortable to seek his advice on the most trivial questions. I am grateful for the fact that he always was very quick in providing his constructive feedback. He was always ready to help when I needed advice, especially in technical questions. I feel privileged to have had a chance to meet Bart, Alireza, and Adrian and work with these admirable people, whose guidance was crucial during my PhD journey, and whose advice I will never forget.

Each chapter of this thesis was very demanding and required me to stay long hours at the university. Luckily, AUT Finance Department has a very stimulating environment, created by academic staff members. In particular, I would like to acknowledge Dr. Olga Dodd, who made me feel like home even though it was on the other side of the equator. Additionally, I would like to thank Dr. Ivan Indriawan and Dr. José Da Fonseca for their responsiveness in many matters. Also, I thank Carol, Eathar, and Tracy for their kind assistance in helping me to run the administrative side of the PhD smoothly. A sincere thanks go to my fellow PhD students, Chris, Anjela, Marinela, Tu and Harris. I am delighted with how much drive these guys have. I truly enjoyed our chats about random things and participation in memorable joint adventures.

I would also like to thank Mark Neal from DairyNZ for helping me to obtain data and for familiarising me with nitty-gritty details of New Zealand farming business.

Last but not least, I want to acknowledge my friends and family. There are two very special people I met in New Zealand. Olga, thank you very much for being a kind, intelligent and adventurous friend. Timothy, thank you very much for all the memorable moments we share, your patience and encouragement. You are my reminder of the most important things in my life. Finally, a sincere thank-you is owed to my mum. Knowledge is not created overnight. She taught me very simple things, which later enabled me to leave my hometown and embrace the world. I am grateful for her kindness, constant care and love she provides for me.

Chapter 1

Introduction

Over time, derivatives markets have experienced a large expansion. The global size of the derivatives market, including OTC derivatives and exchange-traded derivatives, in 2018 is about $688 trillion, whereas thirty years ago it did not exceed $1 trillion.[1] The increased volume of participation partly can be attributed to the introduction of new innovative products, which serve the changing needs of market participants. Without doubt, derivatives bring a number of important benefits to the financial system and its participants. Despite that, whenever a financial crisis or collapse of a corporation occurs, derivatives are blamed for their misuse and sometimes described as "time bombs, both for the parties that deal in them and the economic system" (Buffett, 2002). Therefore, as new products are introduced regularly, it is important to understand the properties and usefulness of these derivatives.

This thesis focuses on two recently introduced products. Firstly, we consider volatility ETNs. Volatility trading has become very popular after the introduction of volatility derivatives, and really took off after VIX ETNs were launched in 2009. The two most liquid volatility ETNs are the S&P 500 VIX short-term futures ETN, or the VXX, and the inverse VIX short-term futures ETN, or the XIV. Briefly, the VXX gives long volatility exposure while the XIV short volatility exposure. These products can be interesting to hedgers and speculators who want to bet on the direction of the future value of the US short-term market volatility. Informed investors can either switch between the VXX and XIV, based on their information

[1] The Bank of International Settlements statistics.

about the future value of volatility, or trade in one market by switching between long and short positions. An interesting fact about the VXX and XIV is that the VXX historically demonstrates a steady loss, and the XIV a growth in value. Only during turbulent times does this dynamic change. This negative performance of the VXX might prompt questioning the information content of the VXX. Understanding which product is the most informative about the future value of volatility and at what times is an interesting question.

Secondly, we consider the NZX dairy derivatives. In 2010, the NZX launched its first dairy derivatives - WMP futures. Since then it has expanded product offerings to other dairy futures and options contracts; however, WMP futures remain the most liquid contracts with the highest trading volume. The WMP futures are the most liquid futures contracts and the WMP options are the most liquid options contracts among dairy derivative products. The proliferation of WMP futures and options is related to the importance of WMP to the New Zealand economy. In fact, New Zealand exports about 95% of its dairy products and dairy exports are dominated by WMP. An interesting fact about WMP traded in global markets is its high volatility.[2] Given the increasing popularity of WMP options and futures markets, we aim to investigate two questions. The first one is related to the information content of this derivatives market, in particular, whether it contains useful information about future volatility of the spot market. The second question is motivated by the high volatility of dairy commodities and its effect on the financial situation of dairy farms. High volatility of WMP translates to volatile payouts for liquid milk, and, thus, has a direct impact on farms' profitability. The absence of risk management practices in such a risky business environment poses significant risks, not only to farming businesses but to the financial stability of New Zealand overall. After the introduction of dairy derivatives by the NZX, New Zealand dairy farms for the first time received access to risk management tools. Given the increase in popularity of WMP futures in recent years, it is timely to investigate benefits these contracts offer to dairy farms.

[2] For example, the annual volatility of WMP futures between 2012 and 2018 is about 37%, while soybeans, gold, cocoa, sugar, and crude oil have the annual volatility of 18%, 24%, 31%, 34%, and 35%, respectively.

In Chapter 3, we perform a study on the intraday price discovery between the VXX and XIV, using granular data with a one-second frequency. We find that, according to several popular price discovery measures, informational leadership exhibits strong time variation, meaning that both ETNs can effectively incorporate relevant information about the future value of volatility. When we estimate the determinants of the price leadership, we find that informed traders prefer to transact in the market that has lower trading costs, and greater liquidity, measured by trading volume or number of trades. We further document that, during market downturns, the price discovery function of the VXX decreases, which we attribute to an increased hedging demand, which is not only driven by information but also is affected by an overreaction to extreme market conditions. The findings of the first essay are important for volatility traders who want to trade on their information about the direction of future US stock market volatility.

In Chapter 4, we continue to explore the informational role of derivative markets, but turn to dairy derivatives products. We investigate how informative WMP options can be about the future volatility of WMP futures. To achieve this goal, we examine the statistical properties and information content of option implied volatility for the dairy market. The measure of the implied volatility is represented by a dairy volatility index that we construct and term it the DVIX. We investigate its forecasting power of future volatility both in-sample and out-of-sample. We find that the DVIX has a better predictive power than historical volatility. However, it does not subsume all historical information, as a combination of volatility forecast based on historical averages or the GARCH-type forecast and the DVIX has a better forecasting power than the DVIX alone. Our findings have broad implications. The ability to predict volatility of the dairy market is of great interest to investors and all the supply chain participants, including dairy processors and farmers. High volatility signals a high level of uncertainty that would require to determine appropriate hedging policies. The estimate of future volatility is also required for portfolio managers who need to estimate a risk measure, such as Value-at-Risk. Additionally, as we find that the volatility estimate based on historical data complements the

DVIX, indicating that the WMP options market may not yet be fully informationally efficient. We suggest how to create the best composite forecast of volatility out of a standard time-series volatility and the DVIX.

In Chapter 5, we investigate the effectiveness of WMP futures to protect the net profit margin of NZ dairy producers, i.e., we focus on the risk-management purpose of derivatives. We develop a profit margin hedging strategy, which aims to protect NZ dairy farmers from downside risk. We perform the analysis both for a representative farm and for individual farms. First, we find that the profit margin hedging strategy can significantly decrease the risk without sacrificing the returns. We find that the results hold even after introducing different levels of commissions. By analysing the risk and return of discretionary cash, we account for all cash expenses a farmer needs to cover during a farming season. Additionally, we document that the adoption of profit margin hedging reduces the chances of financial distress, measured as the likelihood to observe negative discretionary cash during a given year. Another interesting result we find is an association between financial leverage and hedging effectiveness. We find that generally farms with higher levels of debt, sorted by the debt to asset ratio, have the highest percentage improvement in the mean of discretionary cash. The findings of the third essay are of great interest both for the NZ farmers and for the banking system of NZ. NZ farms have a high dependence on debt provided by the banks. In comparison to farms in the EU countries, farms in NZ have the second highest average debt to asset ratio, being behind Denmark only (Loughrey et al., 2018). For example, in 2015 the median debt to asset ratio in NZ was about 50%, while the median debt to asset ratio in the EU countries is just above 20%. The implementation of profit margin hedging can facilitate the improvement of the sustainability of the farming business, and, as a result, reduce risk for the lending institutions.

Overall, the results of the three chapters shed light on the properties and uses of two new derivative products - volatility and dairy derivatives. For volatility ETNs, we find that both of the ETNs contribute to the process of price discovery, which is time-varying. This is an important finding as it shows that both products are quick to incorporate new information about the future value of volatility, at different times. As for dairy derivatives, we document two main interesting findings. Firstly, WMP options contain valuable information about the future short-term volatility of WMP. This finding means that the WMP options market has reached some level of informational efficiency, and incorporates some of the historical information and private information of informed investors about the future dynamics of the spot market. Secondly, WMP futures can be effectively used by primary milk producers in hedging downside risk of the milk price. This use is particularly desirable, as when the NZX launched its dairy derivatives, its motivation was to provide dairy participants with a tool against high volatility of dairy products.

Chapter 2

A Primer on the Derivatives Market and Their Functions

2.1 Introduction

This chapter briefly discusses the original purpose of derivatives and their informational role in relation to volatility forecasting. These background materials form a foundation for the rest of this thesis. We start the discussion by highlighting the purpose of derivatives with references going back as far as antiquity. We then move to the modern finance period and discuss some recent developments in futures, options and ETNs. In particular, we focus on relatively new derivatives markets, such as the volatility derivatives market and the NZ dairy derivatives market. Next, we elaborate on the informational role of derivatives markets and highlight how options markets can be useful in predicting future volatility. The chapter concludes by connecting the discussed topics and formulating the research topics of the following empirical chapters.

2.2 History of Derivative Markets

Modern financial textbooks often misrepresent historical facts related to derivative markets. For example, Hull (2008) states that derivatives have become important during the last 30 years, that is, since the 1980s. In fact, derivatives have a very long history going back as far as antiquity. The Code of Hammurabi, which was created by Babylonian king Hammurabi around 1750 BC, probably makes the first reference to derivatives which insure farmers against a poor harvest season:

> "If any one owe a debt for a loan, and a storm prostrates the grain, or the harvest fail, or the grain does not grow for lack of water; in that year he need not give his creditor any grain, he washes his debt-tablet in water and pays no rent for the year" (Whaley, 2006, p. 11).

The derivative contract implicitly involved here can be interpreted as a binary put option between the farmer and the lender. The option expires worthless if the harvest season is normal, and the farmer has a right not to pay interest if the harvest season is poor.

Another example of an early derivative, namely a forward contract, which is designed to secure a supply of wooden planks, dates back to the nineteenth century BC. In this example a supplier of wood promises to deliver some wooden planks on a future day:

> "Thirty wooden [planks?], ten of 3.5 meters each, twenty of 4 meters each, in the month Magrattum Akshak-shemi will give to Damqanum. Before six witnesses (their names are listed). The year that the golden throne of Sin of Warhum was made" (Weber, 2009, p. 434).

This evidence suggests that the first derivatives were traded over-the-counter (OTC), directly between two parties and could be customised according to parties' needs. Trading had some level of regulation, involving the description of goods, dates and a list of witnesses. Trading often took place in temples, which played not only religious but also a commercial role at the time. This evidence also suggests that, since ancient history, derivatives have played an important role in daily life and were used to eliminate future uncertainties, i.e. for hedging purposes.

Not only were derivatives used for risk-reduction, but also for speculative purposes. A great example of the presence of speculators in derivative markets is the tulip mania in the 17th century in Europe. Tulips were introduced in Holland at the end of the 16th century and the public became fascinated by them, especially the upper class. The prices reached a state of a bubble, as tulip bulbs were traded at a price higher than a house at that time. Speculators didn't even need to hold on to tulip bulbs, as trading on tulip bulbs involved contracts for difference (Weber, 2009). Contracts for difference were created by merchants in Antwerp around 1540. They were similar to the modern cash-settled futures contracts but didn't possess the same safeguard features as modern futures. They were settled by a single and potentially large cash transaction at the settlement date. In 1637, spot prices of tulip bulbs collapsed and many speculators defaulted on their obligations.

The first organised futures exchange emerged in Osaka, Japan in 1710, and was officially acknowledged as a rice exchange in 1730. Trading on the Dojima Rice Exchange was governed by rules similar to modern-day futures exchanges: traders were members of the exchange, contracts were standardised, and a clearinghouse stood between buyer and seller, reducing the risk of default. In the US, the first-ever futures contracts were introduced in Chicago and were listed at the Chicago Board of Trade (CBOT) in 1848. The underlying of the contracts were grains, such as corn, oats, and wheat.

2.3 Development of Derivative Products

Since the first futures exchange was founded in Japan in the 18th century, nowadays more than 75 modern exchanges exist worldwide. Options, futures and Exchange Traded Products (ETPs) are traded on almost any variable. Underlying variables can be standard assets such as commodities, stocks or currencies, or exotic variables, such as weather, electricity or volatility.

Modern technologies replaced the traditional open-outcry trading pits to electronic financial markets. Before the era of electronic trading, all traders converged on a trading floor and verbally communicated their bid and offer prices. Replacement of floor-based open outcry systems with electronic trading brought increased speed and lower transaction costs. The increased market accessibility fostered growth in futures and options market participation. During the last twenty years, 1998-2017, the total volume of futures and options contracts traded globally grew from 1,482 million to 25 billion contracts.[1]

Recent innovations in derivatives are Exchange Traded Notes (ETNs), which were introduced in 2006. ETNs are unsecured debt obligations issued by financial institutions and are traded, like individual stocks, on a stock exchange. ETNs seek to replicate the performance of some index or the inverse of the performance of the index. Thus, they offer investors both long and short exposure to an index, which otherwise could only be accessible through trading in futures, options, or short positions in the index. Leveraged ETNs provide even more flexibility, as they seek to provide investors with returns that are multiples of the underlying index.

The proliferation of derivatives can be partially attributed to the timely introduction of new products to satisfy the changing needs of investors. An example of those products would be volatility derivatives. Market participants have always been looking for an asset class that is negatively correlated with equities. During the recent Global Financial Crisis, it became apparent that commodities, which are traditionally viewed as negatively correlated with equities, failed to deliver their expected diversification benefits (Berkowitz and DeLisle, 2018). In contrast, volatility, which can be measured by the CBOE Volatility Index (VIX), is negatively correlated with equity returns. The VIX was launched in 1993 by the Chicago Board Options Exchange (CBOE) and is defined as the market's estimate of the S&P 500 index's volatility during the next 30-day period. However, it took almost a decade before trading in volatility became possible. The CBOE launched VIX futures in 2004, followed by VIX options in 2006. In 2009, Barclays Bank issued the first VIX ETNs with ticker symbols VXX and VXZ. Volatility ETNs gained huge popularity among retail investors as they share similar trading characteristics with stocks, do not have margin requirements, and need a significantly lower investment capital in comparison to VIX futures. Moreover, some large institutional investors, such as endowments and pension funds, are restricted to trade in futures and options but are allowed to trade in ETPs. The two most successful volatility ETNs are the VXX and the XIV.[2]

[1] https://fia.org/articles/total-2017-volume-252-billion-contracts-down-01-2016

[2] The XIV had $1.9 billion in assets before it lost 97% of its value in a single day on February 5, 2018. Shortly after that, trading in the XIV was terminated, as the issuer announced the early

Another example of the timely introduction of new derivatives are the dairy derivatives launched by the New Zealand Stock Exchange. Historically, the NZ economy has had a strong reliance on the export of dairy products. In 2017, it contributed $NZ14 billion to annual export, 38% of the total primary industry exports' value.[3] Dairy commodity prices are characterised by high volatility, and in October 2010 the NZX launched its first dairy derivative - a WMP futures contract, to help all involved in the dairy supply chain to manage price risk. Later, it expanded its product range with four more futures contracts and three options contracts. Since its inception, the NZX dairy derivatives market has experienced a consistent increase in volumes and participation. The majority of volume traded is in WMP futures contracts. In November 2010, the first full month after trading began, there were only 30 lots of WMP futures contracts traded. During 2017, the monthly average trade volume grew to 16,732 lots. Across all eight derivatives, the average monthly trading volume in 2017 amounted to 25,844 lots.

2.4 Information Content of Derivatives Markets

The view that informed investors might prefer to trade on their information in the options markets, rather than in the stock market, has been supported by academics since the work of Black (1975). The arguments include the leverage inherent to options, as well as the ease of exploiting negative information about the stock's price by buying a put or selling a call option. Recent empirical work documents the information leadership of the options market relative to their underlying. For example, Chakravarty et al. (2004) document the price discovery role of the options market. Pan and Poteshman (2006) find that the option volume, measured by a put-call ratio, predicts future stock price movements, and Cremers and Weinbaum (2010) show that the deviation from put-call parity is useful in predicting future stock performance. These studies provide evidence that the options markets are a venue for informed trading.

redemption of the notes, called acceleration.

[3]https://www.mpi.govt.nz/news-and-resources/open-data-and-forecasting/situation-and-outlook-for-primary-industries-data/

Not only have studies shown that options can be informative about the future price of an underlying asset, but also that they can be used to predict its future volatility (see, for example, Fleming, 1998; Frijns et al., 2010). The implied volatility is the volatility of the underlying asset which equates the observed option price with the theoretical option price. The implied volatility is often viewed as markets' consensus estimate of future volatility of the underlying asset over the course of an option's time to expiration. Therefore, provided that the options market is efficient and the pricing formula is correctly specified, implied volatility is often shown to be a superior estimate of future volatility relative to time-series models.

Early literature often relied on the Black and Scholes (1973) model to extract implied volatility. However, the Black-Scholes model does not hold exactly, and in practice the implied volatility from an option written on a specific underlying asset will depend on time to expiration and the strike price. To arrive at a point estimate of implied volatility a standardised measure would be desirable. In 1993, the CBOE introduced the Volatility Index (VIX) which has become the benchmark for equity market volatility. The VIX, which later was renamed the VXO, is a combination of eight near-the-money S&P 100 option implied volatilities and represents the implied volatility with 22 trading days to maturity.

In more recent research, alternative measures of implied volatility have been examined. A model-free option implied volatility has received considerable interest from both practitioners and academics. As the name implies, it is derived from the set of options' prices without assuming any specific option pricing model. In 2003, the CBOE modified the methodology of the volatility index calculation from using Black-Scholes implied volatilities to a model-free approach. Calculation of the VXO used near-the-money options only, while for VIX a wide range of strikes is required. Additionally, the new VIX replaces the S&P 100 options with S&P 500 options. The VIX has become a widely-watched index of expected future market volatility over the next 30 calendar days. Later, the CBOE expanded the range of volatility indexes to different underlyings, including commodity-related ETFs on crude oil, gold, silver, gold miners and the energy sector.

As we discussed earlier, options markets not only contain information about the future returns on an underlying asset, but also can be used for predicting its future volatility. The important question is whether the forecast based on implied volatility outperforms the forecast based on historical data. A large body of empirical literature, which uses Black-Scholes implied volatility, generally justifies the forward-looking nature of implied volatility and shows that implied volatility has a higher explanatory power than time-series models (historical volatility or GARCH models). However, the conclusion about whether time-series models contain information incremental to implied volatility differs across different markets. For example, for options on currency, Jorion (1995) show that implied volatility subsumes all information contained in estimators based on time-series models. In contrast, Szakmary et al. (2003) find that for some agricultural commodities, such as cocoa, feeder cattle and lean-hogs, the implied volatility forecast does not encompass the historical forecast.

Similarly to Black-Scholes implied volatility, research which uses the model-free approach to extract implied volatility concludes that implied volatility outperforms estimators based on time-series models. For example, in the equity market, Carr and Wu (2006) show that the VIX is an efficient predictor of the S&P 500 index future realized volatility and encompasses the information contained in the GARCH forecast. Moreover, some studies demonstrate that model-free implied volatility is a superior estimate to Black-Scholes implied volatility, as it subsumes all information contained in Black-Scholes implied volatility (see, for example, Jiang and Tian, 2005; Wang et al., 2012).

With this background in mind, the next section connects the above mentioned arguments and outlines the three essays of the thesis.

2.5 Outline of the Thesis

As we discussed in Section 2.3, derivative markets are dynamic and new products are introduced frequently. It is important to understand the properties and usefulness of these new derivatives. The overall aim of this thesis is to build an understanding of the use and properties of two innovative derivatives products, which are volatility and dairy derivatives. As we argued in Section 2.4, options markets are a venue for informed trading. In particular, implied volatility, which is extracted from option prices, provides useful information about future volatility. Two volatility ETNs, the VXX and XIV, allow market participants to speculate on the direction of future US stock market's volatility. It would be interesting to investigate which of these two ETNs is more informative about future volatility and when. In Chapter 3 we answer this question by using price discovery measures of Hasbrouck (1995) and Lien and Shrestha (2014). In Section 2.4, we also discussed that implied volatility forecast outperforms historical volatility forecasts. Motivated by these findings, in Chapter 4, we focus on the NZX dairy derivatives market and construct a dairy volatility index (termed DVIX). We investigate the returns-DVIX relationship and then evaluate the predictive power of the DVIX in predicting subsequent realized volatility of WMP futures. Additionally, we compare its information content relative to historical volatility and GARCH volatility. In Chapter 5, we continue the exploration of the usefulness of the NZX dairy derivatives market, but we move from its informational role to risk-management function. As we have pointed out in Section 2.2, derivatives markets have been used for hedging even before standardised stock exchanges emerged. Given the increase in popularity of WMP futures in recent years, an interesting question would be to investigate the benefits these contracts offer to NZ dairy farmers. In Chapter 5, we explore the effectiveness of profit margin hedging with WMP futures within the New Zealand context.

Chapter 3

Determinants of Intraday Price Discovery in VIX Exchange Traded Notes[1]

3.1 Introduction

With the introduction of VIX derivatives in 2004, volatility, for the first time, became a tradable product. However, it was not until the introduction of VIX Exchange Traded Notes (ETNs) that trading in volatility really took off. These products, first issued in 2009, have gained huge popularity among market participants, with currently 19 VIX-related ETPs.[2] Volatility ETNs allow market participants to take direct, inverse or leveraged positions in S&P 500 volatility. Some ETPs monitor the curvature of the VIX futures term structure and offer direct or inverse exposures to short- or medium-term futures contracts on the VIX. The most popular direct ETN is the iPath S&P 500 VIX Short-Term Futures ETN (VXX) and the most popular inverse ETN is the VelocityShares Daily Inverse VIX Short Term ETN (XIV). Despite the VXX losing about 99% of its value and undergoing four reverse splits since inception, this ETN remains extremely popular, with an average trading daily volume of 60 million contracts by the end of 2016. The XIV, on the other hand, has gained more than 100% since inception and, by the end of 2016, had an

[1]This chapter is based on Fernandez-Perez, A., Frijns, B., Gafiatullina, I., & Tourani-Rad, A. (2018). Determinants of intraday price discovery in VIX exchange traded notes. *Journal of Futures Markets*, 38(5):535-548.

[2]An exchange-traded product (ETP) is a derivative security which is traded on an exchange. ETPs are typically benchmarked to indices, stocks, commodities, or may be actively managed. There are several different types of ETPs, including Exchange-traded funds (ETFs) and Exchange-traded notes (ETNs).

average daily trading volume of close to 30 million contracts.

The demand for the VXX could be due to either diversification benefits or speculation on future volatility changes. Changes in the VIX are negatively correlated with changes in the S&P 500. Thus, adding volatility exposure to a portfolio can potentially be a risk mitigation strategy during market turmoil (Signori et al., 2010). However, several studies document that VIX ETNs do not provide an effective hedge when held as a passive buy-and-hold investment.[3] As a result, when used as a portfolio insurance tool, investors might prefer to use the VXX for short-term, rather than for long-term exposure. However, this kind of trading might be affected by overreactions to extreme market conditions, and thus impair the information content of the VXX's price.

The other reason for trading in VXX and XIV is speculation about future changes in volatility, as the VXX and XIV can be used for betting on the direction of volatility. The VXX increases at times when the XIV falls (and vice versa) and, thus, informed investors might choose to transact in one or the other market, which might result in the flow of information between these two volatility ETNs. If informed traders are more likely to choose one particular market to transact in, then this market dominates the price discovery process and tends to lead prices of the other market. Given the tremendous trading volume in the VXX with exceptional negative returns in the long-run, does that mean that it is dominated by noise traders, who lower the contribution to price discovery for the VXX? Or maybe informed investors switch trading strategies between the VXX and XIV according to changing market conditions, which subsequently leads to time variation in price discovery between the VXX and the XIV?

In this chapter, we are interested in the informational leadership between the direct and inverse short-term volatility ETNs. First, we examine the informational leadership between the direct and inverse short-term volatility ETNs. Second, we examine key determinants of informational leadership. We investigate to what extent the price discovery measures of Hasbrouck (1995) and Lien and Shrestha (2014)

[3]The long-term performances of VIX ETNs versus the VIX show substantial deviations primarily due to the negative roll-over costs during contango markets, which is called "contango trap" (Alexander and Korovilas, 2012a; Whaley, 2009).

are influenced by spread and liquidity measures, and whether a change in market conditions affects the ability of the VXX and XIV to incorporate new information. Using high frequency data, our results show strong time variation in price discovery between the VXX and XIV. This finding indicates that neither the VXX nor the XIV is redundant from the perspective of price discovery and that informed traders opt to switch between these two markets. We document that relative spread- and volume-related metrics are significant determinants of price discovery. We further document that price discovery of the VXX reduces on days with negative returns on the S&P 500 and on days when the level of the VIX increases. This suggests that during those days informed investors choose to trade in the XIV, as the informational content of the VXX might be affected by increasing hedging activity driven by overreaction to the extreme market conditions.

While several prior studies have focused on volatility as an asset class and its diversification benefits, only a few address the question of informational dominance between the VIX and its derivatives, or between the different volatility products. Shu and Zhang (2012) examine the lead-lag dynamics between the VIX and VIX futures on a daily basis and conclude that VIX futures lead the VIX. Frijns et al. (2016), using intraday data, find evidence of bi-directional causality between the VIX and its futures, with VIX futures prices becoming more informative over time. As for volatility ETPs, Bordonado et al. (2016) study the price discovery relationship among direct, leveraged and inverse VIX ETPs employing 1-minute data and identify the price discovery leader within each ETP category.[4] Bollen et al. (2017) expand the analysis of lead/lag relation for the pairs of VIX futures vs. VIX options, VIX futures vs. the VXX and VIX futures vs. the VIX. Using intraday price movements, they find that, for each pair, information about volatility originates in VIX futures, the VXX and VIX futures, respectively, with increasing dominance of VIX futures over time. However, none of the research to date has addressed the price discovery

[4]Bordonado et al. (2016) examine price discovery between three pairs of VIX ETPs and find that, for the direct and indirect ETPs pairs, the older markets and with higher trading volumes are the informational leaders. However, for the leveraged ETPs, the newer market and with lower trading volume impounds information faster, which might be explained by the fact that one of the ETNs in this pair is TVIX which stopped issuance on the 21st of February 2012. Since that time the TVIX has had a significant premium over its indicative value, which might affect the results.

process between different ETN categories, which might be attractive to informed traders during different market conditions, and levels of volatility, and, thus, exhibit switching pattern in informational efficiency.

This study contributes to the existing literature in several ways. First, we contribute to the price discovery literature in volatility derivatives by providing an empirical examination of which volatility ETN plays the dominant role in the mechanism of price discovery in the VIX. Second, our findings help to understand the impact of classical determinants of price discovery as well as market conditions on the efficient pricing of the volatility ETNs.

The remainder of this chapter is organised as follows: Section 3.2 provides background information on the VIX and the two most popular direct and inverse volatility ETNs. In Section 3.3, we describe the methodology adopted in this study. Section 3.4 explains the data and reports summary statistics. In Section 3.5, we present the empirical results. Section 3.6 concludes.

3.2 Background

In this section, we review some of the relevant literature and discuss the pricing methodology and properties of the VXX and XIV. We also show that, by design of the underlying indices of these ETNs, there should exist a cointegrating relation between the VXX and XIV with a cointegrating vector of $(1,1)'$.

3.2.1 Volatility Derivatives

The most common strategy for protecting equity portfolios during market downturns is to use options on the S&P 500. The other strategy, which is documented to bring significant diversification benefits, is to take a long position in implied volatility (Dash and Moran, 2005; Daigler and Rossi, 2006). To allow for direct trading in volatility, VIX futures were introduced in 2004 and VIX options in 2006. However, trading in these VIX derivatives might be too sophisticated for retail investors, and many institutions are restricted from trading in futures and options directly (Whaley, 2013). Thus, the third generation of volatility products - VIX ETPs -

were introduced in 2009. ETPs have an equity-like structure and can easily be traded by both retail and institutional investors.

A number of studies consider the possibility of hedging portfolios with VIX futures and VIX ETNs, and conclude that, for passive buy-and-hold investors, direct VIX ETNs do not provide effective hedges. Husson and McCann (2011), Deng et al. (2012), Alexander and Korovilas (2012a,b), DeLisle et al. (2014), among others, attribute the poor long-term performance of VIX ETNs to the fact that most of them do not provide the performance of the VIX but instead track the performance of constant maturity futures indices. In turn, constant maturity futures indices suffer from roll-over losses, which are associated with the term structure of the VIX futures market, which is typically upward sloping. As an example, the reference index of short-term futures notes measures the return to a portfolio of one- and two-month VIX futures contracts, which are rebalanced daily to achieve an average maturity of one month. Most of the time, VIX futures prices exhibit an upward-sloping term structure and only during market instability change to a downward-sloping term structure. The daily rebalancing implies selling a fraction of holdings in the one-month futures contracts at a lower price than the price which was paid when it was purchased as a two-month contract, and simultaneously buying a fraction of holdings in the two-month futures contracts at a higher price. This strategy incurs a loss in value of the underlying index and, as a result, the value of the direct short-term futures notes, and in particular the VXX, shows a steady loss. Bahaji and Aberkane (2016) go beyond the buy-and-hold strategy and show that uninformed rational risk-averse agents can enhance the performance of their portfolios by dynamically taking short or long positions in short- or mid-term VIX futures indices. Hence, an optimal VIX futures investment strategy can be implemented by dynamically investing in the VXX and XIV ETNs, which are close counterparts of constant short-term long and short positions in VIX futures.

The conflicting evidence of poor long-term hedging performance, decreasing prices of the VXX and enormous trading volume in this product naturally prompts us to question the informational content of the VXX. As Alexander and Korovilas

(2012b) point out, only during times of market instability does the VIX futures term structure swing to backwardation and the VXX stops suffering from the negative roll cost effects, thus becoming an effective diversifier/hedge. The XIV is almost a mirror reflection of the VXX and experiences stable growth when the VIX falls, and loses its value when the VIX goes up. Informed investors might be scared away from investing in the VXX at times of low market volatility and might decide to trade on the information about future value of volatility in the VXX market at times of high market volatility. The opposite strategy should be exercised for the XIV. This switching behaviour of investors between the VXX and XIV during different market conditions and levels of volatility might lead to changes in price discovery between the above-mentioned products.

3.2.2 Cointegrating relation between the VXX and XIV

In 2009, Barclays Bank issued the iPath S&P 500 VIX Short-Term Futures ETNs, or the VXX.[5] The VXX is benchmarked to the S&P 500 VIX Short-Term Futures Index Total Return (TR), or SPVXSTR. The value of the SPVXSTR depends on its value on the previous day and two return components. The first component measures the return from a rolling long position in the first- and second-nearby VIX futures contracts, which creates an average time to maturity of one month. The second component of the index's return includes interest accruals based on the 3-month U.S. Treasury Bill rate. VXX's performance is linked to the performance of the SPVXSTR index minus an investor fee.

In 2010, Credit Suisse designed the VelocityShares Inverse VIX Short Term ETN with ticker XIV. The XIV is benchmarked to the S&P 500 VIX Short-Term Futures Index Excess Return (ER) with ticker SPVXSP, which measures the return only from a rolling long position in the first- and second-nearby futures contracts. The XIV's performance is linked to the performance of the SPVXSP index minus an investor fee.

The closing indicative value (CIV) of the ETN is designed to approximate the

[5] See prospectus of the VXX at `http://www.ipathetn.com/US/16/en/details.app?instrumentId=259118`.

economic value of this ETN. The CIV of each of the ETNs is calculated by the issuers on a daily basis. The value on the inception date was \$100. The CIV of both VXX and XIV are based on the closing levels of the underlying indices but have minor differences in their definitions. In Appendix 4.C, we present detailed calculations which show that both CIVs include the U.S. three-month Treasury rate return component and have either a direct or inverse relation with the return of the rolling VIX futures position. Mathematically,

$$CIV_t^{VXX} = CIV_{t-1}^{VXX} \times (1 + TBR_t + CDR_t) \times (1 - \frac{feeRate}{365}) \tag{3.1}$$

$$CIV_t^{XIV} = CIV_{t-1}^{XIV} \times (1 + TBR_t - CDR_t) \times (1 - \frac{feeRate}{365}), \tag{3.2}$$

where CDR_t is the Contract Daily Return, which is driven by the changes in VIX futures prices, TBR_t is the Treasury Bill Return, $feeRate$ is the investor fee rate which is equal to 0.89% and 1.35% per year for the VXX and XIV, respectively. Comparing Equations (3.1) and (3.2) one can see that, though at the first sight the VXX and the XIV are benchmarked to indices that differ by the return which might be gained through investing in the three-month U.S. Treasury rate, it follows that the CIVs of both of the series include it. This observation gives us grounds to argue that the two ETNs are cointegrated with a cointegrating vector $(1, 1)'$.

The trading price of the ETNs may substantially differ from the stated principal amount, intraday indicative value or CIV, due to the fact that the trading price reflects the investor supply and demand for that ETN. Whaley (2013) and Bordonado et al. (2016) evaluate the tracking performance of the most popular VIX ETPs by comparing the daily returns of market prices with the daily returns of their respective benchmarks and by comparing the observed market prices with the CIV. They conclude that both the VXX and XIV closely, though not perfectly, track the underlying benchmark indices and closing indicative values.

3.3 Methodology

To quantify the process of price discovery, we use the Information Shares (IS) developed by Hasbrouck (1995). As an alternative measure of price discovery we consider the Generalized Information Share (GIS) developed by Lien and Shrestha (2014).

Hasbrouck (1995) defines the price discovery process in terms of the variance of the innovations to the common efficient price. IS measures each market's relative contribution to the variance of the efficient price. If innovations are contemporaneously correlated, then the IS cannot be computed uniquely and is dependent on the order of the individual asset's price in the price vector. Hasbrouck (1995) suggests using a Cholesky factorisation and applying different orders of the prices to compute upper and lower bounds for the IS. Baillie et al. (2002) suggest the use of the mean of the bounds as an ultimate measure of the market's contribution to the price discovery process.

The IS framework is based on the assumption that the system consists of n unit root series and that there is a common stochastic component or efficient price that is shared by all prices, i.e. there are $(n-1)$ cointegrating vectors. In our study, we consider the system of two series $P_t = (P_{1,t}, P_{2,t})'$, a (2×1) vector of log prices for the VXX and the XIV, respectively. The series have the following vector error correction representation (Engle and Granger, 1987):

$$\Delta P_t = \alpha(\beta' P_{t-1} - E(\beta' P_{t-1})) + \sum_{i=1}^{k} \Gamma_i \Delta P_{t-i} + e_t, \tag{3.3}$$

where Γ_i is a (2×2) matrix, e_t is a (2×1) vector of the residuals. The VECM includes two parts: the first part, $\alpha(\beta' P_{t-1} - E(\beta' P_{t-1}))$, represents the long-run or equilibrium relation between the price series; the second part, $\sum_{i=1}^{k} \Gamma_i \Delta P_{t-i}$, captures the short-run dynamics induced by market imperfections. Assuming that the cointegrating vector is known, the VECM can be estimated by Ordinary Least Squares. The $E(\beta' P_{t-1})$ term captures systematic differences in the prices, i.e. $P_{1,t} - P_{2,t}$. In our case, this term mainly results from the indices having different values at inception. This term can be estimated by the sample average prior to the other parameters, corresponding to a “de-meaning” of the data.

It is assumed that there is a common stochastic component or efficient price that is shared by all prices. That means that there is one cointegrating vector $\beta' = (1, -\beta_1)$, such that $\beta' P_t \sim I(0)$. Even though prices are non-stationary, the linear combination $P_{1,t} - \beta_1 P_{2,t}$ is stationary and has a moving average representation:

$$\Delta P_t = \Psi(L)e_t = e_t + \Psi_1 e_{t-1} + \Psi_2 e_{t-2} + ..., \tag{3.4}$$

where $\Psi(L)$ is a matrix polynomial in the lag operator, $\Psi_0 = I_2$, $e_t = (e_{1,t}, e_{2,t})'$ and $e_t \sim iid(0,\Sigma)$ is a zero-mean vector of serially uncorrelated disturbances with covariance matrix Σ. Using the Beveridge-Nelson decomposition, we can present the price levels as:

$$P_t = P_0 + \Psi(1)\sum_{i=0}^{t} e_i + \Psi^*(L)e_t, \tag{3.5}$$

where P_0 is a constant (2×1) vector of initial values, $\Psi(1)$ is the sum of the moving average coefficients and $\Psi^*(L) = \sum_{k=0}^{\infty} \Psi_k^*$ and $\Psi_k^* = -\sum_{j=k+1}^{\infty} \Psi_j^*$. The requirement that $\beta' P_t$ is stationary implies that $\beta'\Psi(1) = 0$. The structure of β implies that two rows of $\Psi(1)$ are proportional to each other and can be expressed as

$$\Psi(1) = \begin{pmatrix} 1 \\ \beta_1^{-1} \end{pmatrix} \begin{pmatrix} \psi_1 & \psi_2 \end{pmatrix} = \begin{pmatrix} \psi_1 & \psi_2 \\ \beta_1^{-1}\psi_1 & \beta_1^{-1}\psi_2 \end{pmatrix}, \tag{3.6}$$

where $\psi = (\psi_1,\psi_2)$ is a common row.

Hasbrouck (1995) considers the term $\Psi(1)e_t$ as the component of the price change that is permanently impounded into the asset price and presumably due to new information. If Σ is diagonal (i.e. the market innovations are not contemporaneously correlated), then the variance of the long-run impact consists of the sum of two terms and each of them represents the contribution from a particular market to the innovation to the random walk component of the price. In this case, the IS of series i is defined as follows:

$$IS_i = \frac{(\psi_i \sigma_i)^2}{\psi \Sigma \psi'}, i = 1, 2. \tag{3.7}$$

If Σ is non-diagonal there is no unique value for the IS, but triangularisation of the covariance matrix can help to determine lower and upper bounds of IS. In this case, the formula is as follows:

$$IS_i = \frac{(\psi F)_i^2}{\psi \Sigma \psi'}, i = 1,...,n, \tag{3.8}$$

where F is the Cholesky factorisation of Σ such that $FF' = \Sigma$ and F is a lower triangular matrix. $(\psi F)_i$ is the i-th element of ψF.

In practice, IS is computed from estimates of a VECM as in Equation (3.3). Johansen (1991) shows that the matrix $\Psi(1)$ can be computed using the following formula:

$$\Psi(1) = \beta_\perp(\alpha'_\perp \Gamma(1)\beta_\perp)^{-1}\alpha'_\perp, \tag{3.9}$$

where $\beta_\perp$ and $\alpha_\perp$ are orthogonal vectors satisfying $\beta'\beta_\perp = 0$ and $\alpha'\alpha_\perp = 0$, respectively, and $\Gamma(1) = I_n - \sum_{i=1}^{k-1} \Gamma_i$.

Lien and Shrestha (2014) modify the IS proposed by Hasbrouck (1995) in two respects. First, the GIS can be applied to calculate a unique measure instead of upper and lower bounds for the IS. Second, it can be used in the situations where the cointegrating vector does not have to be one-to-one. The key idea of their approach lies in the diagonalisation of the correlation matrix instead of the covariance matrix, which results in a unique measure independent of the ordering. The GIS of series i is given by:

$$GIS_i = \frac{(\psi F^M)_i^2}{\psi \Sigma \psi'}, i = 1,...,n, \tag{3.10}$$

where $F^M = [G\Lambda^{-1/2}G^T V^{-1}]^{-1}$, $(\psi F^M)_i$ is the i^{th} element of ψF^M; Λ is a diagonal matrix containing the eigenvalues of the innovation correlation matrix on the diagonal; G is a matrix with the columns represented by the corresponding eigenvectors and V is a diagonal matrix containing the innovation standard deviations on the diagonal.

3.4 Data and Summary Statistics

We obtain intraday trade and quote data for the VXX and XIV from Thomson Reuters Tick History (TRTH) database maintained by SIRCA (Securities Industry Research Centre of Asia-Pacific) with millisecond precision. The starting date of the sample is 3 January 2012 and the ending date is 31 December 2015. We do not start the sample from the inception of the XIV to avoid any liquidity issues. To clean the data, we follow the modified procedure by Barndorff-Nielsen et al. (2009). The cleaning of quotes data is carried out in the following steps. First, we delete entries with a time stamp outside the 9:35 am to 15:55 pm window when the exchange is closed (we eliminate the first and last five minutes of the trading day). Second, we delete entries with a bid or ask price equal to zero. Third, we delete entries with a negative spread. Fourth, we delete entries for which the spread is more than fifty times the median spread on that day. Fifth, we merge entries which correspond to the same second, keeping only the last observation in the group with the same second time stamp. The last step in our cleaning procedure is to delete entries for which the mid-quote deviates by more than five median absolute deviations of the day from a centered median (excluding an observation under consideration) of 50 observations.

Table 3.1 provides summary statistics for the mid-point of the bid and ask quotes. Panel A reports summary statistics for the levels of the VXX and the XIV. The VXX has positive skewness and displays excess kurtosis. The first-order autocorrelation of the VXX is close to one and the Augmented Dickey-Fuller (ADF) test fails to reject the presence of a unit root at conventional significance levels. The XIV has positive skewness and negative excess kurtosis. The series are highly persistent and the ADF test fails to reject the presence of a unit root. The log transformation which is presented in Panel B smooths the data. The log of the VXX preserves the positive skewness but now exhibits negative excess kurtosis and the ADF test again demonstrates that the null hypothesis of a unit root cannot be rejected. The logarithm values of the levels for the XIV have negative skewness and negative excess kurtosis. According to the ADF test statistic, we cannot reject the null of a unit root at conventional significance levels showing that these series are non-stationary. Therefore, the logarithm of the VXX and the XIV are $I(1)$ processes. Panel C shows the descriptive statistics for the first difference of the logarithmic values of the VXX and the XIV. The ADF tests reject the null hypothesis of a unit root.

Table 3.1: Descriptive Statistics

	Mean	St. Dev.	Skewness	Kurtosis	$\rho(1)$	ADF
VXX	183.59	211.09	1.29	0.53	0.9999	-1.78
XIV	24.26	11.79	0.22	-0.97	0.9999	-1.65
log(VXX)	4.51	1.21	0.32	-1.33	0.9999	-0.92
log(XIV)	3.04	0.58	-0.63	-0.51	0.9999	-1.41
Δlog(VXX)	-1.18e-07	2.68e-04	63.62	82285	-0.022	-3854***
Δlog(XIV)	0.26e-07	2.83e-04	-87.11	100326	-0.035	-3898***

Note: This table presents descriptive statistics on the mid quote of the VXX and the XIV at 1-second frequency for the full sample. $\rho(1)$ denotes the first-order autocorrelation. The augmented Dickey-Fuller (ADF) statistics test the null hypothesis that an examined series has a unit root. **, *** indicates significance at the 5% and 1% levels, respectively.

Figure 3.1 shows the 5-day moving averages of the daily volumes of the VXX and XIV since the inception date of the ETN to the end of 2015.[6] After one year from inception, trading volume of the VXX started to grow and in December 2015 reached about 70 million shares per day, showing that this ETN gained acceptance as a volatility trading vehicle. The XIV also gained significant growth within one year after inception, though on a smaller scale. By the end of 2015, the average daily volume reached 30 million shares, almost half of the volume of the VXX. Similarly to the VXX, during big geopolitical events the trading volume in the XIV tends to increase indicating that the traders change their outlook on the future direction of volatility and take either long or short positions in volatility. Peak volumes occur in times of market instability. These days are: 8 August 2011, when "Black Monday" occurred as a response to the USA's credit rating downgrade; mid May 2012 when the event known as "Grexit" during the European debt crisis took place; 20 June 2013 with the Chinese banking liquidity crisis; 15 October 2014 turmoil sparked by weak US economic data and fears over the health of the global economy; 21 and

[6]Inception dates for the VXX and XIV are January 29, 2009 and November 29, 2010 respectively.

24 August 2015 market sell-off triggered by China's stock market crash. Figure 3.2 plots the daily log mid-quote of the VXX and XIV over time. The plot shows that both time series seem to be a mirror reflection of each other.

Figure 3.1: Daily volumes for the VXX and XIV (in 1,000,000s)

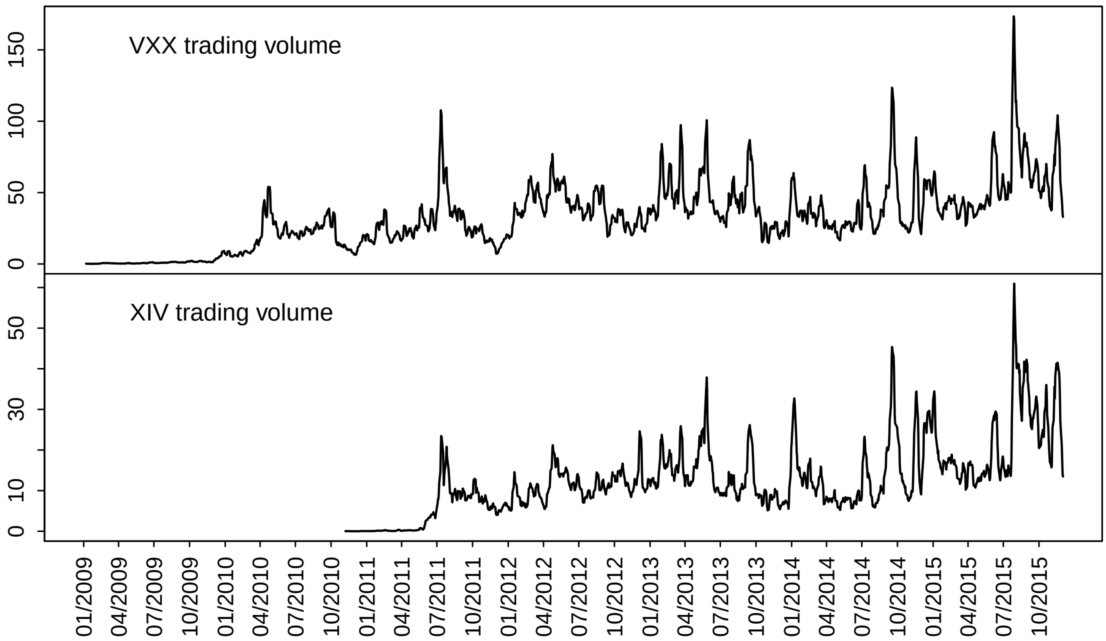

Note: This figure shows the daily trading volumes for the VXX and the XIV for the period between January 29, 2009 and December 31, 2015.

Figure 3.2: Daily closing prices for the VXX and the XIV

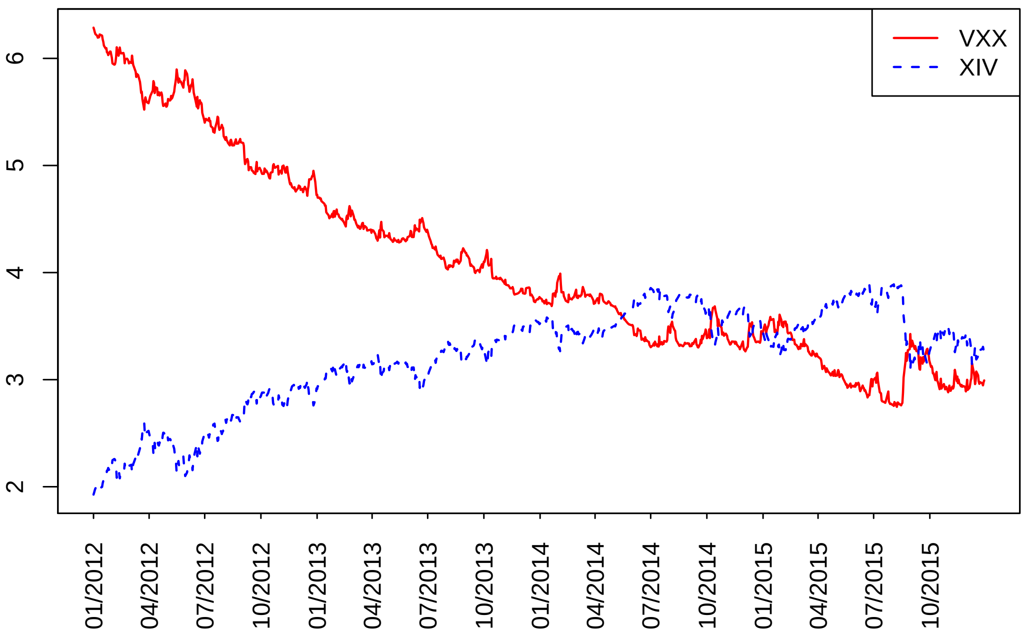

Note: This figure shows a daily log price history for the VXX and XIV from January 3, 2012 to December 31, 2015.

3.5 Empirical results

3.5.1 Cointegration Test

Before estimating the VECM in Equation (3.3), we examine whether the price series are cointegrated. For 92% percent of days the ADF tests show that the price series of both the XIV and VXX are non-stationary. Though they are non-stationary, we expect that they will not drift far apart from each other because they have similar underlying indices with the difference that the performance of the XIV is reversed to the underlying index. Thus, we expect the cointegrating vector to be close to $\beta = (1, 1)'$. We sequentially use the Johansen Likelihood-Ratio test, to examine whether there is one cointegrating vector. The null hypothesis that the number of cointegrating vectors is at most zero is rejected at the 1% level for 94% of the days in the sample. The next test indicates that there is at most one cointegrating vector that cannot be rejected for about 99.5% of those days.[7] The average cointegrating relation is $(1, 1.002)'$ and the null hypothesis that the CE is equal to $(1, 1)'$ cannot be rejected.

3.5.2 VECM and Daily Price Discovery Measures

To identify the lag length of the VECM that will be used to estimate the VECM for every day in the sample, we use the multivariate version of Schwartz's Bayesian Information Criterion (SBIC). After calculating the optimal lag length for each day, we calculate the average value, which is equal to 9. Thus, to compute the daily measures of price discovery, we consider a VECM with a lag-length $k = 9$.[8] Table 3.2 provides the results for the intraday VECM of order 9. The speed of adjustment coefficients α^{VXX} and α^{XIV} from Equation (3.3) represent the speed of convergence to the long-run equilibrium relationship between the VXX and XIV. Schwarz and Szakmary (1994) are the first who proposed using the relative magnitude of the error correction coefficients α^{VXX} and α^{XIV} to assess the contribution of each market to

[7] Based on the 1% level, we discard 6% of days from the analysis.
[8] Results do not change significantly when using different lags.

the price formation process. They argue that the price discovery leader is the market which initiates the mispricing $\beta' P_t$ and the price discovery follower is the market which responds to the disequilibrium. Thus, a smaller α (in absolute terms) indicates a price discovery leadership. In our case, α^{VXX} is smaller than α^{XIV} in absolute terms, thus suggesting that the VXX is the informational leader.

Table 3.2: VEC Model

	ΔVXX$_t$	ΔXIV$_t$
EC term$_{t-1}$	-0.019***	-0.036***
ΔVXX$_{t-1}$	-0.13***	-0.144***
ΔVXX$_{t-2}$	-0.087***	-0.099***
ΔVXX$_{t-3}$	-0.071***	-0.076***
ΔVXX$_{t-4}$	-0.055***	-0.065***
ΔVXX$_{t-5}$	-0.045***	-0.054***
ΔVXX$_{t-6}$	-0.037***	-0.045***
ΔVXX$_{t-7}$	-0.03***	-0.037***
ΔVXX$_{t-8}$	-0.022***	-0.029***
ΔVXX$_{t-9}$	-0.015***	-0.02***
ΔXIV$_{t-1}$	-0.115***	-0.155***
ΔXIV$_{t-2}$	-0.081***	-0.104***
ΔXIV$_{t-3}$	-0.063***	-0.081***
ΔXIV$_{t-4}$	-0.052***	-0.066***
ΔXIV$_{t-5}$	-0.043***	-0.055***
ΔXIV$_{t-6}$	-0.036***	-0.045***
ΔXIV$_{t-7}$	-0.029***	-0.037***
ΔXIV$_{t-8}$	-0.023***	-0.028***
ΔXIV$_{t-9}$	-0.015***	-0.019***
Adjusted R^2(%)	1.81	3.16
	[0.3, 4.83]	[0.4, 7.01]
No. of Days	957	957

Note: This table reports the results for the intraday VEC model difined in Equation (3.3). The VEC model is for estimated every day, and the average coefficients over the sample period, as well as the average adjusted R^2 are reported. 2.5th and 97.5th percentiles of the average adjusted R^2 are reported in square brackets. The *** is used to indicate significance at the 1% level.

When we consider the dynamics of ΔVXX we find evidence of significant negative autocorrelation and that lagged changes in the XIV have a negative and significant effect on current changes in the VXX. Similarly, we find evidence of significant negative autocorrelation in ΔXIV and that lagged changes in the VXX have a negative and significant effect on the current change in the XIV. The average adjusted R^2 for the intraday models of ΔVXX and ΔXIV are 1.81% and 3.16%, respectively.

We use Equation (3.8) to calculate upper and lower bounds of the IS and subsequently calculate mean value between the lower and upper bounds. To calculate the GIS, we use Equation (3.10). Table 3.3 shows distributional properties of both the IS and GIS. According to the IS, the price discovery contribution of the VXX is 52% and that of the XIV is 48%, suggesting that the XIV closely follows the VXX in the price discovery leadership. The daily autocorrelation in the IS is 0.816, showing strong persistence in the IS. The ADF test statistics are significant, suggesting that the IS series are stationary. Similar to the IS, the GIS establishes the VXX to be the price discovery leader. In comparison to the IS, the GIS is more volatile and less persistent, with a standard deviation of 0.202 in comparison to 0.152 and the autocorrelation coefficient drops to 0.77. The difference between 95th and 5th percentiles is wider for the GIS and equals to 0.643 in comparison to 0.514 range for the IS.

Table 3.3: Price Discovery Measures Descriptive Statistics

	VXX		XIV	
	IS	GIS	IS	GIS
Mean	0.515	0.520	0.485	0.480
5%	0.278	0.212	0.208	0.145
Median	0.510	0.518	0.490	0.482
95%	0.792	0.855	0.722	0.788
Std. Dev.	0.152	0.202	0.152	0.202
Skewness	0.253	0.081	-0.253	-0.081
Kurtosis	2.801	2.202	2.801	2.202
$\rho(1)$	0.816	0.770	0.816	0.770
ADF	-6.243***	-6.962***	-6.243***	-6.962***
No. of Obs.	957	957	957	957

Note: This table presents summary statistics of the ISs. $\rho(1)$ denotes the first-order autocorrelation coefficient; the augmented Dickey-Fuller (ADF) statistics test the null hypothesis that an examined series has a unit root. *** is used to indicate significance at the 1% level.

Figure 3.3 shows the 5-day moving averages computed from the VXX's daily IS and GIS. Both graphs exhibit high persistence in daily variation; however, there are also two apparent jumps in the price discovery measures. These days are marked on the graph and happened on 5 October 2012 and 8 November 2013 which coincide with a 1 for 4 reverse split of the VXX. At first sight, the abrupt increase in the

VXX's IS might be puzzling; however, analysis of the relative spread seems to explain the finding. On days when a reverse split occurs the price of the VXX increases 4 times which leads to a decrease of a relative spread to a comparable magnitude. Based on the extant literature (Fleming et al., 1996; Chakravarty et al., 2004), a reduction in trading costs positively affects price discovery, as it makes it cheaper for informed investors to trade on their information. In Section 3.5, we confirm that a decrease in relative spread is associated with an increase in the IS. Hence, reverse splits make the VXX cheaper to trade, which attracts more liquidity and, thus, also informed traders.

Figure 3.3: Price Discovery Measures for the VXX

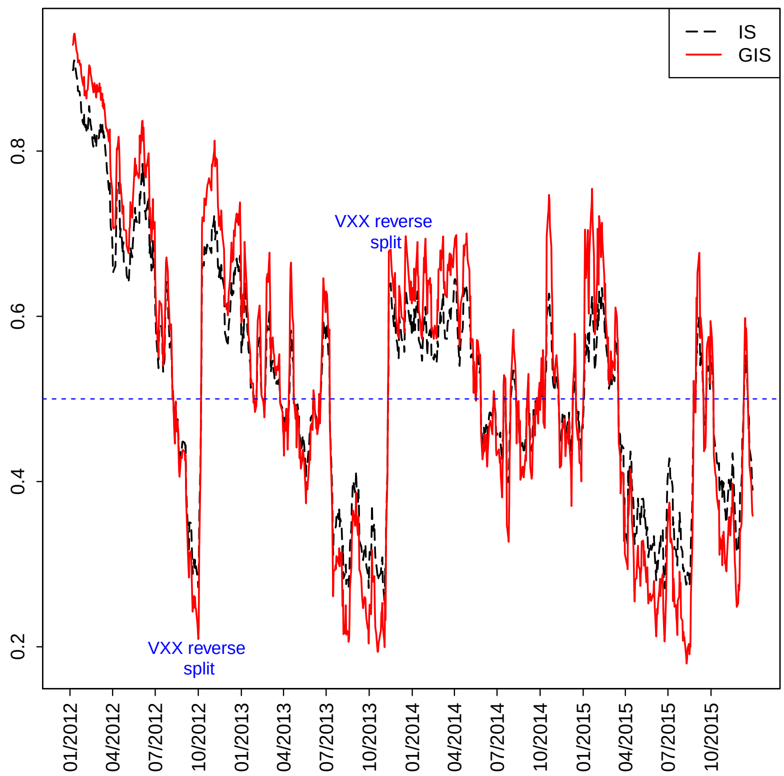

Note: This figure depicts the evolution of the IS and GIS for the VXX between January 3, 2012 and December 31, 2015.

The strong time variation we observe in the process of price discovery between two volatility ETNs is very different from those observed in stocks and stock option markets. Chakravarty et al. (2004) show that, for the period 1988-1993, the contribution of option markets to price discovery is about 17%, on average, with only a slight decrease over time. However, the significant time variation which we document for volatility ETNs is not unique and might be found in other markets.

Schlusche (2009) considers the relative price discovery between the DAX ETF and DAX futures. Though, on average, the futures market is the price leader, the daily measures of price discovery vary over time. Time variation in the price leadership between the VXX and XIV means that neither the VXX nor XIV is the only ETN which effectively can incorporate news about future value of volatility, but, rather, the importance of one or another market changes over time.

3.5.3 Determinants of Price Discovery

Several studies have focused on the determinants of price discovery between the spot and derivatives markets. Fleming et al. (1996) study the lead-lag relation between a stock index, and options and futures contracts on this index. They conclude that the leading market is the one which has the lowest trading cost because informed traders seek to earn the highest profit through execution of their trading strategies. Chakravarty et al. (2004), using stock and option data, show that informed trading in the option market has a significant association with the relative effective bid-ask spread, trading volume and stock volatility. In summary, the main price discovery determinants between the spot and derivatives market are found to be trading costs, market liquidity measures and volatility of the underlying security.

The area of price discovery between volatility products has emerged relatively recently and only a few studies examine the possible determinants of the price discovery process within this asset class. Thus, we investigate how the price discovery measures of the VXX are affected by trading costs, by volume-related metrics and by market conditions.

Frijns et al. (2016) are among the first to study the dynamic relation between the VIX and VIX futures using intraday data. Using Granger causality tests, they find that over the time the VIX futures have become more important in the pricing of volatility and that this dominance is more pronounced on days when the stock market experiences a significant decline and on days with high values of the VIX. Chen and Tsai (2017) take a step further and apply classical measures of price discovery to the VIX and its futures. In line with previous results, they find that VIX futures play a dominant role in the price discovery process and show that the contribution increases with the increase in the futures basis, which is defined as the difference between the VIX and VIX futures. Additionally, they find that VIX futures play a more important role in the price discovery process on days with macroeconomic news announcements related to the consumer price index. Traditional determinants of price discovery, including the trading volume and bid-ask spread, are not found to be the key determinants, which is not surprising as investors cannot directly trade or replicate the VIX and thus no relative measures of liquidity and trading costs can be computed.

To measure trading costs, we use the relative spread, which is defined as the daily average value of the difference between ask and bid quotes relative to the quote midpoint. The volume-related metrics include the total number of trades and the trading volume, all calculated on a daily basis. As a measure of market conditions, we consider the change in the VIX and the return on the S&P 500. The level of the VIX is an easily observed variable, which, according to the pricing formula, should influence the value of the volatility ETNs in opposite directions. During times of high levels of fear, direct ETNs should benefit and increase in value as during these times the futures curve swings from contango to backwardation. At the same time, the inverse ETN should drop in value but has a potential to grow when the market reverts to the normal state. It is important to note that while trading in the VXX might be driven by either the demand for portfolio insurance or speculation on future volatility spikes, trading in the XIV cannot hedge downward movements in the stock market. In the case where investors have a good timing ability about future volatility spikes, they can choose in which market to trade on their information, but if trading in the VXX is driven by the necessity for portfolio insurance, during increased market turbulence, one would expect a decrease in the price discovery of the VXX.

To assess the influence of the discussed variables on changes in the price discovery measures, we run the following regression:

$$\Delta PD_t = \beta_0 + \beta_1 \Delta PD_{t-1} + \beta_2 \Delta Spread_t + \beta_3 \Delta Volume_t + \beta_4 \Delta MarketConditions_t + e_t, \tag{3.11}$$

where ΔPD_t is the daily change in price discovery of the VXX between days $t-1$ and t, $\Delta Spread_t$ is a change in the ratio of effective spreads, $\Delta Volume_t$ is a change in the ratio of one of the volume-related metrics, $\Delta MarketConditions_t$ is either ΔVIX_t or ΔSPX_t, which are a change in the VIX or the log return on the S&P 500 on day t.[9] All the ratios are taken as the VXX value divided by the XIV. We control for the persistence in the price discovery measure by including its lagged value.

Panels A and B of Table 3.4 report the regression results for the daily changes in the IS and GIS, respectively. For the IS, all four model specifications provide a good fit with the lowest adjusted R-square of 29.3% and all coefficients are statistically significant at the 5% level or higher. In line with previous studies, the more liquid market is the informational leader, which is indicated by the significant and positive coefficient for the relative number of trades and the trading volume, and the significant and negative coefficient for the ratio of the relative spreads. This finding implies that, at times when trading costs in the VXX increase relative to the XIV, the IS of the VXX drops. Additionally, an increase in trading in the VXX (measured either by trades volume or trades number), or a decrease in trading in the XIV is associated with an increase of the IS of the VXX.

The change in the VIX has a negative association with the change in the IS of the VXX, while the change in the returns on the S&P 500 has a positive association. This indicates that market downturns, characterised by a decrease in the S&P 500 or an increase in the VIX, are associated with a decrease of the informativeness of the VXX, and an increase in the informativeness of the XIV. These results could be explained by the increased hedging demand for the VXX during times of high levels of fear, which is not driven by information, but expresses a reaction of market

[9] We additionally run a model including a dummy variable controlling for the reverse splits of the VXX. Results stay nearly the same as reported, confirming that abrupt changes in the IS and GIS on days with reverse splits can be explained by the decrease in relative spread.

Table 3.4: Determinants of PD measures

Panel A: Determinants of IS				
	Model 1	Model 2	Model 3	Model 4
Intercept	0.000	0.000	0.000	0.000
	(0.093)	(-0.057)	(0.119)	(-0.024)
ΔIS_{t-1}	-0.452***	-0.450***	-0.451***	-0.448***
	(-13.109)	(-12.958)	(-13.096)	(-12.954)
$\Delta Spread_t$	-0.202***	-0.195***	-0.191***	-0.184***
	(-8.908)	(-8.760)	(-8.622)	(-8.524)
ΔTrades $Volume_t$	0.009***	0.009***		
	(2.579)	(2.624)		
ΔTrades $Number_t$			0.010***	0.010***
			(2.780)	(2.857)
ΔVIX_t	-0.007***		-0.006**	
	(-2.663)		(-2.484)	
ΔSPX_t		1.041***		0.987***
		(2.794)		(2.658)
Adjusted R^2 (%)	29.4	29.3	29.63	29.57
No. of Obs.	912	912	912	912
Panel B: Determinants of GIS				
Intercept	0.000	0.000	0.001	0.000
	(-0.008)	(-0.121)	(0.013)	(-0.094)
ΔGIS_{t-1}	-0.462***	-0.459***	-0.461***	-0.458***
	(-12.283)	(-12.210)	(-12.300)	(-12.232)
$\Delta Spread_t$	-0.277***	-0.270***	-0.263***	-0.257***
	(-8.148)	(-8.158)	(-7.679)	(-7.739)
ΔTrades $Volume_t$	0.012**	0.012**		
	(2.459)	(2.488)		
ΔTrades $Number_t$			0.013***	0.013***
			(2.613)	(2.684)
ΔVIX_t	-0.007*		-0.007	
	(-1.680)		(-1.543)	
ΔSPX_t		1.181**		1.124*
		(1.963)		(1.860)
Adjusted R^2 (%)	28.57	28.57	28.68	28.69
No. of Obs.	912	912	912	912

Note: This table reports regression results of the daily changes in the PD measures of the VXX on various variables described in Equation (3.11). Robust t-statistics is reported in parenthesis and ***, ** and * are used to indicate significance at the 1%, 5% and 10% respectively.

participants to the deterioration of the market conditions. Prior studies have documented the strong negative relation between changes in the VIX and returns on the S&P 500 index (e.g. Whaley, 2009), and, indeed, for our sample the correlation between those two variables is -0.83. This can explain that while the return on the

S&P 500 index has a positive association with the informativeness on the VXX, the change in the VIX has a negative association.

As for the daily changes in the GIS, the results are of a similar kind, with the lowest adjusted R-squared of 28.57%. In summary, we confirm that informed traders prefer to trade in the more liquid market and the one with the lowest transaction costs. Additionally, during market turmoils, the informational leadership of the VXX subsides.

3.6 Conclusions

In this study, we examine the intraday price discovery relation between the VIX short-term futures ETN and inverse VIX short term futures ETN. Using the approaches of Hasbrouck (1995) and Lien and Shrestha (2014), we first conduct the unit root test, cointegration analysis and build the VECM. The Johansen approach suggest that the two series are cointegrated, thus a long-run equilibrium relationship exists. We find that both the information share and the generalized information share are subject to substantial time variation with the latter being more volatile. We conduct a times series regression between the daily changes in the price discovery measures and several possible determinants. We find that price leadership is associated with relative trading costs, relative trading volume and number of trades. The informativeness of the VXX tends to increase, on average, when the relative number of trades (or the trading volume) in the VXX is high and when the XIV number of trades (or the trading volume) is low, when the effective bid-ask spread in the VXX narrows and in the XIV widens. We further document that the informativeness of the VXX market decreases on days when expected future market volatility increases and on days with negative returns in the stock market. This finding suggest that during market downturns informed investors use the XIV to trade on their information, expressing mean-reversion expectations on future volatility, while a trading activity in the VXX market is more driven by an increased hedging demand.

3.A Appendix: Calculation of the CIV for the VXX and XIV

The total (SPVXSTR) and excess (SPVXSP) return versions of the indices are calculated by the following formulas:

$$IndexTR_t = IndexTR_{t-1} \times (1 + CDR_t + TBR_t) \tag{3.12}$$

$$IndexER_t = IndexER_{t-1} \times (1 + CDR_t), \tag{3.13}$$

where CDR_t is the Contract Daily Return and TBR_t is the Treasury Bill Return earned on the notional value of the position. The CDR_t and the TBR_t are given by the formulas:

$$CDR_t = \frac{TDWO_t}{TDWI_{t-1}} - 1 \tag{3.14}$$

$$TBR_t = \left[\frac{1}{1 - \frac{91}{360} \times TBAR_{t-1}}\right]^{\frac{Delta_t}{91}} - 1, \tag{3.15}$$

where $TDWO_t$ is the Total Dollar Weight Obtained on t and $TDWI_{t-1}$ is the Total Dollar Weight Invested on $t-1$, as determined by the following formulas:

$$TDWO_t = \sum_{i=1}^{2} CRW_{i,t-1} * DCRP_{i,t} \tag{3.16}$$

$$TDWI_{t-1} = \sum_{i=1}^{2} CRW_{i,t-1} * DCRP_{i,t-1}, \tag{3.17}$$

where $CRW_{i,t}$ is the Contract Roll Weight of the ith VIX Futures Contract on date t and $DCRP_{i,t}$ is the Daily Contract Reference Price of the ith VIX Futures Contract on date t. $Delta_t$ is the number of calendar days between the current and previous business day and $TBAR_{t-1}$ is the most recent weekly high discount rate for 91-day US Treasury bills effective on the preceding business day. Inspection of Equations (3.16) and (3.17) shows that Contract Daily Return is driven only by the changes in the VIX futures prices and is not dependent on changes in the weights.

At the start of the roll period, all weight is allocated to the first-nearby contract. Then, on each subsequent business day a fraction of the first-nearby VIX futures holding is sold and an equal notional amount of the second-nearby VIX futures is bought. The initial position in the first-nearby contract is progressively rolled to the second-nearby futures contract during the course of the month, until the following roll period starts when the old second-nearby VIX futures contract becomes the new first-nearby VIX futures contract. After that the process repeats.

The closing indicative value is linked to the performance of the underlying index minus an investor fee. On any calendar date the CIV for the VXX is calculated based on the following equations:

$$CIV_t^{VXX} = CIV_{t-1}^{VXX} \times DIF_t - Fee_t \tag{3.18}$$

$$DIF_t = \frac{IndexTR_t}{IndexTR_{t-1}} \tag{3.19}$$

$$Fee_t = \frac{feeRate}{365} \times CIV_{t-1}^{VXX} \times DIF_t, \tag{3.20}$$

where CIV_t^{VXX} is the closing indicative value of the VXX on any given calendar day t, DIF_t is the daily index factor, $IndexTR_t$ is the closing level of the index, $feeRate$ is the investor fee rate which is equal to 0.89% per year.

Combining Equations (3.12), (3.18), (3.19) and (3.20), we arrive at the final formula for the CIV for the VXX:

$$CIV_t^{VXX} = CIV_{t-1}^{VXX} \times (1 + TBR_t + CDR_t) \times (1 - \frac{feeRate}{365}).$$

The closing indicative value for the series of the XIV is equal to:

$$CIV_t^{XIV} = CIV_{t-1}^{XIV} \times DETNP_t - DIF_t \tag{3.21}$$

$$DETNP_t = 1 + TBR_t + DIP_t \times (-1) \tag{3.22}$$

$$DIF_t = CIV_{t-1}^{XIV} \times DETNP_t \times \frac{feeRate}{365} \tag{3.23}$$

$$DIP_t = \frac{IndexER_t}{IndexER_{t-1}} - 1, \tag{3.24}$$

where $DETNP_t$ is the daily ETN performance, DIF_t is the daily investor fee, DIP_t

is the daily index performance, $feeRate$ is equal to 1.35% for the inverse ETN. The daily index performance is adjusted by the leverage amount -1.

Combining Equations (3.13) and (3.21)-(3.24) the CIV for the XIV series is equal to:

$$CIV_t^{XIV} = CIV_{t-1}^{XIV} \times (1 + TBR_t - CDR_t) \times (1 - \frac{feeRate}{365}).$$

Chapter 4

Properties and the Predictive Power of Implied Volatility in the Dairy Market

4.1 Introduction

New Zealand (NZ) is the world's largest exporter of dairy commodities, representing approximately one third of international dairy trade each year. Almost half of global Whole Milk Powder (WMP) exports are sourced from NZ, making dairy products one of the most important agricultural commodities for NZ.[1] Dairy products are also known to display high levels of price volatility. For instance, Figure 4.1 shows that the price of NZ milk powder exhibits considerable fluctuations over time, suggesting high volatility. This high volatility can have a significant impact on the ability of NZ dairy farmers to manage their operations and service debt, and can have serious consequences for the health of NZ's most influential agricultural sector. In this chapter, we estimate the implied volatility of WMP and examine its predictive power for future realized volatility of WMP futures.

Milk powder is one of the most volatile commodities globally.[2] Though it is known that historical volatility is a good predictor of future volatility, in this study

[1] In 2017 dairy products contributed $NZ14.6 billion to annual exports. The top three dairy export products are: whole milk powder (36%), butter (19%), cheese (13%) `https://www.mpi.govt.nz/news-and-resources/open-data-and-forecasting/situation-and-outlook-for-primary-industries-data/`

[2] Figure A1 depicts volatility of different commodities and the S&P 500. For the period between December 2011 and January 2018 the annualized standard deviation of weekly returns for WMP futures is 37.1%, while for the S&P 500, CRB Crude Oil, Gold, Cocoa and Sugar Indices it is equal to 12%, 34.5%, 23.9%, 31.3% and 33.6%, respectively.

Figure 4.1: WMP GDT Auction Prices

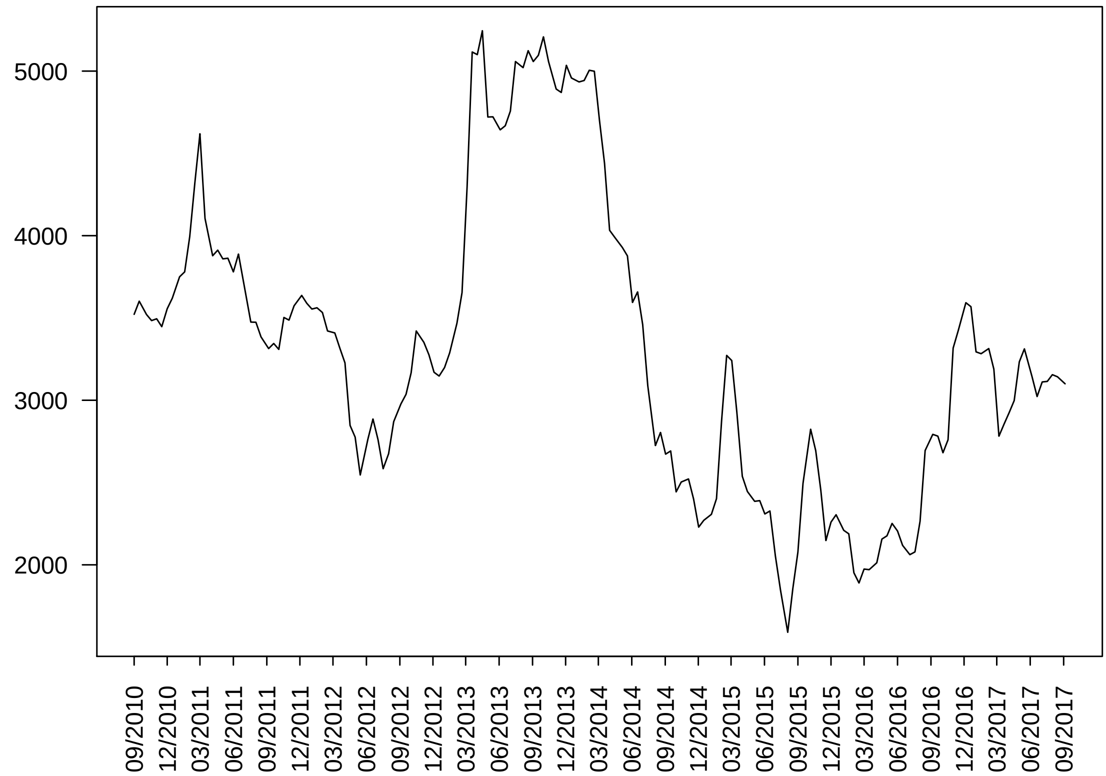

Note: This figure shows GDT auction weighted average price for all contracts of the WMP for the period between September 1, 2009 and September 5, 2017. The price is in USD per metric tonne.

we aim to improve this forecast. Ample evidence suggests that volatility forecasts based on option implied volatility outperforms those that use historical information (see Jorion, 1995; Szakmary et al., 2003; Fleming, 1998; Blair et al., 2001; Triantafyllou et al., 2015, among others). Implied volatility is obtained by inverting an option pricing formula and is often interpreted as the expected volatility over the life of an option. Implied volatility is generally considered to be a superior predictor of future volatility due to the ability of market participants to effectively incorporate all publicly available information as well as any additional information that is relevant for predicting volatility into security prices. By construction, implied volatility is a forward-looking measure, as opposed to historical volatility which relies on historical data. Over the last three decades, extensive research for financial and non-financial products has examined the information content of implied volatility (for an overview, see Poon and Granger, 2003; Gonzalez-Perez, 2015). It has been shown empirically that implied volatility produces superior forecasts of future volatility across different asset classes. For example, Giot (2003) finds that past squared returns (i.e. GARCH effects) provide no significant volatility information in addition to the lagged implied volatility for cocoa and sugar futures. Triantafyllou et al. (2015) compute the implied variance for US wheat, corn and soybean futures. They find that, for corn and wheat, option-implied variance outperforms historical volatility in forecasting volatility, while, for soybeans, both historical and implied variance are significant determinants of future realized variance.

This chapter provides a comprehensive analysis of the information contained in the NZ dairy option market in predicting subsequent realized volatility and our contribution is threefold. First, we construct a dairy volatility index (DVIX) and analyse its statistical properties. The DVIX is a 22-trading day at-the-money implied volatility index constructed from four call and four put options. We find a significant negative and asymmetric relationship between one-day lagged returns and the changes in the DVIX. It means that, at times when returns are positive (negative), the DVIX drops (rises) the next day. Additionally, the next day change in the DVIX is larger when returns are positive, rather than negative. Second, we assess the in-sample forecasting performance of the DVIX in a GARCH-type framework. The results strongly suggest that the DVIX has a high information content regarding conditional variance and that the historical information further improves the model's fit. Third, we conduct an out-of-sample evaluation of the forecast performance of implied volatility, where we consider 1-, 5-, 10- and 22-day ahead forecast horizons. The predictive power of implied volatility is assessed against four alternative time-series forecasts: historical volatility realized during the past 30 trading days and three different GARCH-type forecasts. We find that the DVIX provides substantial information about future realized volatility, although it is not an unbiased estimate of future volatility. We also document that a combination of historical volatility and the DVIX provides the best forecast accuracy for all forecast horizons. Results of Clark and West (2007) test suggest that the inclusion of the DVIX in the predictive regressions improves the out-of-sample performance of all considered time-series forecasters.

To our knowledge, we are the first to construct and examine the predictive power of implied volatility for the NZ dairy market. The construction of the DVIX allows us to visualise the volatility of an important NZ dairy product and highlights the potential need for risk management tools for farmers and manufacturers. The results of this chapter are particularly important for decision makers in the financial and agricultural sectors who require a volatility estimate as an input for pricing and risk management.

The remainder of this chapter is structured as follows: Section 4.2 provides an overview on the measures of volatility and highlights the importance of implied volatility as a predictor of future volatility. Section 4.3 presents the data and the methodology of the DVIX construction. Section 4.4 explores the statistical properties of the DVIX and conducts in- and out-of-sample forecasting tests. Section 4.5 summarises the main results of the chapter.

4.2 Literature Review

4.2.1 Historical Versus Forward-Looking Volatility

There are two main types of volatility which are used to describe fluctuations of an asset's price. The first one is historical, or backward-looking volatility, and the second one is implied or forward-looking volatility.

The simplest measure of historical volatility is the standard deviation of a set of past observations. A more sophisticated type of time-series models is presented by ARCH (Autoregressive Conditional Heteroscedasticity) and GARCH (Generalized ARCH) models introduced by Engle (1982) and Bollerslev (1986), respectively. GARCH models have been developed to account for "stylised facts" documented in financial return series, such as volatility clustering (absolute or squared returns display a positive autocorrelation over several days), excess kurtosis (the distribution of returns displays heavy tails) and leverage effects (negative stock market returns are associated with changes in volatility that are much larger than those associated with positive returns of similar size). In these models, the variance of residuals is not constant and the next period variance is conditional on information this period.

The implied volatility of an asset can be obtained from the market prices of options written on that asset. Option implied volatility is interpreted as the market's expectation about the volatility of an asset over the life of the option. The first volatility index (VIX) was introduced by the CBOE in 1993 and was computed using the implied volatilities from eight near-the-money options written on the S&P 100. In 2003, the CBOE, together with Goldman Sachs, revised their methodology. They incorporated out-of-the-money put and call options over a wide range of strike prices and replaced the S&P 100 options with S&P 500 options. The VIX gained market acceptance and later the CBOE expanded its range of volatility indexes for different stock indexes, interest rates, currency futures, ETFs and single stocks.[3] VIX-type indexes on commodity assets that trade as ETFs include crude oil, gold, silver and energy. In the context of our chapter, the most interesting is the existence of the volatility indexes on agricultural commodities. In June 2011, the CBOE began disseminating volatility benchmarks based on CME Group corn and soybean option prices. In July 2012, the CBOE started disseminating volatility benchmark based on CME Group wheat options. The corn, soybean and wheat volatility indexes are based on the same updated methodology developed by the CBOE for the U.S. equity based VIX.

Shortly after its introduction, the VIX became a benchmark of expected short-term market volatility and is known as the "investor fear gauge" due to its property of spiking at times of market turmoil (Whaley, 2009). Another empirical observation is that the return-VIX relation is asymmetric, meaning that negative stock market moves are associated with much larger moves in the VIX than those associated with positive stock market moves of similar size (Fleming et al., 1995; Whaley, 2009).[4]

[3] The full list of all volatility indexes is available at `http://www.cboe.com/products/vix-index-volatility/volatility-indexes`.

[4] Traditionally, there are two hypotheses which explain negative and asymmetric return-volatility relation. First, Black (1976) and Christie (1982) develop the leverage hypothesis, which argues that when the stock price of a firm declines, the firm's debt-to-equity ratio increases, which makes the firm riskier and increases the volatility of its equity as a result. Second, Campbell and Hentschel (1992) and French et al. (1987) propose the volatility feedback hypothesis. If volatility is priced, an anticipated increase in volatility would raise the required rate of return, which leads to a current stock price decline. Contrary to these fundamental arguments, Hibbert et al. (2008) propose a behavioral approach, which relies on concepts of representativeness, affect, and extrapolation bias.

Padungsaksawasdi and Daigler (2014) examine the return-VIX relation between the commodity ETF price changes (gold and oil) and their associated volatility indexes and document several interesting results. First, they find that co-movement between the price changes of commodity ETFs and their respective VIX changes are much weaker than documented for stock indexes. Second, they find a positive contemporaneous return-VIX relation for gold. They explain it by relating to a "safe-haven" feature of gold during financial turmoil (gold exhibits positive response to negative macroeconomic news). The results highlight differences of the return-VIX relation among different asset classes.[5]

4.2.2 Forecasting of Volatility

The widely-used approach to forecasting volatility is to use time-series models (historical volatility or GARCH models) or to use implied volatilities from options (for an overview of models used in volatility forecasting see Poon and Granger, 2003). While some empirical studies document that GARCH models produce good volatility forecasts over short periods (Andersen and Bollerslev, 1998), there is little evidence to suggest that GARCH models outperform option-implied forecasts of future volatility. Empirical evidence largely shows that implied volatility produces superior forecasts of future volatility across different asset classes. For example, Jorion (1995) examines currency markets and shows that option-implied forecasts outperform Moving Average and GARCH models. Fleming (1998) uses S&P 100 options to compute implied volatilities and shows that the predictive power of implied volatility is superior to GARCH(1, 1) and historical volatility.[6] Blair et al. (2001) follow Fleming et al. (1995) and use the VIX as a more accurate measure of implied market volatility from S&P 100 options. They compare the out-of-sample accuracy of four different volatility forecasts for 1, 5, 10 and 20 days ahead. The four models are the historic volatility, the daily-frequency ARCH forecast, the intraday volatility

[5] Baur and Dimpfl (2018) find a positive (inverted) asymmetric effect in agricultural commodities (softs, grains and livestock), energy commodities and metals (industrial and precious metals) with a tendency to weaken and converge towards an equity-like effect.

[6] The call (put) implied volatility is estimated from all call (put) option transactions within a 10-min window centered around the stock market close via an averaging technique. Fleming (1998) chooses Fleming and Whaley (1994) modified binomial model as an option pricing model.

calculated from 5-min and overnight returns and the implied volatility index, VIX. Results show that the VIX contains all the relevant forecasting information for the forecast horizons of 1, 5, 10 and 20 days ahead.

Another strand of research considers implied volatilities from options written on non-financial assets. Giot (2003) focuses on implied volatilities for agricultural commodities (cocoa, coffee, and sugar futures contracts). In-sample analysis shows that past squared returns (i.e., GARCH effects) provide no volatility information in addition to the lagged implied volatility for the cocoa futures contracts. Triantafyllou et al. (2015) compute model-free implied variance for US wheat, corn and soybean futures markets. Along with macroeconomic data they use historical 2-month realized variance, model-free implied variance, model-free implied skewness to explain 2-month ahead realized variance. They find that for corn and wheat past realized variance does not predict future realized variance, while for soybeans both historical and implied variance are significant determinants of future variance. Szakmary et al. (2003) examine a variety of asset classes, such as futures options on equity indexes, currencies, crude oil, short- and long-term interest rates, agricultural commodities, livestock, metals, refined petroleum products, and natural gas. They find that regardless of how historical measures of volatility are modelled (simple 30-day moving average or GARCH), implied volatility outperforms historical volatility in predicting future realized volatility. Manfredo and Sanders (2004) examine the performance of implied volatility derived from live cattle options to forecast one-week volatility of live cattle futures prices. They find that implied volatility outperforms alternative forecasters and that it has improved its forecasting quality over time.

4.3 Data and the Dairy Volatility Index Calculation

In this section we describe the NZX dairy derivatives market, as well as the data used to construct the DVIX. After discussing the methodology used to construct the DVIX, we present some summary statistics.

4.3.1 Background

In 2010, several exchanges around the world launched dairy derivatives to help the industry hedge against dairy price volatility. In May 2010, the Chicago Mercantile Exchange was the first to launch International Skimmed Milk Powder futures and options contracts with physical delivery points located around the world. In the same month, the Frankfurt-based Eurex launched trading in cash-settled futures on Skimmed Milk Powder and Butter. The NZX launched its first derivative in the agricultural asset class, a Whole Milk Powder (WMP) futures contract, in October 2010. It was followed by Skim Milk Powder (SMP) and Anhydrous Milk Fat (AMF) futures in February 2011. In November 2011, NZX launched WMP options contracts and in December 2014 added Butter (BTR) futures to its derivative product offering. Then followed Milk Price (MKP) futures and options, launched in May 2016 and June 2016, respectively. The most recent dairy derivatives are SMP Option contracts, launched in December 2017. The NZX futures contracts are quoted in US dollars, with one contract representing one tonne of product. NZX Dairy futures contracts are cash-settled rather than physically delivered. One WMP options contract represents the right to buy or sell one WMP futures contract and is also quoted in US dollars. Table A1 summarises contracts specifications of all eight currently available NZX dairy derivatives.

Figure 4.2 depicts aggregated trading volume for each month for all available dairy derivatives traded at the NZX. It shows that the most actively traded dairy derivatives are WMP futures, and the least traded are recently launched MKP options. Across all eight dairy derivatives, the NZX recorded a trading volume of nearly 26,000 lots as of December 2017, where trading in WMP futures accounts for nearly 79%. Since inception, WMP futures experienced a growth in the trading volume and in December 2017 it amounted to about 20,000 lots. WMP options are less actively traded. The first trade occurred in June 2014, which is nearly two and a half years after their launch, and in December 2017 the trading volume amounted to about 1,000 lots.

Figure 4.2: NZX Dairy Derivatives Volume

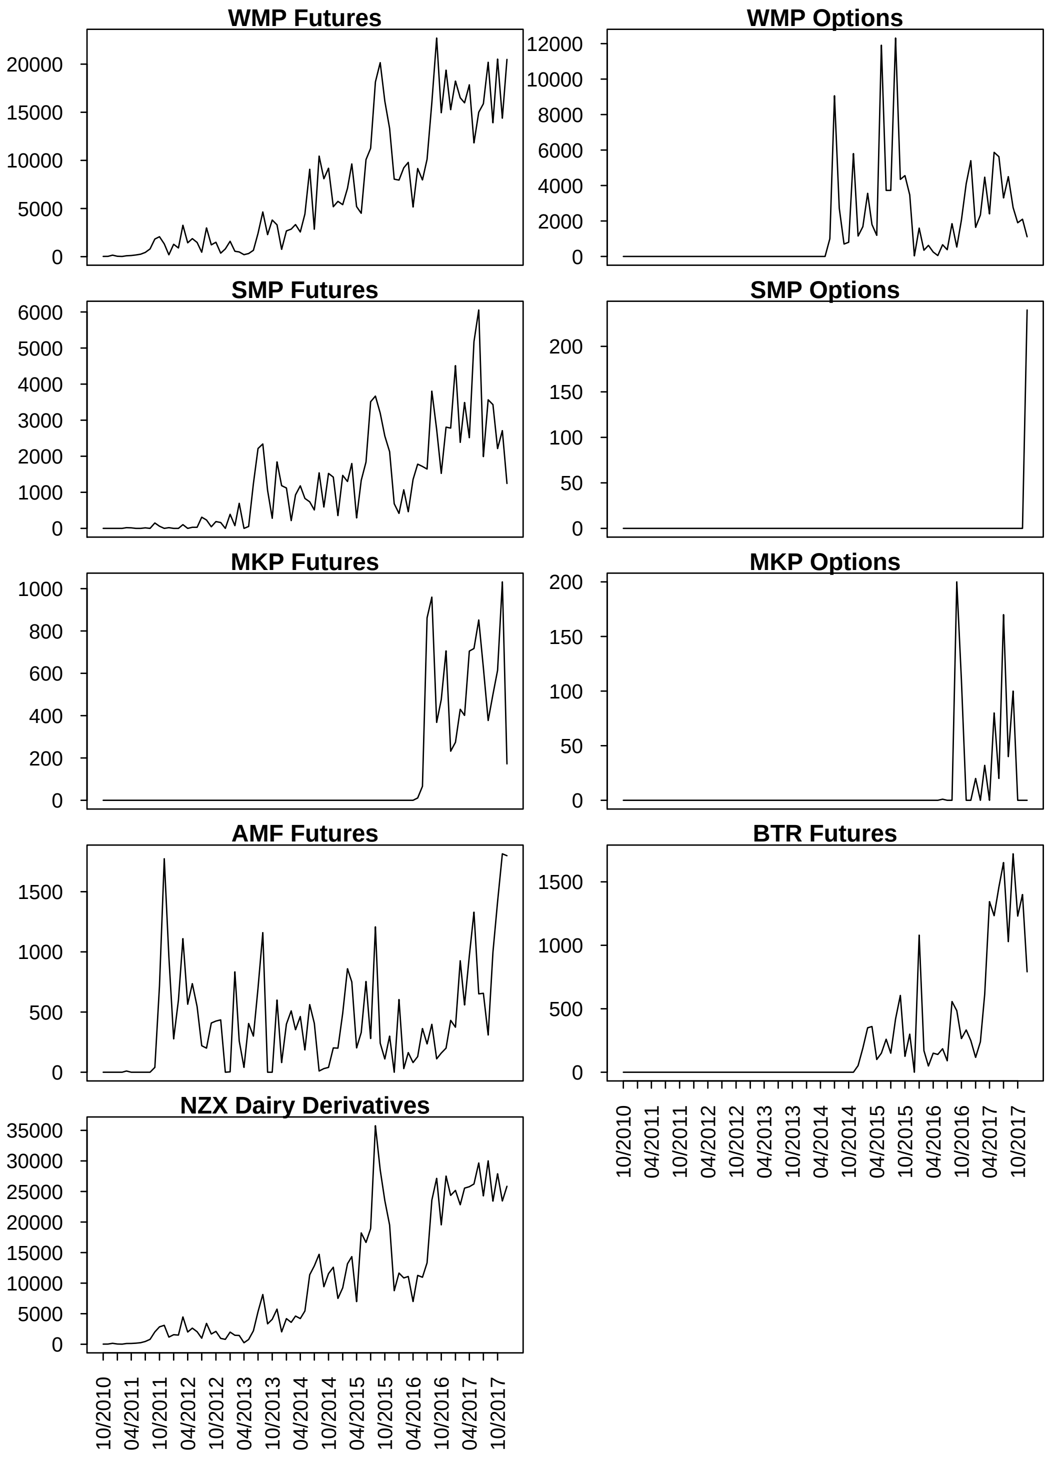

Note: This figure shows the monthly trading volume for all NZX Dairy Derivatives for the period between October 8, 2010 and December 29, 2017.

4.3.2 Data

We use daily data on the NZX WMP futures and options contracts which we obtain from the NZX Research Centre[7]. One of its services, AGRI DATA database, contains

[7]`https://companyresearch.nzx.com/crust/services.php`

information about four sets of variables: trading levels and trading prices on the NZX futures and options; Global Dairy Trade (GDT) auction prices;[8] NZ milk production statistics; pastoral growth index for regions in NZ and on a national level. We collect option and futures daily settlement prices which we use to construct the dairy volatility index. The sample period starts on 30 November 2011 (the first day WMP options were traded on the NZX) and ends 8 January 2018. As a proxy for the risk-free rate, we collect USD Overnight Index Swap (OIS) rates for various maturities available from DataStream (we use a US risk-free rate as the WMP futures and options are settled in US dollars). The estimation period of in-sample and out-of-sample analysis includes the period from 5 January 2015 to 8 January 2018.[9]

4.3.3 Dairy Implied Volatility Computation

We compute the DVIX by closely following the methodology applied to the original CBOE VIX.[10] To construct the DVIX we need three types of information: 1) an option valuation model; 2) the values of the model's determinants (except for volatility); 3) an observed option price. For a given option price, inverting the pricing model yields the implied volatility of that option. We construct the DVIX from eight options, four calls and four puts, written on the WMP futures. The DVIX is constructed in a way that it is at-the-money and has a constant 30 calendar days (22 trading days) to expiration.[11] To achieve the at-the-money implied volatility we combine in- and out-of the money options, and to achieve a constant time to expiry we combine the first and second nearby options. As option prices can be

[8] GDT is the NZ-based spot market for various dairy products. The NZX WMP futures settle against the GDT spot prices.

[9] As subsequent analysis will show, the DVIX values are virtually constant during the years 2011 to 2014, which is largely due to a lack of liquidity in the market. The first trade in WMP options occurred in June 2014. During the period 2011-2014, 21,235 lots were trades; while during 2015-2017, 108,173 lots were traded. Hence, we perform the in- and out-of-sample evaluation starting from 2015 onwards.

[10] The new approach to calculate the VIX requires the availability of many out-of-the-money options across a full range of strike prices. For the WMP options there are only a few contracts that are actively quoted and traded. Thus we choose to follow the original CBOE VIX methodology which relies on only eight near-the-money options at the two nearest maturities.

[11] Thus, our approach is similar to Whaley (1993); however, other time horizons could be considered (for example, two- or three-month time horizons).

very volatile when approaching the expiration date, we use options which have at least eight trading days prior to expiration. To address the early exercise feature of American-style options we use quadratic approximation of American option values proposed by Barone-Adesi and Whaley (1987), which is explained in Appendix 4.B. A detailed explanation of the construction of the DVIX is provided in Appendix 4.C.

Figure 4.3 plots the DVIX over time. We notice a considerable drop in the DVIX value between 17 August 2016 and 23 August 2016 - a trading week from Wednesday to Tuesday (5 trading days). The value plummeted from nearly 40% to 14.5% and reached its minimum of 7.5% on Tuesday. This rapid drop in the DVIX is likely to be caused by the positive news on the price of the WMP, which jumped 18.9% in the GDT auction.[12] Average dairy prices increased by 12.7%. The highest value of 78% and 78.9% is achieved on 7 and 8 September 2015. Although we cannot identify any relevant news in NZ for this day, during that time in mainland Europe farmers protested against falling dairy and meat prices, which may have generated some uncertainty about global milk prices.[13] From Figure 4.3 we observe that the DVIX has little variation in the beginning of the sample period, which is due to the low activity in WMP options, and fluctuates around the value of 49%. Thus, we exclude the period from 30 November 2011 till 31 December 2014 for further analysis. Another interesting observation is a downward trend, indicating a decreasing value of the DVIX. For subsequent analysis it is important to check the stationarity of the DVIX. The daily DVIX level shows a high persistence with a first-order autocorrelation of 0.815. For example, the US Equity VIX shows a first-order autocorrelation of 0.866 during the same time period, indicating similarity between two volatility indices. The Augmented Dickey-Fuller test rejects the null hypothesis on non-stationarity at 1% level. This finding supports the stylised fact that volatility in the dairy market is a persistent but mean-reverting process.

We calculate summary statistics for daily DVIX changes over the period 5 Jan-

[12] `http://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=11694758`

[13] `http://money.cnn.com/2015/09/07/news/economy/europe-milk-prices-protest/index.html` `https://www.theguardian.com/environment/2015/sep/07/farmers-clash-police-brussels-milk-meat-prices-protest`

Figure 4.3: Dairy VIX

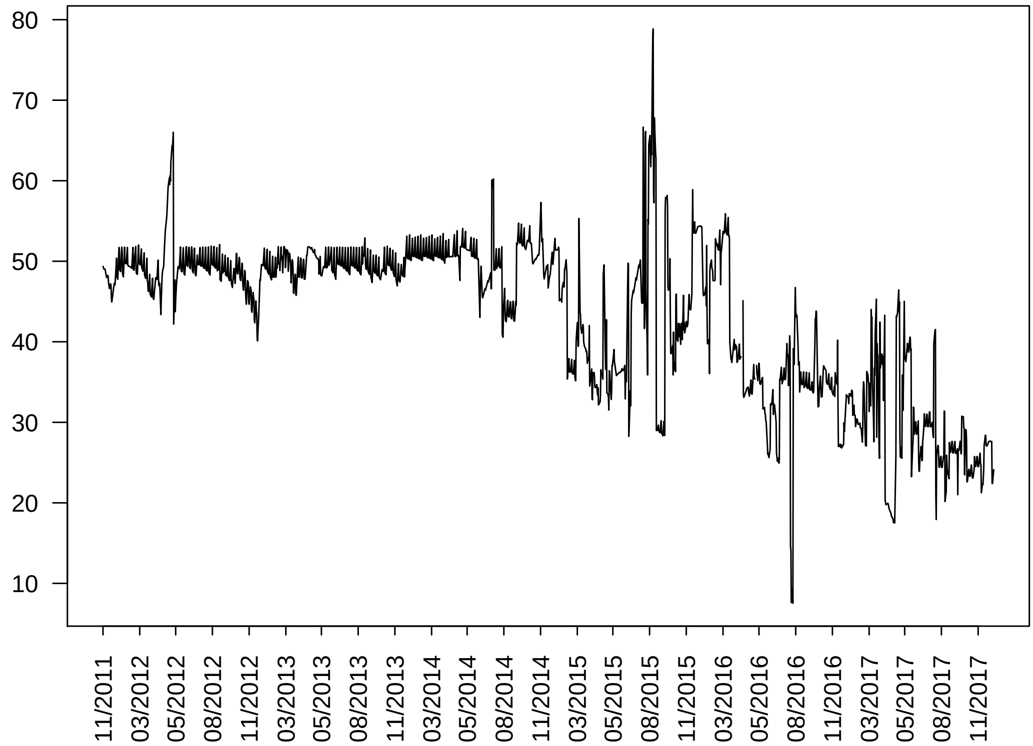

Note: This figure plots the DVIX levels in percentages during the period from 30 November 2011 through 8 January 2018.

uary 2015 to 8 January 2018 and for three sub-periods. The choice of the sub-periods is dictated by the observation on the WMP options volume from Figure 4.2. One can notice a drop in the WMP options volume during the period December 2015 till November 2016. Table 4.1 presents the summary statistics. The mean daily change and the standard deviation remain relatively stable for the three sub-periods. The daily change in the DVIX exhibits positive excess kurtosis, meaning that distribution has fatter tails than a Normal distribution. Skewness varies from negative to positive values, showing that the DVIX has a tendency to both abrupt decreases and increases. The first-order autocorrelation ranges from -0.093 to -0.159. We can conclude that statistical properties of the DVIX exhibit some variation, not significant, however, and we can consider the whole period, starting at 5 January 2015, for the subsequent analysis.

Table 4.1: Descriptive Statistics

	2015/01/05 2018/01/08	2015/01/05 2015/11/30	2015/12/03 2016/10/03	2016/10/10 2018/01/08
Mean ($\times 10$)	-0.004	-0.005	-0.004	-0.004
Median ($\times 10$)	-0.004	-0.002	-0.002	-0.004
Max	0.316	0.279	0.316	0.173
Min	-0.338	-0.338	-0.250	-0.230
St. Dev.	0.046	0.058	0.040	0.040
Skewness	-0.318	-0.489	1.316	-0.870
Kurtosis	15.511	11.076	27.452	10.238
$\rho(1)$	-0.123**	-0.111	-0.093	-0.159**
$\rho(2)$	-0.110**	-0.176**	-0.006	-0.073
$\rho(3)$	-0.075**	-0.115	0.060	-0.099
No. of Obs.	749	230	203	316

Note: This table reports descriptive statistics on the daily DVIX level changes for the full sample and for three sub-samples. ** indicates significance at 5% level.

4.4 The Information Content of Implied Volatility

In this section we test for the presence of seasonal patterns, and examine the intertemporal relationship between WMP futures returns and the DVIX. We also investigate whether the DVIX contains information for future WMP futures volatility in- and out-of-sample.

4.4.1 Statistical Properties of the DVIX

We start by examining whether there are seasonalities in the DVIX. There is quite some evidence on seasonalities in financial time series. For instance, the day-of-the-week effect in the US stock market was first documented by Cross (1973), according to which the mean return between close of Friday and close of Monday is negative. As for the volatility index, Fleming et al. (1995) investigate patterns in the VIX and find that the VIX declines throughout the week. To test the day-of-the-week effect in the DVIX, we estimate the following regression:

$$VIX_t = \alpha + \beta_1 D_{1,t} + \beta_2 D_{2,t} + \beta_4 D_{4,t} + \beta_5 D_{5,t} + \epsilon_t, \tag{4.1}$$

where $D_{i,t}$ is a dummy variable for each day of the week ($D_{1,t}$ for Monday,..., $D_{5,t}$ for Friday). To avoid multicollinearity, we exclude the dummy variable for Wednesday. In Table 4.2, we report the results for the regression as well as Newey-West adjusted t-statistics in parentheses. We do not find any day-of-the-week effects, as all the coefficients on dummy variables are insignificant at any conventional significance level.

Table 4.2: Calendar Anomalies in the Dairy VIX

Day-of-the-Week Effects			Month-of-the-Year Effects		
const	0.3664***	(27.733)	*const*	0.412***	(9.210)
Monday	0.0011	(0.205)	*Feb*	-0.011	(-0.195)
Tuesday	-0.0032	(-0.064)	*Mar*	0.002	(0.053)
Thursday	-0.0022	(-1.060)	*Apr*	-0.084	(-1.600)
Friday	-0.0058	(-1.259)	*May*	-0.061	(-1.214)
$R^2(adj.)$	-0.0047		*Jun*	-0.076	(-1.585)
			Jul	-0.069	(-1.221)
			Aug	-0.021	(-0.376)
			Sept	-0.060	(-0.866)
			Oct	-0.060	(-0.967)
			Nov	-0.077	(-1.315)
			Dec	-0.054	(-0.735)
			$R^2(adj.)$	0.089	

Note: This table presents parameter estimates for the regression of the DVIX on day-of-the-week and month-of-the-year dummy variables. ***, **, * is used to indicate significance at the 1%, 5% and 10% levels.

The second interesting pattern to investigate is a month-of-the-year seasonality. Shadbolt and Apparao (2016) notice that milk production in New Zealand is cyclical and driven by the availability of pasture, which, to a large extent, is determined by rainfall. The milking season starts in August and ends in May. The milk production curve is the lowest during June and July ("winter milk"), and the end of the season, April and May. As a response to such variation, processing plants make long-life products, such as powders, cheeses and whey products at the peak of the production curve. Since the supply of milk may affect the DVIX, we investigate monthly patterns in the DVIX. To analyse the month-of-the-year effect we consider the following regression:

$$VIX_t = \alpha + \sum_{i=2}^{12} \beta_i D_{i,t} + \epsilon_t, \tag{4.2}$$

where $D_{i,t}$ is a dummy variable for each month of the year. To avoid multicollinearity, we exclude the dummy variable for January. Results are presented in Table 4.2. We find that during April, May, June and July the regression coefficients have the smallest values (meaning lower values of the DVIX); however, they were not statistically significant. Therefore, there is no evidence of seasonalities in the DVIX.

Next, we want to run a predictive regression for the DVIX. In this test we want to asses whether the past historical information is informative for future values of the DVIX. We run several regressions with the most general specification of the following form:

$$\begin{aligned} DVIX_t = \alpha_0 + \alpha_1 DVIX_{t-1} + \alpha_2 DVIX_{t-2} + \alpha_3 DVIX_{t-3} + \\ \alpha_4 DVIX_{t-4} + \alpha_5 DVIX_{t-5} + \beta_1 HISTV_{t-1} + \epsilon_t, \end{aligned} \tag{4.3}$$

where $HISTV$ is the sample standard deviation of nearby futures returns over the previous 30 days. We include up to five lags of DVIX values, to consider the whole week. In the first model, we only use historical volatility as a predictor of the DVIX (α_1,...,$\alpha_5 = 0$). Next, we include only the lagged DVIXs to explain current value of the DVIX ($\beta_1 = 0$). Lastly, we include both the historical volatility and the lagged values of the DVIX. We report the results in Table 4.3. When we only use lagged historical volatility to predict the next day DVIX, we find that historical volatility is a significant determinant with the adjusted R^2 equal to 17.9%. The second models reveals that the DVIX is a persistent process with one-day lagged DVIX being positively associated with the next day DVIX, with the explanatory power being 80.6%. When we include both historical volatility and the lagged values of DVIX, the significance of historical volatility drops out, and R^2 remains almost unchanged at 80.7%. Thus we find no evidence that the previous day historical volatility affects the current value of the DVIX, when we control for the lagged values of the DVIX, meaning that the DVIX subsumes information contained in historical volatility.

Lastly, we analyse co-movements of the DVIX with WMP futures returns. Padungsaksawasdi and Daigler (2014) examine the return-volatility relation for the commodity returns, gold and oil. They find that commodity markets behave differently from the

Table 4.3: DVIX Predictive Regression

α_0	α_1	α_2	α_3	α_4	α_5	β_1	$R^2(adj.)$
0.282***						0.278***	0.179
(11.207)						(3.459)	
0.026***	0.789***	0.022	-0.005	0.048	0.064		0.806
(3.207)	(13.454)	(0.292)	(-0.064)	(1.016)	(1.104)		
0.026***	0.787***	0.020	-0.007	0.062	0.052	0.015	0.807
(3.190)	(14.063)	(0.288)	(-0.079)	(1.000)	(1.085)	(1.038)	

Note: This table reports regressions results of Equation (4.3).

stock market, which has the negative and asymmetric return-volatility relation. Following Fleming et al. (1995), Frijns et al. (2010) and Padungsaksawasdi and Daigler (2014) we estimate the following specification:

$$\begin{aligned}\Delta DVIX_t = \alpha_0 + \alpha_1 R_{t-2} + \alpha_2 R_{t-1} + \alpha_3 R_t + \alpha_4 R_{t+1} + \alpha_5 R_{t+2} + \\ \beta_1|R_{t-2}| + \beta_2|R_{t-1}| + \beta_3|R_t| + \beta_4|R_{t+1}| + \beta_5|R_{t+2}| + \epsilon_t,\end{aligned} \tag{4.4}$$

where $\Delta DVIX_t$ is the change in the DVIX from day $t-1$ to t, R_{t-2}, R_{t-1} are two- and one-day lagged returns, R_{t+1}, R_{t+2} are one- and two-day lead returns and R_t is a contemporaneous return. Return is based on the price of the nearby futures contract, F, and daily return is defined as $R_t \equiv ln(F_t) - ln(F_{t-1})$. We use daily log returns of the nearby futures, which have at least eight trading days prior to expiration. We compute daily futures returns always using two consecutive prices of the same contract, to avoid any effect which might result from rollover. The model aims to capture the intertemporal relationship between return and implied volatility and incorporates the possible asymmetric reaction of volatility to positive and negative moves in the WMP futures. The sum of $\alpha_i + \beta_i$ measures the asymmetry of the return-DVIX relationship. Table 4.4 reports the results. The results show that there is no contemporaneous relationship between return and the DVIX, but there is significant negative coefficient for the one-day lagged return and the DVIX. The absolute return variable on a one-day lagged return also turns out to be significant and negative, indicating the asymmetric effect in the return-volatility relation. The impact on the change in DVIX when the return is positive is equal to $\alpha_2+\beta_2$, or -0.4, whereas when the return is negative the impact is equal to $\alpha_2 - \beta_2$, or -0.042. This means that, if today the price of the WMP futures goes up, the next day the DVIX exhibits a decrease in value, while a drop in the WMP futures price is associated with an increase in the next-day option's implied volatility by a much smaller value. This finding supports the asymmetric effect between the return-volatility relation for the dairy market. We confirm the finding of Padungsaksawasdi and Daigler (2014) that for the commodity markets the degree of comovements between returns and associated volatility indexes is much weaker than for the stock markets. The adjusted R^2 for the WMP market is just 2.9%, while for the stock markets the adjusted R^2 is at least 50%.[14]

Table 4.4: Intertemporal Relationship between Daily DVIX Changes and WMP Futures Returns

$Intercept$	0.001	(0.303)
R_{t-2}	0.033	(0.322)
R_{t-1}	-0.221**	(-2.090)
R_t	0.156	(1.020)
R_{t+1}	-0.176	(-1.570)
R_{t+2}	0.036	(0.346)
$\|R_{t-2}\|$	0.005	(0.043)
$\|R_{t-1}\|$	-0.179*	(-1.647)
$\|R_t\|$	0.058	(0.366)
$\|R_{t+1}\|$	0.138	(1.130)
$\|R_{t+2}\|$	-0.115	(-1.016)
$R^2(adj.)$	0.029	

Note: This table reports the estimation results for the regressions described by Equation (4.4). ***, **, * is used to indicate significance at the 1%, 5% and 10% levels.

4.4.2 In-sample Volatility Forecasts

Empirical evidence suggests that, for several financial and non-financial assets, option implied volatility contains all relevant (including historical) information about future volatility. In this section, we evaluate three models which aim to asses whether the DVIX contains information about future volatility of the WMP futures.

Our in-sample models are similar to the models by Kroner et al. (1995) and Giot (2003), Blair et al. (2001). To compare the in-sample performance of several models

[14]Padungsaksawasdi and Daigler (2014) consider the period from August 2008 through March 2012 and report that for the S&P 500 - VIX pair the adjusted R^2 is 72.17%. Similarly, Fleming et al. (1995) consider the period from January 1986 through December 1992 and report the adjusted R^2 of 57.21%.

which use historical information from futures and implied volatilities, we employ GARCH-type models. We first specify the dynamics of the returns:

$$r_t = \mu + \epsilon_t,$$
$$\epsilon_t | \mathcal{F}_{t-1} \sim N(0, h_t),$$

where μ is an average return, $\mathcal{F}_{t-1}$ is the information set at time $t-1$ and ϵ_t is the error term at time t, which has a conditional Normal distribution with zero mean and variance h_t. We consider three specifications to model the variance equation. The first specification is a standard GARCH(1, 1) model and is defined by the following equation:

$$h_t = \omega + \alpha_1 \epsilon_{t-1}^2 + \alpha_2 h_{t-1}. \tag{4.5}$$

In Equation (4.5) the conditional variance ia a function of a constant term, the squared error term from the last period and the last period conditional variance.

The second specification is an augmented GARCH(1, 1) model with the DVIX, which we call the GARCH(1,1)-DVIX. It is defined by Equation (4.6) and uses both historical information and forward-looking information from the option market:

$$h_t = \omega + \alpha_1 \epsilon_{t-1}^2 + \alpha_2 h_{t-1} + \beta_1 DVIX_{t-1}^2, \tag{4.6}$$

where $DVIX_{t-1}$ is the daily implied volatility computed from the $DVIX$ as $DVIX/\sqrt{252}$.

The third model ignores historical information and uses only information from the option market, that is we impose the restrictions $\alpha_1 = \alpha_2 = 0$ in the augmented GARCH(1,1) model. We refer to this as the GARCH(0, 0) - DVIX model and it is defined by Equation (4.7):

$$h_t = \omega + \beta_1 DVIX_{t-1}^2. \tag{4.7}$$

We estimate the parameters of the models by maximising the log-likelihood functions, with the constraints $\omega \geq 0$, $\alpha_1 \geq 0, \alpha_2 \geq 0, \beta_1 \geq 0$, which ensure the nonnegativity of the conditional variance process. We report the results in Table 4.5. In the GARCH(1, 1) specification the coefficient α_2 is close to 1, which is consistent with the documented persistence of volatility. Including the $DVIX_{t-1}$ in the variance equation leads to the decrease in the value of α_2, as the $DVIX_{t-1}$ takes into account a part of the GARCH effect. The GARCH(1, 1)-DVIX model has the highest log-likelihood of 1920.97, then follows the GARCH(0, 0)-DVIX with the log-likelihood of 1891.8 and the log-likelihood of GARCH(1, 1) is 1885.76. To test for the added value of the implied volatility we use the likelihood ratio (LR) test, calculated as twice the difference between the log-likelihoods. For the GARCH(1, 1) versus the GARCH(1, 1)-DVIX, results suggest that we can reject the null hypothesis of $\beta_1 = 0$ at the 1% level (the LR statistics is 35.21), suggesting that the augmented GARCH(1, 1)-DVIX model outperforms the GARCH(1, 1). Next, we compare the models GARCH(1, 1)-DVIX and GARCH(0, 0)-DVIX, the test statistic is equal to 29.17, thus we reject the null hypothesis that the past information does not add significant variance information at the 1% level. The estimation results show that the augmented GARCH(1, 1)-DVIX model performs significantly better than the standard GARCH(1, 1), but the DVIX alone does not subsume all information relevant for predicting future variance.

4.4.3 Out-of-sample Volatility Forecasts

After assessing the in-sample forecast performance of several models, we move to the out-of-sample forecasting assessment. The predictive performance of the DVIX is evaluated against four alternative predictors of volatility, the GARCH, EGARCH, GJR-GARCH (later referred to as GARCH-type models) and historical volatility. To obtain GARCH-type volatility forecast over the T-period horizon we simply find the average of T individual forecasts at horizons $1, 2, ..., T$, annualize it and then take a square root (see Kroner et al., 1995; Jorion, 1995):

$$GARCHV_{t,t+T} = \sqrt{\frac{252}{T}\sum_{j=1}^{T}\widehat{h_j}}. \tag{4.8}$$

Table 4.5: In-sample Estimation of GARCH and DVIX Specifications

	GARCH(1, 1)		GARCH(1, 1) - DVIX		DVIX	
	coeff	t-stat	coeff	t-stat	coeff	t-stat
ω ($\times 10^5$)	0.206	(0.747)	-1.963***	(-3.897)	7.116	(1.333)
α_1	0.014	(1.545)	-0.013	(-1.506)		
α_2	0.980***	(81.243)	0.591***	(7.249)		
β_1			0.346***	(5.245)	0.617***	(4.286)
LL	1885.76		1920.97		1891.8	
$LR-stat$	70.42***				58.34***	

Note: This table reports the estimation results for the in-sample analysis for the period from 5 January 2015 to 8 January 2018. t-statistics are robust following White (1982) and presented in parentheses. *** is used to indicate significance at the 1% level. In addition we report the log-likelihood of each model (LL) and the likelihood ratio (LR) test statistic. LR-stat is calculated as twice the difference between the log-likelihood of the long (GARCH(1,1)-DVIX) and the short models (GARCH(1,1) and DVIX). Three specifications look as follows:

$$\begin{aligned} r_t &= \mu + \epsilon_t, \\ \epsilon_t|\mathcal{F}_{t-1} &\sim N(0, h_t), \\ h_t &= \omega + \alpha_1\epsilon_{t-1}^2 + \alpha_2 h_{t-1}, \\ h_t &= \omega + \alpha_1\epsilon_{t-1}^2 + \alpha_2 h_{t-1} + \beta_1 DVIX_{t-1}^2, \\ h_t &= \omega + \beta_1 DVIX_{t-1}^2. \end{aligned}$$

Another alternative of a time-series model for the volatility is a simple historical average, estimated for example over a 30-day window[15], for each day t (see Szakmary et al., 2003):

$$HISTV_t = \sqrt{\frac{252}{30}\sum_{j=1}^{30} R_{t-j+1}^2}. \tag{4.9}$$

Volatility is an unobservable variable and should be estimated. To proxy for actual ex-post volatility at each date t during the next T trading days, we use a set of daily squared returns according to the following equation:

$$RV_{t,t+T} = \sqrt{\frac{252}{T}\sum_{j=1}^{T} R_{t+j}^2}. \tag{4.10}$$

The realized volatility during the interval T is expressed in annual terms.

[15]Changing the estimation window for 20 or 60 trading days does not significantly affect the results.

GARCH-type specifications are estimated for daily futures returns from 1 December 2011 to 5 January 2015 and the forecasts for the conditional variance are made for the next $T = 22$ days, 10 days, 5 days and 1 day, using estimated parameters. The data are then rolled forward one day, keeping the estimation window of the same size by discarding the last observation from the previous estimation.

Following existing literature (Szakmary et al., 2003; Egelkraut and Garcia, 2006; Brittain et al., 2011) we address three different questions related to the bias and information content of the different volatility forecasts. First, we want to investigate whether any of the forecasts defined previously is an unbiased forecast of future realized volatility. To answer this question we estimate a univariate regression of the form:

$$RV_{t,t+T} = \alpha + \beta X_t + \epsilon_{t,t+T}, \tag{4.11}$$

where X_t is either a time-series forecast of volatility, such as $HISTV_t$, or a more sophisticated GARCH-type volatility, or the $DVIX_t$ volatility forecast. If X_t is an unbiased estimate of future realized volatility, then in Equation (4.11) the intercept should be zero and the slope coefficients should be one. We use a Wald test to test the joint hypothesis $H_0 : \alpha = 0 \, \text{and} \ \ \beta = 1$.

Panel A of Table 4.6 reports the results for unbiasedness tests. For each five models, reported in the first five columns, we present the regression coefficients along with t-statistics, where the standard errors of the estimates are adjusted for serial correlation and heteroscedasticity using the Newey-West (1987) procedure. We discuss unbiasedness of the forecasters only for the forecast horizon of $T = 22$, as the DVIX is the 22 trading days implied volatility of the WMP futures expressed as an annual number. We find that all slope coefficients are smaller than one, ranging from 0.423 for HISTV to 0.725 for DVIX. The intercept is significantly different from zero for HISTV and EGARCHV. A χ^2 test statistics with p-values in brackets strongly rejects the joint hypothesis $\alpha = 0 \, \text{and} \, \beta = 1$ for all the models except for the estimate which is produced by the standard GARCH model. The results suggest that HISTV, EGARCHV, GJR-GARCHV and DVIX are biased estimates of the future realized volatility.

Table 4.6: Out-of-Sample Estimation of Forecasting Regressions

Parameter	Model								
	HISTV	GARCHV	EGARCHV	GJR-GARCH	DVIX	HISTV+DVIX	GARCHV+DVIX	EGARCHV+DVIX	GJR-GARCH + DVIX
Panel A: Forecast for T = 22									
const	0.169***	0.060	0.133***	0.084	0.031	0.013	-0.023	-0.004	0.002
	(3.897)	(0.708)	(2.081)	(0.909)	(0.555)	(0.225)	(-0.293)	(-0.067)	(0.032)
TSV	0.423***	0.714***	0.451***	0.621***		0.262*	0.281	0.183	0.169
	(2.998)	(2.859)	(2.811)	(2.317)		(1.897)	(1.061)	(1.077)	(0.571)
DVIX					0.725***	0.559***	0.618***	0.640***	0.645***
					(4.315)	(3.264)	(3.215)	(3.411)	(3.024)
$R^2(adj.)$	0.157	0.110	0.090	0.104	0.210	0.258	0.222	0.221	0.214
χ^2	10.837***	3.728	22.121***	5.859***	12.054***				
	(0.004)	(0.155)	(0.000)	(0.053)	(0.002)				
Panel B: Forecast for T = 10									
const	0.140***	0.039	0.119*	0.057	0.014	-0.001	-0.035	-0.003	-0.015
	(3.694)	(0.540)	(1.850)	(0.774)	(0.308)	(-0.021)	(-0.562)	(-0.045)	(-0.241)
TSV	0.373***	0.647***	0.376***	0.571***		0.224*	0.269	0.084	0.177
	(2.912)	(2.972)	(2.144)	(2.627)		(1.802)	(1.147)	(0.487)	(0.710)
DVIX					0.651***	0.506***	0.542***	0.613***	0.564***
					(4.716)	(3.676)	(3.427)	(4.033)	(3.191)
$R^2(adj.)$	0.095	0.076	0.041	0.072	0.134	0.161	0.142	0.134	0.137
Panel C: Forecast for T = 5									
const	0.128***	0.011	0.083	0.039	0.007	-0.006	-0.047	-0.015	-0.022
	(3.187)	(0.161)	(1.333)	(0.632)	(0.157)	(-0.126)	(-0.732)	(-0.250)	(-0.385)
TSV	0.386***	0.597***	0.356**	0.494***		0.204*	0.308	0.118	0.183
	(2.708)	(2.772)	(1.975)	(2.638)		(1.678)	(1.279)	(0.668)	(0.780)
DVIX					0.546***	0.414***	0.419***	0.495***	0.456***
					(3.944)	(2.848)	(2.599)	(3.447)	(2.428)
$R^2(adj.)$	0.059	0.052	0.027	0.043	0.073	0.090	0.082	0.074	0.076
Panel D: Forecast for T = 1									
const	0.074**	-0.003	0.059	0.046	-0.011	-0.022	-0.053	-0.029	-0.018
	(2.374)	(-0.049)	(1.058)	(0.862)	(-0.287)	(-0.521)	(-0.886)	(-0.537)	(-0.349)
TSV	0.263**	0.475**	0.271	0.316*		0.166	0.242	0.091	0.041
	(2.236)	(2.294)	(1.593)	(1.906)		(1.353)	(0.980)	(0.505)	(0.218)
DVIX					0.447***	0.339**	0.344**	0.408***	0.426***
					(3.726)	(2.520)	(2.177)	(2.788)	(2.722)
$R^2(adj.)$	0.015	0.014	0.006	0.007	0.021	0.025	0.023	0.020	0.020

Note: This table presents the out-of-sample estimation results of Equations (4.11) and (4.12) for the period from 5 January 2015 to 8 January 2018. GARCH-type models are estimated using daily returns from December 2011 to January 2015 and then rolling forward by one observation. We use the Newey and West (1987) correction to adjust the coefficient standard errors to account for the the heteroscedastic and autocorrelated error structure. ***, **, * is used to indicate significance at the 1%, 5%, 10% levels respectively.

Although unbiasedness is a desired property, a bias of a known form does not affect a predictive power of a forecaster. In our next question, we want to compare the predictive power of the various forecasters for different horizons. For that, we again estimate regressions of the realized volatility against the various forecasters, described by Equation (4.11), and compare the R^2. For example, if the DVIX forecasts the future realized volatility better than time-series forecasts, then the R^2 from the corresponding regression needs to be the highest. First five columns of Table 4.6 present estimation results for all the models and for the different forecast horizons. Results imply that all five different volatility measures contain important information about future volatility, as all slope coefficients are positive and significant. Comparing the explanatory power for the different forecast horizons reveals an interesting pattern. The degree of predictability offered by different volatility measures decreases as we decrease the forecast horizon from 22 to 1 trading day (from the highest of 21% to 0.6%). We attribute it to the fact that the realized volatility expressed by Equation (4.10) for short horizons might be a noisy proxy for the true volatility, as the prices of futures contracts which we use to calculate daily returns often remain unchanged for the two consecutive days, thus yielding zero returns. We find that the DVIX outperforms all alternative forecasters for all the forecast horizons, supported by the highest adjusted R^2 value. The HISTV shows the second best result in forecasting the subsequent realized volatility. At the 22-day horizon, the adjusted R^2 coefficient is equal to 21.0% for the model which includes the DVIX as a forecast for future realized volatility in comparison to 15.7% for the model with the HISTV. At the 10-day horizon, the difference in the adjusted R^2 is 13.4% vs. 9.5%. At the 5-day horizon, the R^2 is 7.3% vs. 5.9% and at the 1-day horizon 2.1% vs. 1.5%. The GARCH-type forecasters do not produce better forecast than HISTV or DVIX at any of the horizons.

In the third question we want to assess whether the DVIX efficiently impounds all information about future realized volatility, including what is represented by time-series forecasts. To answer this question, we need to consider the results from the encompassing regressions involving the DVIX and either the HISTV or GARCH-type forecasts:

$$RV_{t,t+T} = \alpha + \beta_1 TSV_t + \beta_2 DVIX_t + \epsilon_{t,t+T}. \tag{4.12}$$

If the DVIX is an informationally efficient predictor of the subsequent realized volatility and the time-series forecasts contain no information beyond what is al-ready included in the DVIX, then we should expect β_1 in Equation (4.12) to be insignificant. If both β_1 and β_2 are significant, then the time-series forecast complements the DVIX. The estimation results for four different forecast horizons and four models are reported in columns six through nine of Table 4.6. At the 22-day horizon we find that the DVIX efficiently impounds all information when the alternative volatility forecast is modelled with GARCHV, EGARCHV or GJR-GARCHV estimate, indicated by the insignificant β_1 coefficients. However, when the alternative volatility forecast is modelled with HISTV, we find that the approaches complement each other. When we look at the shorter horizons, we again find that the DVIX efficiently impounds all information when the alternative volatility forecast is modelled with the GARCH-type models. However, HISTV offers some additional information which is not captured by the DVIX alone. We also notice that the model which includes both HISTV forecast and the DVIX produces the highest adjusted R^2, representing the best combination of the time-series and implied volatility estimates for predicting subsequent realized volatility. For the 22-day horizon, movements in the DVIX alone can explain 21.0% of the variability in subsequent realized volatility, while the combination of the DVIX and HISTV can explain 25.8%, suggesting that the HISTV forecast offers some additional information which is not captured by the DVIX.

Previous results show that the combination of two volatility forecasts (a time-series forecast with the DVIX) has a better predictive power, measured by the adjusted R^2, than a time-series forecast alone. Following Rapach et al. (2010) we compare the mean squared prediction error (MSPE) for the predictive regression with one forecaster against a combination of two volatility forecasters. We construct the out-of-sample R^2 statistic, R^2_{OS}, which measures the reduction in MSPE for a long model relative to a parsimonious model. R^2_{OS} is defined as follows:

$$R^2_{OS} = 1 - \frac{\sum_{i=1}^{q}(\sigma_i - \hat{\sigma}_i^{large})^2}{\sum_{i=1}^{q}(\sigma_i - \hat{\sigma}_i^{parsimonious})^2}, \tag{4.13}$$

where σ is realized volatility and $\hat{\sigma}^{parsimonious}$ ($\hat{\sigma}^{large}$) is a forecast of σ constructed by using a time-series forecast (a time-series forecast and the DVIX), q is a number of data points in the out-of-sample forecast. As in Clark and West (2007) we generate out-of-sample forecasts of the realized volatility using an expanding estimation window. More specifically, we divide the total sample from 5 January 2015 to 8 January 2018 into two equal sub-samples, an in-sample portion of the first m observations and an out-of-sample portion of the last q observations. We construct a forecast using only the data available up to the time at which the forecast is made. To form initial out-of-sample forecasts, we use regression coefficients from Equations (4.11) and (4.12), estimated on the evaluation period of m observations, the next out-of-sample forecast uses regression estimates based on $m + 1$ observations. The last out-of-sample forecast is based on $m + q - 1$ observations. When $R^2_{OS} > 0$, a long model forecast outperforms a parsimonious forecast. We conduct a Clark and West (2007) test to find out whether a long model has a significantly lower MSPE than the parsimonious model. The statistic of Clark and West (2007) is an adjusted version of the Diebold and Mariano (1995) and West (1996) statistic, and can be used for comparing forecasts from nested models. The null hypothesis of the test is that there is no difference in the accuracy of two forecasts (equal MSPE). Under the alternative, MSPE from a larger model is less than that of a parsimonious model. Table 4.7 shows the results for four different forecasting horizons. Almost each entry of Table 4.7 is positive and statistically significant, which means that combining the DVIX with a time-series forecaster improves an accuracy of realized volatility forecasting. We also notice that the statistical significance of the R^2_{OS} statistic falls as we shorten the forecasting horizon. This finding is intuitive, as the DVIX is a 22-trading days estimate of future volatility, and its forecasting performance deteriorates on lower forecasting horizons.

4.5 Conclusions

In this chapter, we evaluate the ability of option implied volatility to forecast future realized volatility in New Zealand's largest goods export sector - the dairy sector. To conduct both in- and out-of-sample analyses, we use data for the most actively traded NZX Dairy Derivative - WMP futures and options contracts. We

Table 4.7: R^2_{OS} statistics

Parsimonious vs Long Model	Forecast Horizon			
	T = 22	T = 10	T = 5	T = 1
HISTV vs. HISTV+DVIX	0.266***	0.180***	0.099***	0.010**
	(3.282)	(3.806)	(3.826)	(2.293)
GARCHV vs. GARCHV+DVIX	0.331***	0.242***	0.143***	0.028***
	(4.410)	(4.997)	(4.969)	(3.462)
EGARCHV vs. EGARCHV+DVIX	0.352***	0.272***	0.166***	0.040***
	(5.002)	(5.636)	(5.524)	(4.218)
GJR-GARCH vs. GJR-GARCH + DVIX	0.284***	0.237***	0.120***	-0.147
	(3.787)	(3.386)	(2.593)	(-2.216)

Note: This table reports R^2_{OS} statistic. Statistical significance for the R^2_{OS} is based on the p-value for the Clark and West (2007) statistic. Clark and West (2007) test statistic is presented in parentheses. The statistic corresponds to a one-sided test of the null hypothesis that there is no difference in the accuracy of two forecasts (equal MSPEs). Under the alternative, MSPE from a long model is less than that of a parsimonious model. We use the Newey and West (1987) standard error estimate to control for autocorrelation. ***, **, * is used to indicate significance at the 1%, 5% and 10% levels.

compare two approaches in a volatility forecasting exercise - time-series models and the market-based forecast recovered from the option market. The time-series predictors include the historical volatility and GARCH-type forecasts. To construct the dairy implied volatility index, we closely follow the CBOE VIX methodology. Before investigating the forecasting performance of the DVIX, we assess its time series properties, seasonalities and the relation between the DVIX and WMP futures returns.

The analysis of the intraweek pattern does not reveal any day-of-the-week effect in the DVIX. Similarly, investigation of seasonalities at the monthly level does not reveal any month-of-the-year effect. When analysing the degree of comovement between the changes in the WMP futures prices and the changes in the DVIX, we find that an increase (decrease) in the implied volatility at time t is associated with a decline (increase) in the WMP futures returns at time $t-1$, suggesting an inverse one day lagged return-DVIX relationship. Further investigation shows that this relationship is asymmetric. Positive moves in the WMP futures prices are associated with larger absolute changes in the DVIX than negative moves in the WMP futures prices.

Next, we compare the in-sample performance of three different models designed for predicting conditional future variance: a standard GARCH model which only uses historical information, a GARCH model which uses both historical information and implied volatility, and a model which uses the information of the DVIX only. The results strongly suggest that the DVIX has a high information content regarding conditional variance and that the inclusion of historical information further improves the model's fit.

Finally, we perform the out-of-sample forecast of the future realized volatility using the DVIX and the other alternative volatility forecasts. The forecast horizons range from 22 trading days to 1 day. We find that the DVIX provides substantial information about future realized volatility and beats alternative forecasters, while being a biased estimate of future realized volatility. We also document that the combination of historical information and the DVIX provides the best forecast accuracy for all forecast horizons, meaning that the historical volatility forecast and the DVIX complement each other.

To our knowledge, we are the first to construct and examine the predictive power of implied volatility for the NZ dairy sector. The results of our chapter are particularly important for decision makers in the financial and agricultural sectors who require the estimate of volatility for pricing and risk management purposes. By constructing the implied volatility index we have created a measure of volatility in the dairy sector. The DVIX quantifies volatility and, by comparing its current level with some historical values, one can gauge the behaviour of the NZ dairy market.

4.A Appendix: Figures and Tables

Figure A1: Volatility in Commodity and Other Asset Prices (in a percentage)

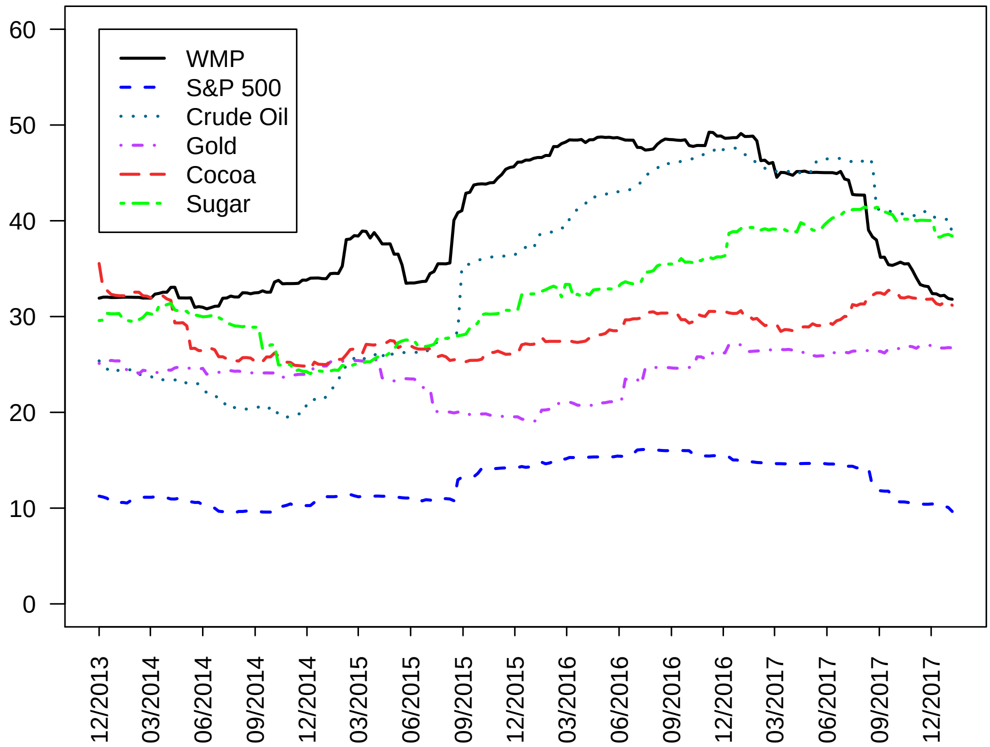

Note: This figure shows annualized volatility of WMP Futures, S&P 500, CRB Crude Oil, Gold, Cocoa and Sugar Indices for the period between December 5, 2011 and January 8, 2018. Volatility is measured as the standard deviation of weekly returns over preceding 2 years.

Table A1: NZX Dairy Derivatives Products

Contract/ Commodity	Contract Size	Terminal Price/ Settlement Method	First Traded	Contract Month
Whole Milk Powder (WMP) Futures*	1 metric tonne	Cash settled against Global-DairyTrade WMP prices	October 8, 2010 *November 30, 2011	18 months are available for trading *12 months are available for trading
Skim Milk Powder (SMP) Futures*	1 metric tonne	Cash settled against Global-DairyTrade SMP prices	February 18, 2011 *December 4, 2017	18 months are available for trading *12 months are available for trading
Anhydrous Milk Fat (AMF) Futures	1 metric tonne	Cash settled against Global-DairyTrade AMF prices	February 18, 2011	18 months are available for trading
Butter (BTR) Futures	1 metric tonne	Cash settled against Global-DairyTrade Butter prices	December 12, 2014	18 months are available for trading
Milk Price (MKP) Futures*	6,000 kilograms of milk solids	Cash settled against Fonterra's Farmgate Milk Price	May 26, 2016 *June 28, 2016	Every September such that up to 5 calendar years are available for trading

Note: This table summarises all currently available dairy derivatives at the NZX. *denotes options are also available, and their specifications.

4.B Appendix: Approximation of American Option Values by Barone-Adesi and Whaley and Its Inversion

Let F denote the current futures price, T is the time to expiration of the futures contract, σ is volatility and X is the strike price of an American option on futures.

Let c denote the value of a European call option (Black and Scholes, 1973). According to Barone-Adesi and Whaley (1987), the value C of an American futures call option is approximated by Equations (4.14) - (4.20)

$$C = \begin{cases} c(F,T,X) + A_2\left[\frac{F}{F^*}\right]^{q_2}, & F < F^* \\ F - X, & F \geqslant F^* \end{cases} \tag{4.14}$$

$$c(F,T,X) = Fe^{-rT}N(d_1) - Xe^{-rT}N(d_2) \tag{4.15}$$

$$d_1 = \frac{ln(\frac{F}{X}) + 0.5\sigma^2 T}{\sigma\sqrt{T}} \tag{4.16}$$

$$d_2 = d_1 - \sigma\sqrt{T} \tag{4.17}$$

$$A_2 = [\frac{F^*}{q_2}]\{1 - e^{-rT}N[d_1(F^*)]\} \tag{4.18}$$

$$q_2 = \frac{1}{2}[1 + \sqrt{1 + \frac{8r}{\sigma^2(1 - e^{-rT})}}]. \tag{4.19}$$

The critical value F^* is defined as a solution of

$$F^* - X = c(F^*) + A_2(F^*). \tag{4.20}$$

The approximation formulas for an American futures put option are similar and are described by the following set of Equations 4.21 - 4.27:

$$P = \begin{cases} p(F,T,X) + A_1\left[\frac{F}{F^{**}}\right]^{q_1}, & F > F^* \\ X - F, & F \leq F^* \end{cases} \tag{4.21}$$

$$p(F,T,X) = Xe^{-rT}N(-d_2) - Fe^{-rT}N(-d_1) \tag{4.22}$$

$$d_1 = \frac{ln(\frac{F}{X}) + 0.5\sigma^2 T}{\sigma\sqrt{T}} \tag{4.23}$$

$$d_2 = d_1 - \sigma\sqrt{T} \tag{4.24}$$

$$A_1 = -[\frac{F^{**}}{q_1}]\{1 - e^{-rT}N[-d_1(F^{**})]\} \tag{4.25}$$

$$q_1 = \frac{1}{2}[1 - \sqrt{1 + \frac{8r}{\sigma^2(1 - e^{-rT})}}]. \tag{4.26}$$

The critical value F^{**} is determined by solving

$$X - F^{**} = p(F^{**}) + A_1(F^{**}). \tag{4.27}$$

In our paper we are interested in extracting implied volatility and for that we use numerical methods, as there is no closed-form inverse solution of Barone-Adesi and Whaley (1987) formula. Each iterative step initiated to find implied volatility involves another iterative procedure, which solves Equation (4.20) for the critical price F^*. Thus we have two nested iterative procedures.

4.C Appendix: Calculation of the Dairy VIX

The Chicago Board of Options Exchange (CBOE) introduced the CBOE Volatility Index (VIX) in 1993, which later was renamed the VXO. The VXO is constructed from the Black and Scholes (1973) option implied volatilities of the eight near-the-money, nearby, and second nearby options on the S&P 100 index. The VXO is based on trading days, meaning that instead of directly using the calendar implied volatility, the implied volatility is transformed to a trading-day basis in the following manner:

$$\sigma_t = \sigma_c\Big(\frac{\sqrt{N_c}}{\sqrt{N_t}}\Big), \tag{4.28}$$

where σ_t (σ_c) is the trading-day (calendar-day) implied volatility rate and N_t (N_c) is the number of trading (calendar) days to option expiration, computed as:

$$N_t = N_c - 2 \times \text{int}(N_c/7). \tag{4.29}$$

In constructing the DVIX we closely follow the methodology described in Whaley (1993). Because options on WMP futures are American style, to extract implied volatilities we use the option pricing approximation of Barone-Adesi and Whaley (1987), which is detailed in Appendix 4.B. Next, we apply the trading-day adjustment using Equations 4.28 and 4.29. The construction of the Dairy VIX is based on the eight near-the-money, nearby, and second nearby options on WMP futures contracts. We denote the option strike price just below the current underlying futures price, F_i, as X_d^i, and the strike price just above the current settlement futures price as X_u^i, where i corresponds to maturity $i \in 1, 2$. We can arrange the implied volatilities of the nearby and second nearby options in the following array:

	Nearby Conrtact			Second Nearby Conrtact	
	Call	Put		Call	Put
$X_d^1 \, (< F_1)$	$\sigma_c^{X_d^1}$	$\sigma_p^{X_d^1}$	$X_d^2 \, (< F_2)$	$\sigma_c^{X_d^2}$	$\sigma_p^{X_d^2}$
$X_u^1 \, (\geq F_1)$	$\sigma_c^{X_u^1}$	$\sigma_p^{X_u^1}$	$X_u^2 \, (\geq F_2)$	$\sigma_c^{X_u^2}$	$\sigma_p^{X_u^2}$

The next step is to average the put and call implied volatilities in each of the four categories of options (at each of the four strike prices), that is:

$$
\begin{aligned}
\sigma^{X_d^1} &= (\sigma_c^{X_d^1} + \sigma_p^{X_d^1})/2 \\
\sigma^{X_u^1} &= (\sigma_c^{X_u^1} + \sigma_p^{X_u^1})/2 \\
\sigma^{X_d^2} &= (\sigma_c^{X_d^2} + \sigma_p^{X_d^2})/2 \\
\sigma^{X_u^2} &= (\sigma_c^{X_u^2} + \sigma_p^{X_u^2})/2.
\end{aligned}
$$

Next, interpolate between the two near-the-money average implied volatilities to obtain at-the-money implied volatilities. More specifically:

$$
\begin{aligned}
\sigma_1 &= \sigma^{X_d^1}\Big(\frac{X_u^1 - F_1}{X_u^1 - X_d^1}\Big) + \sigma^{X_u^1}\Big(\frac{F_1 - X_d^1}{X_u^1 - X_d^1}\Big) \\
\sigma_2 &= \sigma^{X_d^2}\Big(\frac{X_u^2 - F_2}{X_u^2 - X_d^2}\Big) + \sigma^{X_u^2}\Big(\frac{F_2 - X_d^1}{X_u^2 - X_d^2}\Big).
\end{aligned}
$$

Lastly, interpolate between the nearby and second nearby implied volatilities to create a 22 trading days implied volatility as follows:

$$
DVIX = \sigma_1\Big(\frac{N_{t_2} - 22}{N_{t_2} - N_{t_1}}\Big) + \sigma_2\Big(\frac{22 - N_{t_1}}{N_{t_2} - N_{t_1}}\Big),
$$

where N_{t_1} and N_{t_2} are the number of trading days to expiration of the nearby and second nearby contract, respectively.

Chapter 5

Profit Margin Hedging in the New Zealand Dairy Farming Industry

5.1 Introduction

New Zealand (NZ) is the eighth largest milk-producing country in the world, exporting about 95% of its dairy production.[1] The dairy sector is the largest goods export sector of NZ, with an average annual export revenue of NZD 13.2 billion over the past five years to 2017 (Ballingall and Pambudi, 2017). In 2017, Whole Milk Powder (WMP) accounted for 36% of total dairy export revenue, the highest proportion amongst all dairy products.[2] The dairy farming sector is the second most profitable farming sector in NZ (Ballingall and Pambudi, 2017), however, recent milk payouts received by dairy farms have shown considerable variations. The dairy sector in NZ is free from government interventions, i.e. the government does not provide any price support mechanisms or subsidies and, hence, farms are exposed to shocks in global milk prices.[3] For most farms, the milk price per season is set by Fonterra, a farmer-owned cooperative, which controls about 80% of the NZ milk supply. The price of milk per season depends on five reference commodities which are WMP, Skim Milk Powder (SMP), and their by-products (butter, Anhydrous Milk Fat (AMF) and Buttermilk Powder (BMP)). Amongst those five commodities,

[1] `https://www.dcanz.com/about-the-nz-dairy-industry/`

[2] `https://www.mpi.govt.nz/news-and-resources/open-data-and-forecasting/situation-and-outlook-for-primary-industries-data/`

[3] The only programme in existence is called the Income Equalization Scheme which was designed to smooth out taxable income and hence reduce the tax obligations.

WMP plays the most important role, as historically its contribution to the price of milk is about 62%.[4] In Figure 5.1, we show milk and WMP prices per season in NZD, where the WMP price is weighted by production during a season (Fonterra, 2017). As can be seen from the graph, during the 2014-2015 season, the dairy sector experienced a downturn and the milk price dropped by 48%. The next season it further declined by 11% to 3.9 NZD per kg of milksolids. This extreme volatility of milk prices led to a decline in operating profits of many dairy farms.[5] During the 2015-2016 season, total cash expenses exceeded dairy cash income, resulting in negative profit margins for dairy producers (DairyNZ, 2017). Such an inherently risky operating environment poses at least two problems. The first concerns the sustainability of farming businesses and the second is the inability of farms to service their debts. The Reserve Bank of NZ underlines the second problem in its financial stability report, stating that the dairy sector's indebtedness is one of the top three most important domestic risks to NZ's financial system (RBNZ, 2018). Given the obvious importance of protecting the financial position of dairy farms, reducing milk price risk should be a key focus of financial institutions and farms.

Traditional literature on hedging (e.g. Stein, 1961; Johnson, 1960) sets its objective to minimize the variability of returns. For a producer of an agricultural commodity, this approach generally dictates routinely taking short positions in futures contracts to achieve a minimum-variance hedge ratio. However, Collins (1997) argues that hedging should be used to avoid bankruptcy, rather than to minimize the variability of returns. In the so called "profit margin hedging" strategy, a producer hedges for production only when a futures price is above expected variable and/or fixed costs, or more generally is above a target, which can deliver a predetermined fixed margin. Thus, the objective of profit margin hedging is to assure profitable production by locking in favourable prices in futures markets when they appear. While the protection against downside price risk for a future cash sale is the main concern of profit margin hedging, the debate on whether profit margin hedging can

[4] `https://www.fonterra.com/nz/en/investors/farmgate-milk-prices/milk-price-methodology.html`

[5] According to NZ media, during this downturn mental health workers saw increased suicide rates, domestic violence, alcohol and drug use amongst farmers.

Figure 5.1: Milk and WMP prices per season in NZD $/kg

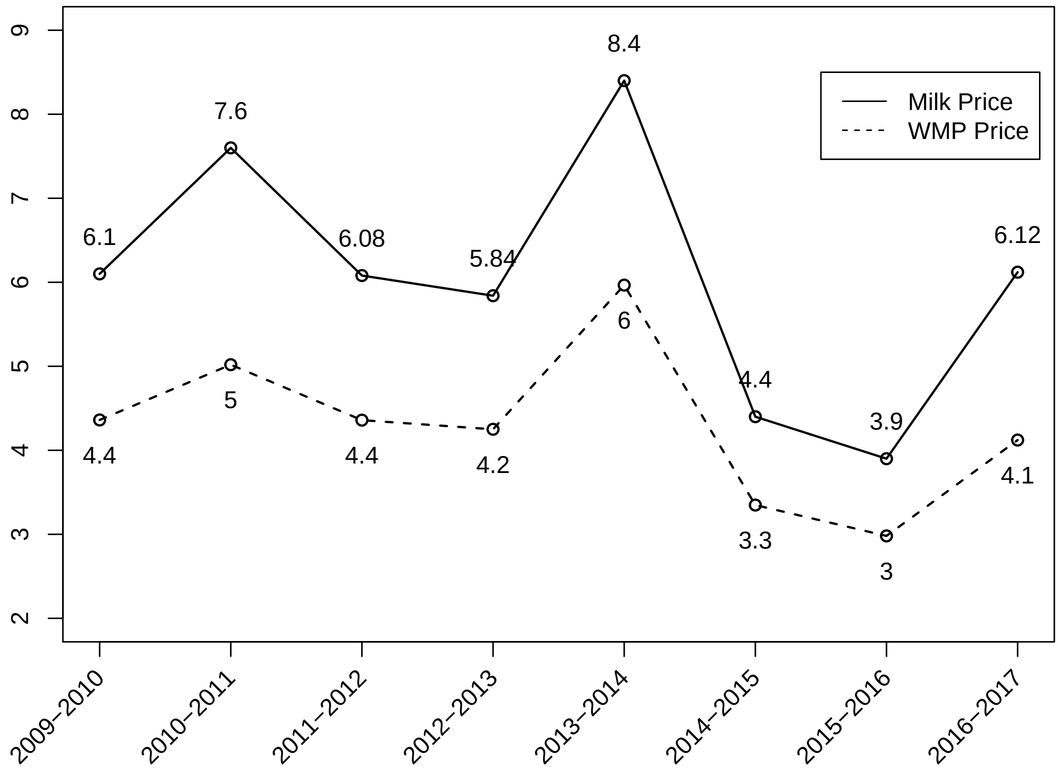

Note: Milk prices are Fonterra's prices per kilogram of milk solids and WMP price is the weighted average USD contract prices converted to NZD and divided by 1000.

be used to enhance margins still goes on. Conceptually, in efficient markets future price changes are unpredictable and, thus, hedging in futures markets should not generate speculative profits. Nevertheless, some empirical evidence suggests that profit margin hedging can generate an increase in farms' average returns. Kenyon and Clay (1987), for instance, find that profit margin hedging for hog producers increases average returns and reduces return variance. Yoon and Brorsen (2005) argue that multiyear rollover hedging[6] can lead to increased expected returns if futures prices follow a mean reverting process. Kim et al. (2010) develop theoretical arguments that justify the profit margin hedging strategy over continuous hedging or selling at harvest. They show that when futures prices are mean reverting, profit margin hedging has a higher expected profit over alternative strategies.[7]

In this chapter, we examine the performance of a profit margin hedging strategy for NZ dairy farms. We implement this strategy for six seasons covering 2011 to

[6]Multiyear rollover hedging is similar to profit margin hedging with the difference that the former considers the possibility to lock in favourable prices for multiple years, instead of a single period.

[7]Their derivation is based on a static one-period model without basis and yield risk.

2017. A target value of a milk payout which avoids a financial failure is called a break-even milk price (BEMP). BEMP defines what level of milk income is required to meet farm working expenses, taxes, interest and rent payments, and drawings. We develop a profit margin hedging strategy where the expected milk production is hedged whenever the price is above the BEMP. To construct this strategy, we use WMP futures contracts, which are the most liquid dairy derivatives traded on the NZ stock exchange (NZX). In the first part of our analysis, we apply the profit margin hedging strategy to a representative farm. Specifically, we compare the strategy's impact on risk and return of the farm's monthly revenue, relative to a no-hedging strategy of cash sales and to continuous hedging. In the second part of the analysis, we implement the strategy to a unique sample of real farm data, obtained from the DairyBase database, which contains NZ dairy farms financial and physical data. We measure the benefits of the profit margin hedging strategy by analysing the change in discretionary cash. Low discretionary cash signals liquidity/solvency problems, and we define an occurrence of a negative discretionary cash position as financial distress.

Both parts of the analysis demonstrate significant benefits of profit margin hedging, and the results do not substantially change after the incorporation of basis risk. When we apply the profit margin hedging strategy to a representative farm, we find that, even after accounting for brokerage fees, the strategy increases the farm's average payout, and reduces its volatility and semivariance. Of the three strategies, profit margin hedging delivers the highest average farm payout, followed by continuous hedging, and not hedging. Paired differences tests confirm that the average price of the profit margin hedging strategy is significantly different from continuous hedging and not hedging. We also find that the profit margin hedging strategy delivers a greater decrease in semivariance relative to the continuous hedging scenario.

When we apply the strategy to a sample of real farm data, after accounting for a high level of brokerage fee, we find that the mean value of discretionary cash increases by 36%, volatility reduces by 35% and downside risk, measured by semivariance, reduces by 74%. We find that the largest improvement in discretionary cash, by

62%, occurs for the highly leveraged farms. Additionally, we show that profit margin hedging reduces the probability of financial distress during any year by more than half, from 35% to 16%. To estimate the economic effect of the profit margin hedging strategy, we scale profits generated by this strategy across the sample of farms to all NZ dairy farms. We estimate that the strategy could have generated an additional NZD 0.49 billion annually, about 3.7% of the yearly dairy export revenue. Our findings suggest that profit margin hedging can increase the sustainability of the farm businesses by decreasing the chances of their financial distress.

Our study has two important implications. First, our findings show that the WMP futures are not redundant and highlight the usefulness of the futures market to dairy producers. Second, the findings suggest that the futures market allows farms to lock in favourable prices and thus futures contracts could be used by farms or by financial and government institutions aiming at providing risk management services to farms. Our findings are an important reminder of the benefits of hedging.

This study is positioned within two streams of literature, cross hedging and selective hedging. Cross hedging means that we hedge exposure to milk price by not trading in milk price futures, but instead in WMP futures. Historically, the WMP price contributes to the price of milk with a weight of around 62%, which makes the correlation between milk price and WMP high, motivating us to explore cross hedging techniques.[8] In addition, it relates to the literature on selective hedging in the sense that an agent enters a futures position only when prices are favourable. We concentrate our study on hedging output prices only, while some studies develop risk management strategies for hedging both input and output prices (Peterson and Leuthold, 1987; Kim et al., 2009).

The remainder of this chapter is organized as follows: Section 5.2 reviews the literature on dairy farm risk management, profit margin hedging, and explains the structure of the NZ dairy market. Section 5.3 starts by detailing how profit margin hedging is applied to NZ dairy farms and then proceeds with the empirical analysis. First, it examines the predictability of the WMP futures to determine whether profit margin hedging could be used to increase expected returns for dairy farms. Second,

[8]The correlation between yearly returns in milk price and WMP price is 0.96.

it presents the empirical results for a representative dairy farm. Third, it expands the strategy to the individual farms. We present concluding comments in Section 5.4.

5.2 Literature Review

5.2.1 Dairy Farm Risk Management

Previous research on dairy farm risk management has primarily focused on the US and mainly deals with managing the risk of output prices. However, some studies concentrate on hedging farms' input costs. Bosch and Johnson (1992) consider the variability in feed prices and crop yields as the main risk for net returns of dairy farms and find that hedging and crop insurance lower expected net returns but reduce risk. Maynard et al. (2005) focus on output price variation and evaluate hedging effectiveness of futures and put options in minimizing downside price risk. They estimate minimum semivariance futures and options hedge ratios and find that, when futures are used, the semivariance of the net price received for milk is reduced by 24 - 59% depending on the region.

A few studies analyse the effects of various risk management strategies for dairy farms through the use of Monte Carlo simulations. Manfredo and Richards (2007) evaluate the effect of various risk management strategies on the financial performance of a representative US dairy cooperative and its members. They document that placing a hedge when futures prices are greater than the variable costs of milk production results in a reduction of semivariance of milk revenue by 27%, but increases the standard deviation by 8%. Neyhard et al. (2013) incorporate an individual's debt position and analyse the performance of futures and options contracts as hedges for a dairy farm with three different levels of debt. They simulate both milk and feed prices, and implement different risk management techniques aimed at meeting all expense and debt obligations of the farm. They find that, in the case where both milk and feed are hedged with futures, the net farm income standard deviation decreases by 5.6%, 6.8% and 7.8% for low, average and high debt levels, respectively.

Another strand of literature aims to identify the factors that explain the adoption of risk management tools by farms. In an expected utility framework, Turvey and Baker (1989) show that the capital structure of a farm is an important variable determining the amount of commodity hedged. Empirical work by Wolf and Widmar (2014) supports this finding. Wolf and Widmar (2014) collect survey data of dairy farms in the US and estimate a multinomial logit model on forward pricing adoption of milk or feed. They find that managers of larger herds, with more education and higher solvency risk, are more likely to use both feed and milk forward pricing methods. Among the reasons that dairy farm managers provide for not using any forward pricing tools are lack of knowledge, reliance on cooperative to adopt forward pricing, costs, basis risk, and lack of time to manage finances.

In sum, the research on dairy farm risk management suggests that a reduction of volatility or semivariance of milk revenue can be achieved through the use of derivatives. However, in reality a lack of understanding of hedging techniques, among other reasons, prevents farms from an active use of forward pricing tools.

In contrast to the US, milk price risk management in NZ is a relatively new topic. In 2010, the NZX introduced WMP futures, followed by SMP and AMF futures, WMP options and butter futures in December 2014. In 2016, the NZX developed milk price futures and options contracts. Although NZ has an existing dairy derivative market, only one academic study has tested the effectiveness of dairy derivatives as hedgers (Koeman and Białkowski, 2015). They estimate the static minimum-variance hedge ratio through a regression model and control for the possibility of cointegration between spot and futures WMP prices. They use data from October 2010 to March 2013 and find that the optimal hedge ratio is equal to 0.6 and that the variability of a hedged portfolio can be reduced by 70%. Their paper is different from the current study in several aspects. Koeman and Białkowski (2015) find the optimal hedge ratio as the slope coefficient in the regression of changes in spot prices on changes in futures prices. This classical approach evaluates the ability of WMP futures to minimize the variance of the hedged portfolio, where the effectiveness is measured by the R-squared. Their analysis is beneficial to producers of WMP. In contrast, we design a strategy specifically for dairy farms and measure a direct impact from hedging in WMP futures on farm revenue. We acknowledge that sustainability of farming business depends on the ability to receive milk payouts which are above break-even prices. We estimate how the WMP futures market can improve the level of cash that is available for drawings, debt repayments, capital development, and purchases. Additionally, we extend the analysis to actual farm-level data, where we incorporate individual farms' cash expenses and production data.

In this chapter, we aim to expand our knowledge about the usefulness of WMP futures for NZ dairy farms when we apply a selective hedging technique, namely profit margin hedging. This strategy allows dairy farms to protect themselves from financial distress and opens a possibility to increased profits. Kim et al. (2010) have shown, theoretically, that, if futures prices follow a mean-reverting process, then a hedging rule of selling futures above the long-run mean leads to an increase in profits relative to continuous hedging and not hedging. We test this theory by examining the time-series properties of WMP futures to see whether we can expect increased profits in Subsection 5.3.2, but first we explain profit margin hedging in more detail.

5.2.2 Profit Margin Hedging

We start this subsection with an example that explains how profit margin hedging works. Suppose at time t a producer decides to hedge WMP to be sold at $t+1$ because there is a risk that the price could fall below the break-even WMP price. Assume that the break-even price is USD 3,000, there is no basis risk and no transaction costs. We plot a hypothetical price of $t+1$ futures in Figure 5.2. Because the producer wants to sell WMP above USD 3,000, every day he compares the $t+1$ futures price with the break-even price. On a given day the price of the futures contract is USD 3,030, the first time above the break-even price, and the producer sells one futures contract at USD 3,030. To offset the position the producer can buy the contract back, or wait till the expiration; we assume that he waits till the expiration. Since the contract is cash settled, she does not need to worry about delivery. At $t+1$, the price of the futures is USD 2,230, and thus the producer makes a profit of USD 800 on the futures contract. The WMP spot price is also equal to USD 2,230 (i.e. zero basis), and the net sell price is made up of the spot price and the gain on futures, totalling to USD 3,030. By adopting the profit margin hedging strategy, the producer has created price certainty ahead of time, and knows that she can meet her financial obligations. Next, we review specific approaches which have been examined in the literature dedicated to profit margin hedging strategies for different agricultural commodities.

Figure 5.2: Profit Margin Hedging Example

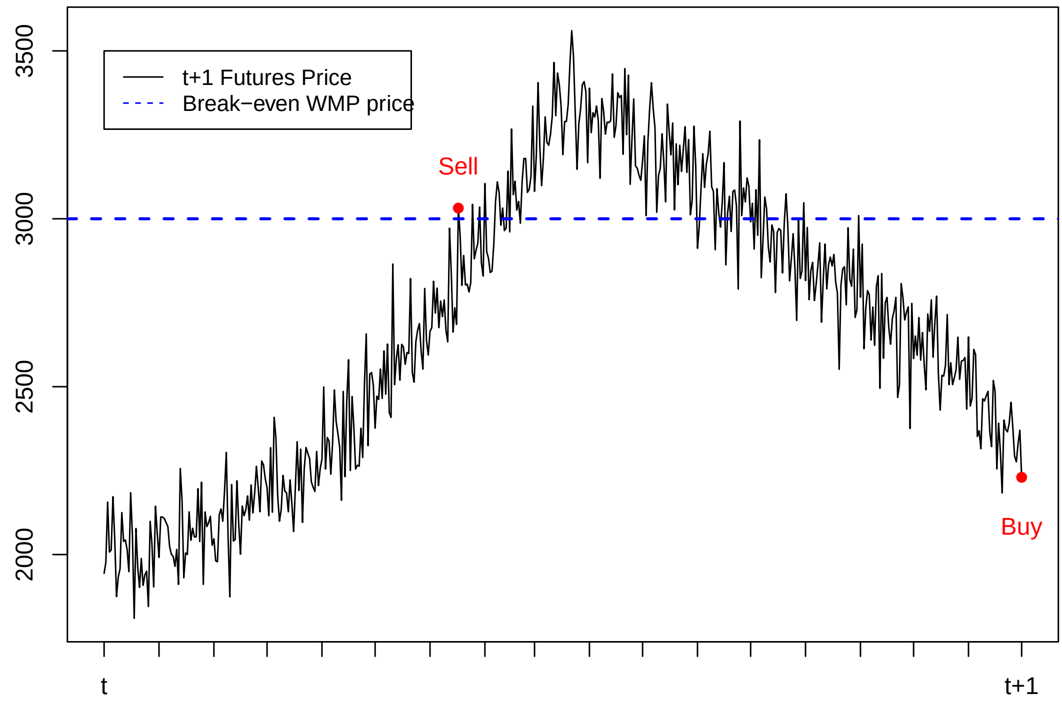

Note: The Figure illustrates how a producer can hedge a production of WMP above the break-even price.

Martin and Hope (1984) define an approach where a proportion of the crop is hedged at a base target set at production costs, and the rest is hedged if the futures price moves up or down by the predetermined level relative to a base target. Wood et al. (1989) explore profitable margin opportunities for cotton producers with an objective of locking in futures prices above total production costs. They find that cash sales at harvest generally provide lower profit margins than margins attained with futures contracts. Kenyon and Clay (1987) adapt a slightly different trigger of futures market activity for hog producers where a hedge is placed if the current futures price is above production costs plus a predetermined fixed margin. They document that, when hedging at low expected profit margins, the strategy does not yield an increased profit, but at several higher levels of expected profit margin, the strategy increases average profit and decreases its variance. Schroeder and Hayenga (1988) choose a similar approach for cattle feedlot producers and again find that hedging with futures can increase average returns and reduce variance compared to cash-market returns. Kim et al. (2009) use a local polynomial forecasting technique to predict cash prices and adopt hedging only when the forecasted cash profit margin is negative. They consider both a one-to-one and risk-minimizing hedge ratio. They show that such selective hedging dominates continuous and unhedged strategies in terms of mean and variance. Kim et al. (2010) implement a profit margin hedging strategy for wheat, corn and soybeans, respectively. They document that only in the case of soybeans, the strategy generates a significant increase in returns in comparison to continuous hedging and selling at harvest.

To summarize, there is no precise rule in profit margin hedging for choosing a target price and a hedge ratio. Some studies choose to cover only a part of production costs, some full production costs, and others extra positive profit margin. If futures prices are above a target, more commonly, producers sell futures contracts in a one-to-one ratio to spot market, but sometimes a risk-minimizing hedge is adopted.

5.2.3 The New Zealand Dairy Market

To further understand the dairy market in NZ, we briefly explain the role of Fonterra, Global Dairy Trade (GDT) events and the NZX dairy derivatives market. The Fonterra co-operative was formed in 2001 from a merger of the country's two largest dairy co-operatives, Kiwi Co-operative Dairies and New Zealand Dairy Group, with the New Zealand Dairy Board. Upon its creation, Fonterra collected approximately 96% of NZ's milk production. As of today, Fonterra still has a dominant position in the dairy product markets and collects about 80% of milk production, 95% of which is exported in the form of dairy ingredients, which makes Fonterra the world's largest dairy exporter. Fonterra is owned by around 10,000 dairy farmers, whose proportion of ownership depends on the volume of milk they supply to Fonterra. Fonterra buys raw milk from its farmer shareholders at a rate per kilogram of milk solids, which is called the Farmgate Milk Price. The Farmgate Milk Price is calculated in accordance with the Farmgate Milk Price Manual. In broad terms, the Farmgate Milk Price is the theoretical price that Fonterra would derive if it converted all milk into the 'Reference Commodity Products', which are WMP, SMP and their by-products (butter, BMP, and AMF). This theoretical price is adjusted for costs, such as those which would be incurred to transport raw milk to Fonterra's NZ factories, produce these same commodities in an efficient way, freight them to the point of export from NZ and make a market return on investment. The prices for the 'Reference Commodity Products' are USD prices achieved by Fonterra on the twice-monthly GDT events platform, converted to NZD at Fonterra's actual average monthly foreign-exchange conversion rate. The GDT events trading platform was formed in 2008 to facilitate the global trading of dairy products. It connects sellers and buyers from over 80 countries, who during an online auction process discover reference prices for globally traded dairy ingredients. GDT offers six different forward contracts. Contract 1 is for shipment in one month, Contract 2 for shipment in two month and so forth.

The NZX launched its first dairy derivative - WMP futures - in 2010, shortly after the introduction of the GDT platform. WMP futures are cash settled to an average of the two winning prices for WMP, Contract 2, determined in GDT events in the same month. The futures trading terminates on the day before the second GDT event of the month. The NZX later added SMP, AMF and BMP futures and options. The most recent contracts are milk price futures and options, which are cash settled against Fonterra's farmgate milk price and were launched in 2016.

5.3 Profit Margin Hedging Strategy

5.3.1 Specifications of the Profit Margin Hedging Strategy for NZ Dairy Farms

In this section, we detail the steps in a profit margin hedging strategy, which are: an objective, decision rule to determine the time when a position in the futures market should be established, and the role of brokerage fees and basis risk. We also describe how we implement each of these steps for the case of a representative NZ dairy farm.

The objective of profit margin hedging is to lock in favourable prices when they occur in futures markets. We define prices as favourable if, after adjusting for basis risk and fees, they are higher than the target price which we set to the NZ break-even milk price (BEMP) reported in the Dairy NZ Economic Surveys (DairyNZ, 2017). This approach in selecting a target price aims to guarantee the economic viability of the farm's business. BEMP indicates how much milk income is required to meet farm working expenses, interest, rent, tax and drawings. Table 5.1 reports BEMPs for the seasons 2011 - 2017 for owner-operated farms. The data shown in Table 5.1 are the averages for groups of farms that closely match the average regional herd size, hectares and milksolids production, as described in the New Zealand Dairy Statistics for a particular year (DairyNZ, 2017). Because for hedging milk prices we use WMP futures, we need to convert BEMP to break-even WMP prices. To do so, we collect annual farmgate milk prices in NZD, weighted-average USD prices of WMP and average USD:NZD conversion rates from Fonterra's Farmgate Milk Price Statements for the seasons 2009-2017 (Fonterra, 2017). We use these data to identify the relation between milk prices and WMP, and then use the estimated parameters to find break-even WMP prices. A linear regression of weighted-average NZD prices of WMP on farmgate milk prices yields the relation:

$$
\underset{}{Milk_t^{NZD}} = \underset{(-1.540)}{-0.733} + \underset{(14.544)}{1.578}WMP_t^{NZD}, \qquad R^2 = 0.97, \tag{5.1}
$$

where t runs through eight seasons from 2009 to 2017 (the data frequency is annual) and t-statistics are presented in parentheses. We use the BEMPs reported in Table 5.1 as an input for the left-hand side of Equation (5.1) and solve it for break-even WMP prices in NZD. The last step is to convert NZD prices back to USD. Results are reported in Table 5.2. This method gives us an estimate of the break-even WMP price given the BEMP. Ideally, we would incorporate prices of the other four Reference Commodities (SMP, Butter, AMF and BMP). The WMP break-even price is calculated for a representative farm and indicates how much WMP income is required to meet farm working expenses, interest, rent, tax and drawings. It accounts for livestock and other dairy cash income received in the season. The WMP break-even price increases as farm working expenses, interest, rent, tax and drawings increase, and livestock and other cash income decrease.

Table 5.1: New Zealand Break-even Milk Price (NZ$ per kg MS)

	2011-12	2012-13	2013-14	2014-15	2015-16	2016-17
Farm working expenses	3.95	4.13	4.33	4.07	3.64	3.73
Interest and rent	1.31	1.39	1.29	1.36	1.36	1.35
Tax	0.32	0.25	0.38	0.21	0.05	0.1
Drawings	0.57	0.65	0.77	0.69	0.49	0.51
Total cash expenses	6.14	6.42	6.77	6.33	5.53	5.7
Livestock & other dairy cash	0.4	0.44	0.42	0.56	0.6	0.53
Break-even milk price	5.74	5.98	6.35	5.77	4.93	5.17

Note: This table reports total cash expenses and break-even milk price for the representative farm between 2011 and 2017.

Table 5.2: Break-even WMP Prices

Season	Break-Even Milk Price NZD/kg MS	Break-Even WMP Price USD/kg MS
2011-2012	5.74	3.160
2012-2013	5.98	3.397
2013-2014	6.35	3.629
2014-2015	5.77	3.248
2015-2016	4.93	2.541
2016-2017	5.17	2.590

Note: This table reports the break-even WMP price given the break-even milk price. WMP prices are obtained from the equation $Milk_t^{NZD} = -0.733 + 1.578WMP_t^{NZD}$ and then converted to USD.

In the simulations, we assume that a dairy farm receives payments each month.

To fix a date, we choose the next day after the first GDT event each month, which usually happens to be the first Wednesday of each month. We assume that a farm enters into a short hedge for a contract which expires the same month as the payment is due. Additionally, we assume that there is a one month lag between milk collection and payment. For example, for milk collected in June, a farm would have received payment at the beginning of July (next day after the first GDT event in July), and it also would have been looking to hedge its June production by entering a short position in July WMP futures contracts.

The next point we want to address is a decision rule that triggers a transaction in futures and its timing. The net price received for milk sales is the value which consists of spot sales of milk, gain/loss in futures contracts minus transaction costs. The final price received by a farm equals the first Wednesday spot price of WMP plus any net gain or loss from the futures trade minus transaction costs. The net price received by a producer using short hedging is defined as:

$$NP = S_1 + F_0 - F_1 - C, \tag{5.2}$$

where S_1 is the WMP spot price, F_0 and F_1 are the opening and the closing futures price, and C is the futures transaction costs. We introduce a transaction cost at three different levels of 30, 50 and 70 USD per contract (round-trip transaction), which is comparable to the indicative fees charged by NZ brokerage firms.[9] The net price also can be expressed as $NP = F_0 + B_1 - C$, where B_1 is the closing basis ($B_1 = S_1 - F_1$). Because the closing basis is uncertain at the time the hedge is placed, a farm can only form expectations about the net price received. After defining the expected net price received, we can formulate the decision rule and its timing: a producer places a short hedge any time after the contract is available for trading, provided that the net price is at least as high as the target price. Because WMP futures for each calendar month are available 18 months into the future, dairy farms can hedge their production 18 months before monthly payments are received, if it is profitable to do so. Once hedges are placed they are not lifted until the cash

[9] `https://www.omf.co.nz/legal/omf-disclosure-statement`

sale of milk.

Lastly, we want to emphasize the role of basis risk. Basis risk is the risk of experiencing the realized closing basis different from the expected one. If a farm follows a decision rule described earlier and enters the position in the futures market, the expected profit margin (EPM), the realized profit margin (RPM) and their difference can be defined as follows:

$$
\begin{aligned}
EPM &= (F_0 + B_1^{exp} - C) - target \\
RPM &= (F_0 + B_1^{real} - C) - target \\
RPM - EMP &= B_1^{real} - B_1^{exp},
\end{aligned}
$$

If RPM is greater or equal to zero, that means that a farm can cover all cash expenses (farm working expenses, interest and rent, tax and drawings) from milk sales and the objective of profit margin hedging is achieved. In the case where EPM is zero (net price just enough to cover cash expenses) and B_1^{real} is less than B_1^{exp}, then RPM is smaller than all cash expenses. Thus, in the case when the realized basis is lower than expected, the farm might receive less than the target price for the WMP sale, making profit margin hedging unable to secure a positive cash flow requirement.

The net price depends on the expected closing basis, and thus we need to make some assumptions about it. Kim et al. (2010) consider two scenarios of incorporating expected basis in hedging decisions. The first is to assume no basis risk, that is, to assume that the actual closing basis is known at the time of the decision. The second is to model an expected closing basis as a five-year moving average. Hatchett et al. (2010) conduct a study in which they try to find the best length of moving average to use. They find that different moving average lengths have similar forecast accuracy. However, if a structural break occurs, a previous year's basis should be used as a forecast. The main difference between our study and theirs is that they analyse commodities which are harvested once a year, while we need to forecast a basis monthly. Because of the limited data available, we choose to forecast the basis as a moving average of past historical monthly values available up to the hedging decision time. Thus, in a perfect foresight basis scenario, a farm opens a trade in WMP futures if the sum of the futures price and the actual closing basis less brokerage fee is greater or equal to a target price. In the second scenario, instead of an actual closing basis, he uses an estimated closing basis.

After having outlined the details of the profit margin hedging implementation, we now introduce various risk measures, which we use to assess hedging performance. We use the standard deviation and semivariance (standard deviation considers both positive and negative deviations from a mean as risk, while semivariance focuses on downside risk). Semivariance is the expected value of the squared negative deviations of possible outcomes from target returns. The semivariance is defined as

$$SV(X,T) = E\{min(X-T,0)^2\}, \tag{5.3}$$

where T is the target price and X is a random variable. In subsection 5.3.3, we will apply Equation (5.3) to realized NPs with the target as the WMP break-even price. In subsection 5.3.4, we use Equation (5.3) to calculate semivariance of discretionary cash below zero. We report semivariance as the square root of the semivariance measure defined by Equation 5.3. The square root is taken in order to express the semivariance in dollar units.

For each risk measure (RM), we follow Conlon and Cotter (2013) and define hedging effectiveness as:

$$HE_{RM} = 1 - \frac{RM(NP_h)}{RM(P_s)}, \tag{5.4}$$

where NP_h is the net price received by a producer if he chooses to hedge and P_s is the price received in the case of no-hedge.

As mentioned in the introduction, the primary objective of hedging is not to make money but to minimize risk. In an efficient market, new information is rapidly incorporated into futures prices. Because new information arrives randomly to the market, price changes should be unpredictable, leaving no possibility for speculative profits. Nevertheless, the question of whether profit margin hedging can increase expected returns is debatable. Kim et al. (2010) posit theoretical arguments that would justify profit margin hedging over continuous hedging or selling at harvest.

They show that when futures prices are mean reverting, profit margin hedging has a higher expected profit over alternative strategies. In the next subsection, we resort to a standard test - variance ratio test - to examine mean reversion in WMP futures.

5.3.2 Mean reversion in WMP Futures

In our study, we use WMP futures contracts, which are based on the Fonterra product, Regular NZ, Contract 2, that is the GDT auction prices of a WMP contract with a delivery in two months. We obtain data from the Agri Data database, which is a part of the NZX Research Centre. The NZX launched WMP futures on October 8, 2010, and, therefore, we consider the sample period from October 8, 2010 to January 3, 2018. To test the random walk hypothesis, we use the second nearby contracts, which are the most active contracts[10] and use weekly observations referring to Wednesday. We consider Wednesdays' observations because GDT events usually happen on Tuesdays at 12:00 UTC, and information about a change in WMP prices is incorporated during the next trading session which is Wednesday in NZ. We define continuously compounded returns as $r_t \equiv log(P_t) - log(P_{t-1})$, and make sure that returns are always taken for a contract expiring in the same month. For example, on the 27th October 2010 the second nearby futures was the November contract, but in a week's time, on 3rd November 2010 the second nearby futures contract is the December contract. To calculate the return between 27th October and 3rd November, we take prices of the November contract only.

Panel A of Table 5.3 presents descriptive statistics of weekly returns on the WMP futures. Returns are negatively skewed and have excess kurtosis. Moreover, the Jarque-Bera test rejects normality at the 1% level. These findings are in line with prior research on futures prices of other agricultural commodities, which find that futures price changes are not well approximated by a log-normal distribution and often leptokurtic (see Hudson et al., 1987; Yang and Brorsen, 1994; Khalifa et al., 2011, among others). We use the Engle ARCH test to establish the presence of conditional heteroskedasticity in returns. The results indicate that we can reject

[10] During our sample period, traded volume in the nearby contract is equal to 69,997 contracts, while the second nearby contract is 110,000 contracts.

the null hypothesis of no conditional heteroskedasticity and conclude that there are significant ARCH effects in the return series.

Table 5.3: WMP Futures: Summary Statistics and Autocorrelation for Weekly Returns

Panel A: Summary Statistics						
Sample Size	Mean	SD	Skewness	Kurtosis	Jarque-Bera	Engle ARCH
372	-0.003	0.047	-0.04	8.37	447.44***	29.02***
Panel B: Autocorrelation						
ρ_1	ρ_2	ρ_3	ρ_4	ρ_5	LBQ_5	LBQ_{10}
0.13**	0.24***	0.01	0.06	-0.11**	32.36***	45.67***

Note: This table presents summary statistics, autocorrelation coefficients and Ljung-Box Q-test for WMP weekly returns. ***, ** indicates significance at the 1% and 5% levels, respectively.

To establish the random walk nature of WMP futures prices we first check for serial correlation. Under the random walk hypothesis, returns of the time series must be uncorrelated at all leads and lags. Panel B of Table 5.3 reports autocorrelation and the Ljung-Box Q-statistic for weekly WMP futures returns. Results show that the first, second and fifth order autocorrelation coefficients of 13%, 24% and -11% are significant at the 5% level. Moreover, the Ljung-Box Q-statistic with five (ten) lags has a value of 32.36 (45.67), which is significant at the 1% level. These findings indicate that we can reject the random walk hypothesis, which means that future price changes can be forecasted using the past price changes.

Next, we follow Lee et al. (2000), Smith and Rogers (2006), Kim et al. (2010), among others and perform the Variance Ratio Test of Lo and MacKinlay (1988). The idea behind the test is that if the natural logarithm of a price series is a random walk, then the variance of q-period returns should equal q times the variance of one-period returns:

$$VR(q) = \frac{\sigma^2(q)}{q\sigma^2(1)}.$$

The sampling distribution of $VR(q)$ under the null hypothesis of uncorrelated return innovations in the presence of general heteroskedasticity is provided in Appendix 5.A. Under the null hypothesis, the Z-statistic is asymptotically standard normal. We perform the variance ratio test for return horizons of 2, 4, 8 and 16 weeks. Table 5.4 shows the variance ratios and test statistics. Results show that the Z-statistic is significantly different from 1 at the 1% level for return horizons of 2, 4 and 8 weeks, meaning the rejection of the random walk hypothesis for the WMP futures prices. The reported Z-statistics are adjusted for heteroskedasticity, which means the rejections of the random walk hypothesis are not due to a changing variance.

Table 5.4: Variance Ratio Test for Futures Prices

Return Horizon (q-weeks)	$VR(q)$	Z-statistic
2	1.14	2.50***
4	1.45	3.74***
8	1.52	2.62***
16	1.14	0.46

Note: Under the random walk null hypothesis the variance ratio $VR(q)$ is one and Z-statistics have a standard normal distribution. *** indicates significance at the 1% level.

The test of Lo and MacKinlay (1988) focuses on the hypothesis that an individual variance ratio for some q is one; however, the null hypothesis requires the variance ratio to be one for any q. The multivariate variance ratio test of Chow and Denning (1993) addresses this issue. The ZV-statistic takes the maximum value among different Z-statistics and asymptotically follows the studentized maximum modulus distribution. From Table 5.4, we deduce that the ZV-statistic is 3.74, rejecting of the null hypothesis $VR(q) = 1$ for all q at the 5% level (the studentized maximum modulus distribution with 20 degrees of freedom at the 5% level is 3.64). The results of the variance ratio test provide further evidence that the random walk hypothesis can be rejected for the WMP futures prices.

Given the results of predictability of WMP futures prices, based on Kim et al. (2010), we expect that profit margin hedging not only reduces risk of the low milk revenue, but also enhances average milk payout. In the next section we present the findings of the simulations for a representative farm and a sample of individual farms.

5.3.3 Results for a Representative Farm

In this section, we simplify the setting of the farm's operations and assume a situation where a representative farm sells one tonne of WMP monthly over six seasons 2011 - 2017, which totals 72 months. Because the underlying asset of a single WMP futures is one tonne of WMP, such simplification means that if a farm chooses to hedge, it sells one futures contract. As discussed in Subsection 5.3.1, for profit margin hedging we set a target price to WMP break-even prices realized during the years 2011-2017. For each season, the target is different and defined in Table 5.2. We model two scenarios to account for basis risk: we assume a perfect foresight on the closing basis (no basis risk) or make a forecast of the closing basis as the average realized basis available at the decision time. The strategy is selective in a sense that a farm only hedges when futures prices are favourable. This scenario only arises when the net price is greater or equal to the target.

We conduct simulations to compare the profit margin hedging strategy with continuous hedging and no hedging strategies. To make the continuous hedging strategy comparable to the profit margin hedging strategy, we set up a rule where a farm hedges the entire position in the cash market, i.e. one tonne of WMP. We choose the hedging horizon to be 15 months, because it is equal to the average hedging horizon in the profit margin hedging strategy, as shown in the next paragraph.

In Table 5.5, we show the effect of hedging on risk and return of the payout to a farm, measured in WMP prices (USD per one tonne). The results show the superiority of profit margin hedging versus continuous hedging. If fees are zero, continuous hedging leads to an increase in the mean payout by 7.3%, while profit margin hedging increases the mean by 14.1%. An effect on the standard deviation of the payout is similar, a reduction by 28.9% and 31%, but a reduction in semivariance clearly demonstrates the difference. While profit margin hedging sets its objective to select the WMP futures prices which can deliver the dairy revenue above the total cash expenses, continuous hedging dictates to mechanically enter the futures market without any consideration of the WMP futures price. Results show that continuous hedging reduces semivariance by 57%, while profit margin hedging completely eliminates the downside risk. In Table 5.6, we conduct a paired t-test to assess whether an increase in means of monthly WMP prices between each pair of strategies is statistically different from zero. In the continuous hedge vs. no hedge pair (Panel A), the t-ratio is between 1.7 and 1.2, meaning that the price difference is not statistically different from zero at the 5% level. At the same time, in the profit margin hedge vs. continuous hedge pair (Panel C), the profit margin hedging average price is statistically higher than the price of continuous price at any conventional significance level. Based on these findings, we conclude that profit margin hedging leads to a higher mean value of the payout and lower downside risk in comparison to continuous hedging. In the subsection 5.3.4, where we perform the analysis for a sample of farm-level data, we concentrate on profit margin hedging only, based on the findings for the representative farm data.

Table 5.5: Alternative Hedging Strategies, Prices are in US$/MT

Risk/Return Characteristics of monthly WMP Prices					Effectiveness of Hedging			
Panel A: No Basis Risk								
Strategy	Mean	St Dev	Semivariance	CV	Mean	St Dev	Semivariance	CV
No Hedge	3,250	1,032	410	0.32				
Routine Hedge No Fee	3,488	734	176	0.21	7.3%	28.9%	57.0%	33.7%
Routine Hedge Low Fee	3,458	734	190	0.21	6.4%	28.9%	53.7%	33.2%
Routine Hedge Mid Fee	3,438	734	199	0.21	5.8%	28.9%	51.5%	32.8%
Routine Hedge High Fee	3,418	734	208	0.21	5.2%	28.9%	49.2%	32.4%
PM Hedge No Fee	3,708	712	0	0.19	14.1%	31.0%	100.0%	39.5%
PM Hedge Low Fee	3,686	713	0	0.19	13.4%	30.9%	100.0%	39.1%
PM Hedge Mid Fee	3,668	713	0	0.19	12.9%	30.9%	100.0%	38.8%
PM Hedge High Fee	3,659	712	0	0.19	12.6%	31.0%	100.0%	38.7%
Panel B: Basis Risk								
Strategy	Mean	St Dev	Semivariance	CV	Mean	St Dev	Semivariance	CV
No Hedge	3,250	1,032	410	0.32				
Routine Hedge No Fee	3,473	796	176	0.23	6.9%	22.8%	57.0%	27.8%
Routine Hedge Low Fee	3,443	796	190	0.23	5.9%	22.8%	53.7%	27.1%
Routine Hedge Mid Fee	3,423	796	199	0.23	5.3%	22.8%	51.5%	26.7%
Routine Hedge High Fee	3,403	796	208	0.23	4.7%	22.8%	49.2%	26.3%
PM Hedge No Fee	3,709	717	8	0.19	14.1%	30.5%	98.0%	39.1%
PM Hedge Low Fee	3,679	717	14	0.19	13.2%	30.5%	96.5%	38.6%
PM Hedge Mid Fee	3,659	717	19	0.20	12.6%	30.5%	95.4%	38.3%
PM Hedge High Fee	3,639	717	24	0.20	12.0%	30.5%	94.2%	37.9%

Note: This table presents the mean, standard deviation and semivariance of monthly WMP prices for cash sales, always hedge and profit margin hedging strategies. Panel A reports the results without basis risk and Panel B with basis risk.

Now, we discuss the results of the profit margin hedging strategy in more detail. Simulations show that the decision rule was satisfied 72 out of 72 months, and the short position was opened on average 15.5 (15.3) months before the payout is

Table 5.6: Paired Differences t-Ratios of Average Prices

No Basis Risk					Basis Risk				
Panel A: Routine Hedge vs. No Hedge									
	No Fee	Low Fee	Mid Fee	High Fee		No Fee	Low Fee	Mid Fee	High Fee
Mean	238	208	188	168	Mean	238	208	188	168
SD	1222	1222	1222	1222	SD	1222	1222	1222	1222
t-ratio	1.7	1.4	1.3	1.2	t-ratio	1.7	1.4	1.3	1.2
Panel B: Profit Margin Hedge vs. No Hedge									
Mean	458	436	418	409	Mean	459	429	409	389
SD	1294	1284	1282	1273	SD	1290	1290	1290	1290
t-ratio	3.0***	2.9***	2.8***	2.7***	t-ratio	3.0***	2.8***	2.7***	2.3***
Panel C: Profit Margin vs. Routine Hedge									
Mean	221	228	230	242	Mean	222	222	222	222
SD	353	363	362	360	SD	354	354	354	354
t-ratio	5.3***	5.3***	5.4***	5.7***	t-ratio	5.3***	5.3***	5.3***	5.3***

Note: This table presents paired differences of the expected price between always hedge and cash sales, profit margin hedging and cash sales, profit margin and always hedge for three different levels of fees. Under the null hypothesis t-statistic follows t-distribution with 71 degrees of freedom. The *** indicates significance at the 1% level.

due, in the case of no basis risk and basis risk scenarios, respectively. The results show that irrespective of how basis risk is modelled, profit margin hedging always increases returns and decreases risk. In the case of no basis risk and zero brokerage commission, the average price during 6 seasons is improved by 14.1%. The incorporation of brokerage commissions (USD 30, USD 50, USD 70) reduces the mean WMP price, but still results in an improved mean in comparison to the no hedge scenario. Implementation of profit margin hedging allows a substantial reduction in the volatility of prices, demonstrated by a decrease in the standard deviation of 31%. The coefficient of variation (CV) measures the ratio of standard deviation relative to returns. The strategy with the smallest CV is preferred. Results show that profit margin hedging provides the best risk-return trade-off. The semivariance, a measure of downside risk below the target, also strengthens the benefits of hedging. Because hedging was triggered for each month, in the case of no-basis risk the semivariance is reduced to zero, and when basis risk is taken into account, the semivariance is slightly above zero. To assess whether an increase in means of monthly WMP prices between two strategies is statistically different from zero, we conduct a paired t-test. From Table 5.6 Panel B, we can see that t-ratios range between 2.3 to 3, meaning that for each pair the price difference is significant at conventional levels.

In Table 5.7, we aggregate the resulting monthly payouts into average milk prices for each of the six seasons. We achieve this by averaging 12 payouts from July to June (we assume that milk collected in June is paid in July) for each season, converting the WMP USD averages to NZD and then converting it to the relevant milk price using Equation (5.1), which was used to convert BEMPs to break-even WMP prices. The main conclusion we draw is that profit margin hedging allows the representative farm to avert the turbulent times in the NZ dairy industry which occurred during seasons 2014-2015 and 2015-2016, when milk prices were lower than BEMPs. For example, the BEMP during the 2014-2015 season was NZD 5.77 per kgMS, which is higher than the milk cash price of NZD 4.46. However, in the no-basis risk, no fees scenario, the milk price with profit margin hedging is NZD 8.69. Thus, our findings demonstrate that profit margin hedging has the potential to support the financial viability of dairy business, without sacrificing average returns. The representative farm would have earned, on average, NZD 0.98 and 0.87 per kgMS in the case of no-basis risk, zero and high fees, respectively.

Table 5.7: Milk Price Per Season

Season	BEMP	Cash Sale	Hedge No Fee	Hedge Low Fee	Hedge Mid Fee	Hedge High Fee
Panel A: No Basis Risk						
2011-2012	5.74	6.08	6.61	6.55	6.51	6.47
2012-2013	5.98	6.88	6.79	6.73	6.70	6.66
2013-2014	6.35	8.36	6.48	6.51	6.48	6.56
2014-2015	5.77	4.46	8.69	8.63	8.59	8.55
2015-2016	4.93	3.99	7.42	7.36	7.31	7.27
2016-2017	5.17	5.99	5.64	5.57	5.52	5.49
Average	5.66	5.96	6.94	6.89	6.85	6.83
Panel B: Basis Risk						
2011-2012	5.74	6.08	6.61	6.55	6.51	6.47
2012-2013	5.98	6.88	6.78	6.72	6.68	6.64
2013-2014	6.35	8.36	6.53	6.48	6.44	6.40
2014-2015	5.77	4.46	8.69	8.63	8.59	8.55
2015-2016	4.93	3.99	7.42	7.36	7.31	7.27
2016-2017	5.17	5.99	5.60	5.53	5.49	5.44
Average	5.66	5.96	6.94	6.88	6.84	6.79

Note: This table presents the estimated price of milk per season with and without profit margin hedging. The price of milk is calculated as an average WMP price per season and then converted to milk price using Equation (5.1).

5.3.4 Results for Individual Dairy Farms

Individual farms may have BEMPs different from a representative farm as production costs, debt structure, and cash flow requirements vary across individual farms. The advantage of profit margin hedging is that it can be tailored to the specific needs of a farm, based on its unique cash flow requirements. In this section, we test our hedging strategy at the farm level. We begin by discussing the data, implementation of the strategy, followed by the results.

The DairyBase database was established by DairyNZ in 2005 and contains NZ dairy farms' financial and physical data. Participation is voluntary, meaning that farms may not report their information in all years. Our dataset contains owner-operator farm data from the Waikato and Marlborough-Canterbury regions for six seasons between 2011 and 2017. We choose these regions because they are the biggest regions of the North and South Islands of NZ, respectively, measured by the number of herds. For the 2016-2017 season, the Waikato region makes up 31.5% of all owner-operator farms and Marlborough-Canterbury makes up 11.4% (DairyNZ, 2017). For the six season period, there are 608 farms that reported financial information for at least one of the seasons. However, only 50 farms consistently reported for all six seasons. If we remove the 2011-2012 season, we almost double the number to 92 farms. This motivates us to conduct further analysis using only data for the five seasons from 2012 to 2017. To assess how representative the selected farms are, we compare several profitability statistics and physical characteristics of the regional data from DairyNZ Economic Surveys to the averages of the sample. We find that the selected farms have slightly larger herds and milksolids production per cow; however, farm working expenses and operating profits are very close to regional averages. For the hedging analysis, profitability statistics are more important than physical characteristics, and thus, we conclude that the selected sample is a good representation of the regional data.

To implement the strategy, we need to define the break-even WMP prices and the quantity of produced milksolids for the cross-section of farms for each of the five seasons. We resort to ex-post analysis, that is, we find the break-even prices and output from realized data.[11] For example, to find the break-even WMP price and output for the season 2012-2013 for a specific farm, we use information reported by the farm for the 2012-2013 season. This approach allows us to concentrate only on price uncertainty, keeping output fixed. The DairyBase database reports total milksolids produced for a season, but does not specify the production for each month during the season. We assume that for each farm the distribution of milk production during a season is the same as the national average, provided by the Agri Data database.[12] To calculate monthly production for each farm, we extract monthly NZ milk production data from the Agri Data database for the seasons between 2012 and 2017 and find the fraction of each month's production relative to the total of the season. We then find averages for each month across five seasons and use them to distribute individual farm milk production across a season.[13] We also assume that futures are completely divisible, that is, a farm can sell any number of futures contracts.

Similar to the analysis in the previous section, we consider scenarios with no, low, medium and high levels of brokerage fees: USD 0, USD 30, USD 50, and USD 70 per contract per round-trip transaction, respectively. The average annual milksolids production per farm is 216 tonnes, which translates into NZD 6,480, NZD 10,800 and NZD 15,120 annual brokerage expense for a farm. To assess the effect of hedging, we analyse the change in discretionary cash for each farm with and without a hedge. Discretionary cash is what is left after farm working expenses, rent, interest, and tax are paid and net income from non-dairy farming activities is added. Discretionary cash is available for drawings, debt repayments, capital development, and purchases. We define an occurrence of negative discretionary cash during a season as financial distress. To calculate the value of discretionary cash after the implementation of a

[11] We can infer the break-even prices only from the ex-post analysis for the following reasons: in case we set the break-even price for a new season, for example, to the same as in a previous season, but a farm decides to expand, for example, his farm working expenses, when, the discretionary cash in the new season is likely to be negative. For the profit margin hedging strategy to work, a farm needs to prepare a budget before the start of the season and control the spending according to the prepared budget.

[12] The Agri Data database is a part of the NZX Research Centre database.

[13] For example, to find the fraction for January, for each of the five seasons we divide NZ total production for January by NZ total production during that season. Then we find the fraction for January as the average of January's fractions across five seasons. We use the obtained fractions for each month to distribute the total production of each farm across months.

profit margin hedge for a given farm, and for a given season, we add the profit/loss from profit margin hedging during the season to the discretionary cash realized during that season. WMP futures are priced in USD and, therefore, we need to convert a profit/loss to NZD. We use the annual average USD:NZD conversion rates from Fonterra's Farmgate Milk Price Statements.[14]

5.3.4.1 The Effect of Profit Margin Hedging

To assess the effect of profit margin hedging on the average discretionary cash over the 2012-2017 seasons, we start with a visual assessment of its distribution and compare it to the no hedge case. In Figures 5.3 and 5.4, we plot the empirical cumulative distribution functions (CDFs) as well as kernel density estimations for the average discretionary cash in NZD for the no hedge and profit margin hedging strategies, respectively. Four observations about the CDF plots deserve to be noted. First, the graphs show that the CDF for the no hedge lies to the left of the CDFs for profit margin hedging. This means that for each fixed value of discretionary cash, the probability to observe a value smaller than that is higher in the case without hedging. For example, the probability to observe discretionary cash below or equal to NZD 200,000 is equal to 55% under profit margin hedging and 72% for the no hedge case. Second, no hedge has more negative outcomes and less positive outcomes. Third, fees do not outweigh the benefits of hedging. Fourth, basis risk does not substantially affect the results. Similar conclusions can be drawn from the density estimate graphs, which are presented in the bottom panels of Figures 5.3 and 5.4. First, with profit margin hedging, the distribution moves to the right, indicating an increase in the mean of discretionary cash. Second, the basis risk does not substantially change the results, and lastly, fees do not qualitatively change the conclusions about the benefits of hedging.

After this preliminary assessment of the differences between two alternatives, we

[14] As a robustness check, we use forex spot rates to convert profits/losses from monthly futures settlements. We obtained the exchange rates from Datastream. We find that this approach does not impair, but in fact improves the results. We attribute the finding to the fact that during 2014-2015 and 2015-2016 seasons, Fonterra's average exchange rate was 1% and 6% lower than the spot rates. While during 2012-2013, 2013-2014 and 2016-2017 seasons Fonterra locked better exchange rates than spot rates, 2014-2015 and 2015-2016 are the most important seasons, as the gains from hedging during these seasons were the biggest. Results are available upon request.

Figure 5.3: Kernel Density Estimation and Cumulative Distribution Function of Discretionary Cash without Basis Risk

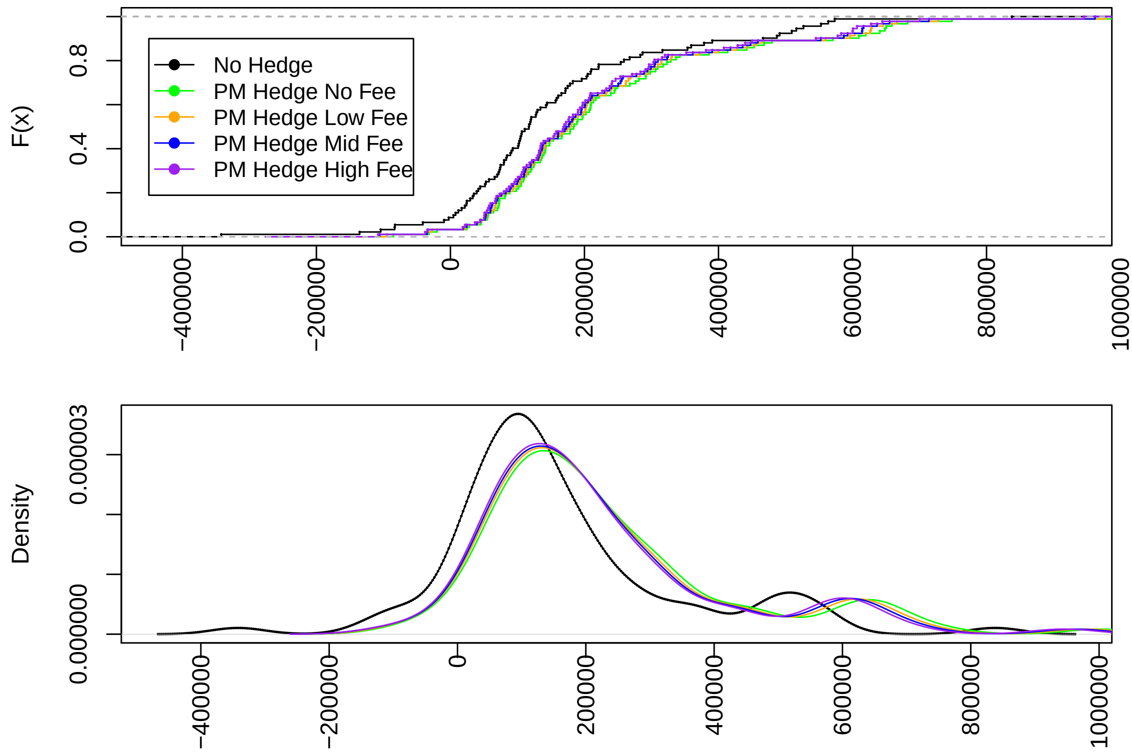

Note: The x-axis of the plots shows the mean value of discretionary cash over 2012-2017 seasons in NZD. The y-axis of the bottom graph shows the probability $F(x) = P(Discretionary\,Cash\ \leq x)$.

perform a Kolmogorov-Smirnov (K-S) test to compare the two alternatives. In Table 5.8, we show the results for no hedge versus profit margin hedging with different levels of commissions. The null hypothesis of K-S is that the two data sets are drawn from the same distribution. The results show that, irrespective of fees and basis risk, we reject the null hypothesis that discretionary cash for no hedging strategy and profit margin hedging is from the same distribution. Therefore, we conclude that the observed shifts in distributions between the no hedge and profit margin hedging strategies in Figures 5.3 and 5.4 are statistically significant.

5.3.4.2 The Effect of Profit Margin Hedging on Risk and Return of Discretionary Cash

As a next step we calculate the mean, volatility and semivariance of discretionary cash for each farm with and without profit margin hedging. Mean, volatility and semivariance values are calculated for discretionary cash during five seasons and,

Figure 5.4: Kernel Density Estimation and Cumulative Distribution Function of Discretionary Cash without Basis Risk

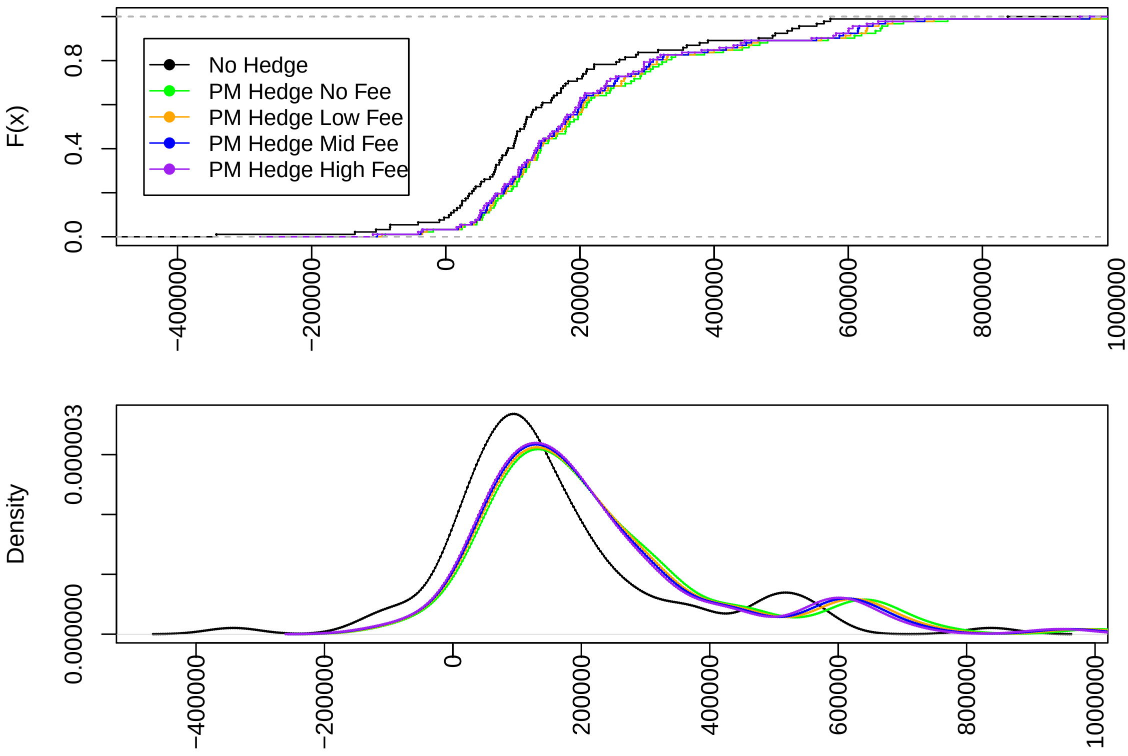

Note: The x-axis of the plots shows the mean value of discretionary cash over 2012-2017 seasons in NZD. The y-axis of the bottom graph shows the probability $F(x) = P(Discretionary\ Cash\ \leq x)$.

Table 5.8: Kolmogorov-Smirnov Test for Discretionary Cash with and without Profit Margin Hedging

	No hedge vs. hedge no fee	No hedge vs. hedge low fee	No hedge vs. hedge mid fee	No hedge vs. hedge high fee
Panel A: No Basis Risk				
K-S statistic	0.239***	0.228***	0.228***	0.217***
p-value	0.010	0.016	0.016	0.026
Panel B: Basis Risk				
K-S statistic	0.239***	0.228***	0.217***	0.217***
p-value	0.010	0.016	0.026	0.026

Note: This table presents the results of the Kolmogorov-Smirnov test which measures the difference between two cumulative distribution functions. The null hypothesis is that two samples are drawn from the same distribution. *** indicates significance at the 1% level.

thus, represent annual values. We present the results in Tables 5.9 and 5.10 where we group individual farms into quintiles based on the volatility (Panel A), semivariance (Panel B), and mean (Panel C) of discretionary cash without hedge. When we group individual farms into quintiles, we show the average of the individual farm's volatility, semivariance and mean of discretionary cash. As an example, in the first row and the first column of Panel A of Table 5.9 we show the average volatility of discretionary cash without hedge among farms with the smallest 20% volatility (i.e. low quintile) of discretionary cash without hedge. The average volatility of discretionary cash for the low quintile of individual farms is equal to NZD 109,309 per year. In the following rows we show how volatility of discretionary cash changes for the same group of farms if the profit margin hedging strategy is used. In the last column, we present the average volatility of discretionary cash for all individual farms grouped together.

Table 5.9: Hedging Effectiveness for Individual Farms in the No Basis Risk Scenario

Quintiles	Low	2	3	4	High	All	Low	2	3	4	High	All
Panel A:	Volatility of Discretionary Cash						Reduction in Volatility of Discretionary Cash					
No hedge	109,309	180,562	255,860	410,700	678,701	328,482						
Hedge no fee	86,362	135,074	194,525	244,391	403,944	213,561	21%**	25%***	24%***	40%***	40%***	35%***
Hedge low fee	85,818	134,620	193,803	243,709	402,909	212,872	21%**	25%***	24%***	41%***	41%***	35%***
Hedge mid fee	85,614	133,993	193,284	243,366	402,300	212,412	22%**	26%***	24%***	41%***	41%***	35%***
Hedge high fee	85,415	133,570	192,442	242,379	401,942	211,857	22%**	26%***	25%***	41%***	41%***	36%***
Panel B:	Semivariance of Discretionary Cash						Reduction in Semivariance of Discretionary Cash					
No hedge	745	27,739	80,066	135,241	316,757	113,124						
Hedge no fee	2,474	7,726	14,968	24,073	76,406	25,209	-232%	72%***	81%***	82%***	76%***	78%***
Hedge low fee	2,631	7,726	16,221	24,491	81,314	26,814	-253%	72%***	80%***	82%***	74%***	76%***
Hedge mid fee	2,736	8,559	16,726	25,172	85,272	28,048	-267%	69%***	79%***	81%***	73%***	75%***
Hedge high fee	2,785	9,650	17,492	25,841	88,076	29,131	-274%	65%***	78%***	81%***	72%***	74%***
Panel C:	Mean of Discretionary Cash						Increase in Mean of Discretionary Cash					
No hedge	- 28,953	71,274	116,562	188,264	443,343	159,165						
Hedge no fee	48,298	128,875	178,067	259,654	550,718	234,566	267%***	81%***	53%***	38%***	24%***	47%***
Hedge low fee	43,640	124,296	171,479	252,082	536,296	226,959	251%***	74%***	47%***	34%***	21%***	43%***
Hedge mid fee	40,646	120,442	167,796	247,478	526,718	221,987	240%***	69%***	44%***	31%***	19%***	39%***
Hedge high fee	38,355	116,930	163,134	243,027	517,152	217,068	232%***	64%***	40%***	29%***	17%***	36%***

Note: This table presents the effect of hedging on individual farms in the no basis risk scenario. We group farms into quintiles of annual volatility, semivariance and mean of discretionary cash for unhedged position. We use a paired t-test to compare the difference in means of volatility, semivariance and mean of discretionary cash between no hedge and profit margin hedging strategies. **, *** indicates significance at the 5% and 1% levels, respectively.

We first discuss the results for the no basis risk scenario which are presented in Table 5.9. Panel A shows that the volatility of yearly discretionary cash is reduced by between 21% to 40%, with a mean value of 35% for the whole sample, ignoring fees. Volatilities of discretionary cash with hedge and different fees are very similar, as the fees expense is fairly constant for each season, and hence only slightly affects the volatility of discretionary cash. We conduct a paired t-test to find out if, on average, profit margin hedging leads to an improvement in volatility within each quintile and for the whole sample. We find that the average differences in means are statistically different from zero at the 5% level. Panel B of Table 5.9 presents the semivariance of discretionary cash. We measure semivariance relative to a threshold of zero. Discretionary cash below zero means that milk sales do not cover all farm working expenses, rent, interest, and taxes, and therefore the farm does not have any cash to make withdrawals, debt repayments or capital developments. We find that for the quintiles two to high the semivariance is reduced by 72% to 82%, with an average of 78%, and, only for the low group, we observe an increase in the semivariance. Intuitively, a profit margin hedge with a target price equal to break-even price, discretionary cash should always be positive. However, as our study uses a cross-hedge, the dynamics of the milk price are not fully explained by the WMP price and so there remains some risk of receiving a milk payout below BEMP. We find that commissions increase the semivariance and reduce the hedging effectiveness, on average by 2%, 3%, and 4% (moving from low to high fees). A paired t-test shows that the average differences in means of the semivariance between profit margin hedging and no hedging are statistically different from zero at the 5% level in all quintiles, except for the low quintile. Therefore, we conclude that the reduction in the semivariance of discretionary cash is significant at the 5% level. Panel C of Table 5.9 shows the effect of profit margin hedging on the mean of discretionary cash. We can see that without hedging the mean of discretionary cash for the low group is NZD -\$28,953. After the hedge is in place, the mean goes up by 267% to NZD 48,298. Fees reduce the mean by roughly NZD 5,000, NZD 8,000 and NZD 10,000 (moving from low to high fees), respectively. For other quintiles, the mean increases by 81%, 53%, 38% and 24% (moving from second to high quintile) in the case of no fees, and by 64%, 40%, 29% and 17% in the case of high fees. This shows that the highest percentage increase in the mean occurs for the low group and effectiveness decreases as percentile group increases. These results further show that profit margin hedging delivers the strongest benefits for the most vulnerable group of farms. The increase for the whole sample ranges between 36% and 47%, depending on the level of fees. A paired t-test shows that, on average, profit margin hedging leads to an increase in the mean of discretionary cash within each quintile and for the whole sample.

Table 5.10 presents the results when a farm does not know the value of the closing basis in advance and has to predict it. When we compare Tables 5.9 and 5.10, we find that a change in basis risk expectations does not change the results. In fact, the numbers are very similar, which means that the basis risk in WMP futures does not diminish the value of hedging. A paired t-test gives similar results to a no basis risk case, i.e. we find that the average difference between profit margin hedging mean, variance and semivariance of discretionary cash and no hedging mean, variance and semivariance of discretionary cash are statistically different from zero at the 5% level in all quintiles, except for the low quintile for semivariance of discretionary cash.

Table 5.10: Hedging Effectiveness for Individual Farms in the Basis Risk Scenario

Quintiles	Low	2	3	4	High	All	Low	2	3	4	High	All
Panel A:	Volatility of Discretionary Cash						Reduction in Volatility of Discretionary Cash					
No hedge	109,309	180,562	255,860	410,700	678,701	328,482						
Hedge no fee	86,745	134,938	194,599	245,513	404,170	213,895	21%**	25%***	24%***	40%***	40%***	35%***
Hedge low fee	86,041	134,491	192,713	243,972	403,901	212,936	21%**	26%***	25%***	41%***	40%***	35%***
Hedge mid fee	85,550	133,941	192,237	243,616	403,462	212,473	22%**	26%***	25%***	41%***	41%***	35%***
Hedge high fee	85,477	133,619	192,201	242,770	402,349	211,993	22%**	26%***	25%***	41%***	41%***	35%***
Panel B:	Semivariance of Discretionary Cash						Reduction in Semivariance of Discretionary Cash					
No hedge	745	27,739	80,066	135,241	316,757	113,124						
Hedge no fee	2,474	7,902	15,582	23,536	77,835	25,515	-232%	72%***	81%***	83%***	75% ***	77%***
Hedge low fee	2,631	7,902	16,109	23,185	81,816	26,674	-253%	72%***	80%***	83%***	74%***	76%***
Hedge mid fee	2,736	8,568	16,748	22,799	85,849	27,709	-267%	69%***	79%***	83%***	73%***	76%***
Hedge high fee	2,841	9,479	17,946	24,006	89,050	29,040	-281%	66%***	78%***	82%***	72%***	74%***
Panel C:	Mean of Discretionary Cash						Increase in Mean of Discretionary Cash					
No hedge	- 28,953	71,274	116,562	188,264	443,343	159,165						
Hedge no fee	48,262	127,609	177,336	259,662	550,587	234,142	267%***	79%***	52%***	38%***	24%***	47%***
Hedge low fee	44,081	124,150	170,933	251,764	536,120	226,816	252%***	74%***	47%***	34%***	21%***	43%***
Hedge mid fee	40,965	121,041	166,471	246,894	526,369	221,725	241%***	70%***	43%***	31%***	19%***	39%***
Hedge high fee	37,890	117,464	163,539	241,599	516,814	216,807	231%***	65%***	40%***	28%***	17%***	36%***

Note: This table presents the effect of hedging on individual farms in the scenario with basis risk. We group farms into quintiles of annual volatility, semivariance and mean of discretionary cash flow for unhedged position. We use a paired t-test to compare the difference in means of volatility, semivariance and mean of discretionary cash between no hedge and profit margin hedging strategies. **, *** indicates significance at the 5% and 1% levels, respectively.

5.3.4.3 The Economic Effect of Profit Margin Hedging

The previous analysis is based on data of 92 farms, which make up 1.09% of the milk production across NZ. To assess the economic effect of profit margin hedging for the NZ economy, we want to know what the dollar value that the profit margin hedging strategy could have generated over the 2012-2017 seasons if it were adopted by all NZ dairy farms. Our approach is to scale up the profit/loss generated by hedging during each season for 92 farms where the scaling factor is the milk produced by the 92 farms relative to the total milk produced in NZ during each season. Specifically, for each season, we calculate the total profit/loss from the futures position across all farms in the sample and scale it to the profit/loss of all NZ farms by the proportion of milk production in the sample to the total milk production of NZ. Then, we take the average across all five seasons. We find that profit margin hedging could have generated NZD 0.49, NZD 0.54 or NZD 0.58 billion average per year (moving from high to low commission), with perfect knowledge of the basis, or NZD 0.49, NZD 0.53, NZD 0.58 billion average per year without knowledge of the basis. Given that the average yearly dairy export revenue is NZD 13.2 billion, the average gain of NZD 0.49 billion translates to 3.74% (0.49/13.2) of the yearly dairy export revenue.

5.3.4.4 The Effect of Leverage

Literature shows that capital structure plays an important role in explaining the adoption of futures by farms (e.g. Turvey and Baker, 1989; Shapiro and Brorsen, 1988; Wolf and Widmar, 2014). Low-leveraged farms are less likely to hedge, as they are financially more secure. Based on this literature, we want to address two questions. The first is whether the level of leverage affects the level of volatility, semivariance and mean of discretionary cash. The second is whether there is a relation between leverage and hedging effectiveness.

To address the first question, we sort farms into quintiles by their leverage ratio and calculate the mean volatility, semivariance and mean of discretionary cash in each quintile group. We use a two-sample t-test between high and low quintiles to find whether the differences are significant. Tables 5.11 and 5.12 show that the difference in mean semivariance between the high-minus-low leverage ratio quintiles is NZD 113,434 (t-stat 6.42) and the difference in mean of discretionary cash is -NZD 112,941 (t-stat -2.29). We find that the difference in volatility is not statistically significant. We find that after implementing the profit margin hedging strategy, the difference remains significant for the semivariance, but not for the mean. For the no fees hedge, an increase in the mean for the high leverage quintile is 66%, while for the low leverage it is 23%. This finding indicates that hedging reduces the gap between the mean of discretionary cash of different quintile groups.

Table 5.11: Effect of Hedging on Discretionary Cash for Individual Farms Grouped by Leverage without Basis Risk

Quintiles	Low	2	3	4	High	All	High - Low	t-stat	Low	2	3	4	High	All
Debt to Asset	31%	41%	52%	64%	90%	47%			31%	41%	52%	64%	90%	47%
Panel A:	Volatility of Discretionary Cash								Reduction in Volatility of Discretionary Cash					
No hedge	217,762	324,713	302,947	487,880	315,956	328,482	98,194	1.73						
Hedge no fee	169,658	178,898	216,106	270,959	233,516	213,561	63,857	1.32	22%***	45%***	29%***	44%***	26%***	35%***
Hedge low fee	169,110	178,352	215,799	269,880	232,555	212,872	63,445	1.31	22%***	45%***	29%***	45%***	26%***	35%***
Hedge mid fee	168,939	177,717	215,305	269,488	231,943	212,412	63,003	1.30	22%***	45%***	29%***	45%***	27%***	35%***
Hedge high fee	168,679	177,242	214,962	268,277	231,436	211,857	62,757	1.29	23%***	45%***	29%***	45%***	27%***	36%***
Panel B:	Semivariance of Discretionary Cash								Reduction in Semivariance of Discretionary Cash					
No hedge	10,525	105,606	103,565	227,062	123,959	113,124	113,434***	6.42						
Hedge no fee	2,009	15,508	24,718	54,658	30,168	25,209	28,158**	2.53	81%**	85%***	76%***	76%***	76%***	78%***
Hedge low fee	2,148	16,310	26,217	58,916	31,581	26,814	29,433**	2.56	80%**	85%***	75%***	74%***	75%***	76%***
Hedge mid fee	2,145	17,057	26,857	62,590	32,765	28,048	30,620**	2.61	80%**	84%***	74%***	72%***	74%***	75%***
Hedge high fee	2,239	17,805	27,879	64,816	34,133	29,131	31,893**	2.67	79%**	83%***	73%***	71%***	72%***	74%***
Panel C:	Mean of Discretionary Cash								Increase in Mean of Discretionary Cash					
No hedge	228,204	170,924	155,504	124,535	115,263	159,165	-112,941**	- 2.29						
Hedge no fee	281,252	237,935	222,602	239,509	191,339	234,566	-89,912	- 1.53	23%***	39%***	43%***	92%***	66%***	47%***
Hedge low fee	275,489	230,758	215,498	228,820	183,924	226,959	-91,565	- 1.61	21%***	35%***	39%***	84%***	60%***	43%***
Hedge mid fee	271,456	226,326	210,996	221,797	179,001	221,987	-92,454	- 1.64	19%***	32%***	36%***	78%***	55%***	39%***
Hedge high fee	267,453	221,852	206,364	215,429	173,846	217,068	-93,606	- 1.69	17%***	30%***	33%***	73%***	51%***	36%***

Note: This table presents the effect of hedging on individual farms in the scenario with no basis risk. We group farms into quintiles of average leverage ratio. We use a paired t-test to compare the difference in means of volatility, semivariance and mean of discretionary cash between no hedge and profit margin hedging strategies. We use a two-sample t-test to compare the difference in means of volatility, semivariance and mean of discretionary cash between high and low quintile of leverage ratio. **, *** indicates significance at the 5% and 1% levels, respectively.

Table 5.12: Effect of Hedging on Discretionary Cash for Individual Farms Grouped by Leverage with the Basis Risk

Quintiles	Low	2	3	4	High	All	High - Low	t-stat	Low	2	3	4	High	All
Debt to Asset	31%	41%	52%	64%	90%	47%			31%	41%	52%	64%	90%	47%
Panel A:	Volatility of Discretionary Cash Flow								Reduction in Volatility of Discretionary Cash					
No hedge	217,762	324,713	302,947	487,880	315,956	328,482	98,194	1.73						
Hedge no fee	169,851	179,358	216,306	271,999	233,325	213,895	63,474	1.31	22%***	45%***	29%***	44%***	26%***	35%***
Hedge low fee	169,325	178,468	215,722	270,254	232,258	212,936	62,933	1.30	22%***	45%***	29%***	45%***	26%***	35%***
Hedge mid fee	169,003	178,272	215,426	269,536	231,487	212,473	62,484	1.29	22%***	45%***	29%***	45%***	27%***	35%***
Hedge high fee	168,784	177,654	215,427	268,386	231,054	211,993	62,270	1.28	22%***	45%***	29%***	45%***	27%***	35%***
Panel B:	Semivariance of Discretionary Cash Flow								Reduction in Semivariance of Discretionary Cash					
No hedge	10,525	105,606	103,565	227,062	123,959	113,124	113,434***	6.42						
Hedge no fee	2,026	15,452	25,520	56,090	29,565	25,515	27,539**	2.47	81%**	85%***	75%***	75%***	76%***	77%***
Hedge low fee	2,161	16,357	26,068	59,001	30,909	26,674	28,748**	2.48	79%**	85%***	75%***	74%***	75%***	76%***
Hedge mid fee	2,259	17,108	26,893	61,587	31,878	27,709	29,618**	2.50	79%**	84%***	74%***	73%***	74%***	76%***
Hedge high fee	2,359	17,892	27,967	64,769	33,450	29,040	31,091**	2.58	78%**	83%***	73%***	71%***	73%***	74%***
Panel C:	Mean of Discretionary Cash Flow								Increase in Mean of Discretionary Cash					
No hedge	228,204	170,924	155,504	124,535	115,263	159,165	-112,941**	-2.29						
Hedge no fee	281,076	237,609	221,868	238,619	191,309	234,142	- 89,767	-1.52	23%***	39%***	43%***	92%***	66%***	47%***
Hedge low fee	275,097	230,385	215,907	228,239	184,140	226,816	- 90,957	-1.60	21%***	35%***	39%***	83%***	60%***	43%***
Hedge mid fee	271,496	225,705	210,851	220,935	179,231	221,725	- 92,265	-1.64	19%***	32%***	36%***	77%***	55%***	39%***
Hedge high fee	267,359	221,278	205,991	214,704	174,256	216,807	- 93,103	-1.68	17%***	29%***	32%***	72%***	51%***	36%***

Note: This table presents the effect of hedging on individual farms in the scenario with basis risk. We group farms into quintiles of average leverage ratio. We use a paired t-test to compare the difference in means of volatility, semivariance and mean of discretionary cash between no hedge and profit margin hedging strategies. We use a two-sample t-test to compare the difference in means of volatility, semivariance and mean of discretionary cash between high and low quintile of leverage ratio. **, *** indicates significance at the 5% and 1% levels, respectively.

As for the second question, we want to assess whether we can establish a relation between the level of leverage and hedging effectiveness. From Panels A and C of Table 5.11, we conclude that farms with low debt to asset ratios benefit the least from hedging, as the reduction in volatility and increase in the mean of discretionary cash is the smallest in comparison to other quintiles. Another observation is that for the fourth and highest quintiles, the improvement in the mean of discretionary cash is the strongest. We conclude the same from Table 5.12, which groups data into quintiles of the debt to asset ratio in the case of profit margin hedging with basis risk. From Table 5.11, we find that highly leveraged farms in the fourth and highest quintiles (i.e. farms with debt to asset ratios between 52% and 90%), have the biggest increase in the mean of discretionary cash by 92% and 66% (average of 79%) without fees, respectively, and by 73% and 51% (average of 62%) with high fees. Farms with low debt to asset ratios, i.e. below 31%, experience the smallest improvement in mean discretionary cash by 23% without fees and by 17% with high fees. These farms also experience the smallest reduction in volatility by about 22%. A paired t-test indicates that the improvements in the mean of volatility, semivariance and mean of discretionary cash are significant at the 5% level.

Based on our results, we conclude that the level of leverage is an important variable for farms, that adopt the hedging strategy. While we find that profit margin hedging decreases risk and increases returns for farms at all levels of leverage, farms with low leverage are the least advantaged by hedging, while farms with high leverage benefit most.

5.3.4.5 Probability of Financial Distress

Table 5.13 presents a simple measure which allows us to evaluate the effect of profit margin hedging on the probability of financial distress during a given year. We define financial distress as the inability to cover farm working expenses, interest, rent and tax payments from dairy cash income. Quantitatively this is measured by the occurrence of negative discretionary cash. If a farm chooses not to hedge, 159 out of 460 observations (92 farms during 5 seasons) are characterized by negative discretionary cash, i.e. a proportion of 34.6%. If profit margin hedging is implemented, the number of observations with negative discretionary cash decreases by more than double, to 63, which is 13.7% of the total sample. Depending on the magnitude of fees, this proportion increases to 14.8, 15.4 and 16.3 for the low, mid and high fees scenarios, respectively. The results do not substantially change after incorporating basis risk.

Table 5.13: Frequency of Negative Discretionary Cash Occurrence

	No Basis Risk	Basis Risk
No hedge	34.6%	34.6%
Hedge no fee	13.7%	13.9%
Hedge low fee	14.8%	14.8%
Hedge mid fee	15.4%	15.4%
Hedge high fee	16.3%	16.5%

Note: This table presents the frequency of negative discretionary cash among 92 farms during five seasons 2012-2017, totalling to 460 observations.

Overall, the results of this subsection demonstrate that our profit margin hedging strategy decreases risk and improves returns for a sample of NZ dairy farms. We also find that WMP futures do not bear high basis risk, and when we model different scenarios of basis risk we find qualitatively similar results.

5.4 Conclusion

In this study, we examine the effectiveness of profit margin hedging for NZ dairy farms. We demonstrate how the WMP futures can be used to protect farms from price risk. We base our results on historical data available for the period 2011 to 2017. We start by showing that prices of WMP futures do not follow a random walk. According to Kim et al. (2010) this result means that profit margin hedging can also be used as a tool to increase the average milk price.

In the first part of the analysis, we evaluate profit margin hedging from the perspective of a representative farm. We compare the risk and return of the average monthly payout expressed in WMP price between profit margin hedging, no hedging and continuous hedging strategies. We find that profit margin hedging delivers the highest average payout and lowest semivariance. We further find that depending on fees and basis risk, the expected return is increased by between 12% to 14.1%, the variance is reduced between 30.5% to 31.0%, and that almost all downside risk is eliminated. We find that profit margin hedging shows especially reliable results in reducing the downside risk, thus helping us to maintain the financial viability of dairy farm operations.

In the second part of the analysis, we implement profit margin hedging, using actual data, for a sample of individual farms. The results show that in the case of no basis risk and zero fees, the mean value of annual discretionary cash for all farms is increased by 47%, volatility is reduced by 35% and downside risk, measured by semivariance, is reduced by 78%. Although the introduction of fees reduces the increase in returns and reduction in risk, profit margin hedging still offers a significant improvement over no hedging. We find that highly leveraged farms, which have debt to asset ratios above 52%, see the largest increase in mean discretionary cash by 79% without fees and by 62% with high fees. Additionally, we show that profit margin hedging reduces the probability of financial distress during a given year by more than half, from 35% to 16%. To estimate the economic effect of the profit margin hedging strategy, we scale up the profit generated by this strategy across the sample of farms to all NZ dairy farms. We estimate that the strategy could have generated NZD 0.49 billion yearly average over a five year period, which is 3.7% of the yearly dairy export revenue.

This study has several important implications. We document that WMP futures offer significant benefits for NZ dairy farms. We demonstrate that profit margin hedging enhances the sustainability of the farming business, by reducing uncertainty about future profit. Reduced certainty about profit can negatively impact investment and production planning decisions, restrict access to capital and threaten solvency. High indebtedness of the NZ dairy farm sector makes it vulnerable to low dairy prices, and the results of our study can be of interest for policy-makers who are concerned with financial stability.

67

5.A Appendix: Variance Ratio Test

Let p_t denote the log price process and let a sample consist of $nq + 1$ observations, where p_0 and p_{nk} are the first and the last observations and q is any integer greater than one. Lo and MacKinlay (1988) show that the variance ratio statistic of q-period returns can be calculated as:

$$VR(q) = \frac{\sigma^2(q)}{\sigma^2},$$

where $\sigma^2(q)$ is an unbiased estimator of $1/q$ of the variance of the q-period returns and σ^2 is an unbiased estimator of the variance of the one-period returns and defined by:

$$\begin{aligned}
\sigma^2(q) &= \frac{1}{m}\sum_{k=q}^{nq}(p_k - p_{k-q} - q\mu)^2 \\
\sigma^2 &= \frac{1}{nq-1}\sum_{k=1}^{nq}(p_k - p_{k-1} - \mu)^2 \\
\mu &= \frac{1}{nq}(p_{nq} - p_0) \\
m &\equiv q(nq - q + 1)(1 - \frac{q}{nq}).
\end{aligned}$$

A test statistic $Z(q)$ is adjusted for heteroscedasticity in returns and defined by:

$$Z(q) = \frac{\sqrt{nq}(VR(q) - 1)}{\sqrt{\theta}} \overset{a}{\sim} N(0,1),$$

where θ is asymptotic variance of variance ratio $VR(q)$:

$$\begin{aligned}
\theta &\equiv 4\sum_{k=1}^{q-1}\left(1 - \frac{k}{q}\right)^2 \delta_k \\
\delta_k &= \frac{nq\sum_{j=k+1}^{nq}(p_j - p_{j-1} - \mu)^2(p_{j-k} - p_{j-k-1} - \mu)^2}{\left[\sum_{j=1}^{nq}(p_j - p_{j-1} - \mu)^2\right]^2}.
\end{aligned}$$

$Z(q)$ is asymptotically normally distributed with mean zero and standard deviation of one.

Chow and Denning (1993) derived the multivariate variance ratio test where the null hypothesis is that $VR(q_i)$ equals one for all i = 1,...,l. The test statistic is defined by:

$$ZV = \max_{1 \leq i \leq l} |Z(q_i)|,$$

which asymptotically follows the studentised maximum modulus distribution under the random walk null hypothesis.

Chapter 6

Concluding Remarks

In the last fifty years, derivatives markets saw unprecedented growth. While the US stock exchanges continue to attract the majority of trading volume, exchanges worldwide also demonstrate an expansion in the variety and complexity of offered derivatives. This variety of products, further fostered by decreased costs and increased speed of trading, allows different market participants to use derivatives for their benefits. However, as new products are launched frequently, it is important to evaluate the usefulness and properties of these products. This thesis focuses on relatively new derivatives, which were introduced in the last ten years: volatility and dairy derivatives. In the first essay, we investigate the informational leadership between the two most liquid volatility ETNs - the VXX and XIV. We find that a price leadership in these two competing markets is time-varying and tends to fluctuate from one market to the other. The results suggest that informed traders prefer to trade in the more liquid market and with the lowest transaction costs. The other interesting finding is that the informativeness of the VXX decreases on days when the "the market's fear gauge" (VIX) increases and on days with negative returns in the S&P 500. In the second essay, we investigate the informational content of the WMP options market. We construct the dairy implied volatility index (the DVIX) using prices of options and futures on WMP. In brief, we find that the DVIX contains useful information about future realized volatility. However, the market is not yet fully informationally efficient, as forecasts based on historical price data complement the DVIX in forecasting future realized volatility. In the third essay, we investigate the role of WMP futures from a risk-management perspective. We develop a profit margin hedging strategy, with the purpose of protecting the financial position of dairy farms from financial distress. To analyse the effect of the hedging strategy on risk and return of dairy farms, we employ data of the representative farm and a sample of individual farms. In brief, we find that the profit margin hedging strategy leads to a decreased level of risk and increased returns. The results are robust to the presence of basis risk and to different levels of transaction fees. We also show that the developed strategy outperforms the routine hedging, which assumes selling futures contracts regardless of their price relative to the break-even price of milk.

Although this thesis makes several important contributions, there are some avenues for extending the research questions considered in this thesis. The first question relates to the collapse of the XIV. Trading in the XIV was terminated in February 2018. It is interesting to see how it affected the information content of other inverse volatility ETFs. Three candidates, which have similar objective to XIV, could be considered: the ProShares Short VIX Short-Term Futures ETF (SVXY), the REX VolMAXX Short VIX Futures Strategy ETF (VMIN) and VelocityShares Daily Inverse VIX Medium-Term ETN (ZIV). In regards to dairy derivatives, the third empirical chapter concentrates on output price risk only, but an extension would be to incorporate input cost price risk. The NZ farming system is predominantly grass-based but still partially relies on purchased supplementary feeds, such as wheat, barley, and palm kernel. As of 2018, the NZX does not offer feed futures, so derivatives contracts from other stock exchanges could be considered.

In general, as the development of financial markets is not expected to slow down, it is important to evaluate new research questions about properties and uses of new derivative products and apply new methods and techniques to traditional questions. It can help not only point out potential risks, but can bring benefits to market participants. Benefits start with more informed speculating/hedging strategies and extend to the use of derivative products to infer useful information later applied to underlying assets.

Bibliography

Alexander, C. and Korovilas, D. (2012a). Diversification of equity with VIX futures: Personal views and skewness preference. Working paper.

Alexander, C. and Korovilas, D. (2012b). Understanding ETNs on VIX futures. Working paper.

Andersen, T. G. and Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4):885–905.

Bahaji, H. and Aberkane, S. (2016). How rational could VIX investing be? *Economic Modelling*, 58:556–568.

Baillie, R. T., Booth, G. G., Tse, Y., and Zabotina, T. (2002). Price discovery and common factor models. *Journal of Financial Markets*, 5(3):309–321.

Ballingall, J. and Pambudi, D. (2017). Dairy trade's economic contribution to New Zealand. Technical report, NZIER, `https://nzier.org.nz/publication/dairy-trades-economic-contribution-to-new-zealand`.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2009). Realized kernels in practice: Trades and quotes. *Econometrics Journal*, 12(3):1–32.

Barone-Adesi, G. and Whaley, R. E. (1987). Efficient analytic approximation of American option values. *The Journal of Finance*, 42(2):301–320.

Baur, D. G. and Dimpfl, T. (2018). The asymmetric return-volatility relationship of commodity prices. *Energy Economics*, 76:378–387.

Berkowitz, J. P. and DeLisle, R. J. (2018). Volatility as an asset class: Holding VIX in a portfolio. *The Journal of Alternative Investments*, 21(2):52–64.

Black, F. (1975). Fact and fantasy in the use of options. *Financial Analysts Journal*, 31(4):36–72.

Black, F. (1976). Studies of stock market changes. *Proceedings of the 1976 American Statistical Association, Business and Economical Statistics Section*, pages 177–181.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

Blair, B. J., Poon, S.-H., and Taylor, S. J. (2001). Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. *Journal of Econometrics*, 105(1):5–26.

Bollen, N. P., O'Neill, M. J., and Whaley, R. E. (2017). Tail wags dog: Intraday price discovery in VIX markets. *Journal of Futures Markets*, 37(5):431–451.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Bordonado, C., Molnár, P., and Samdal, S. R. (2016). VIX exchange traded products: Price discovery, hedging, and trading strategy. *Journal of Futures Markets*, 37(2):164–183.

Bosch, D. J. and Johnson, C. J. (1992). An evaluation of risk management strategies for dairy farms. *Journal of Agricultural and Applied Economics*, 24(2):173–182.

Brittain, L., Garcia, P., and Irwin, S. H. (2011). Live and feeder cattle options markets: returns, risk, and volatility forecasting. *Journal of Agricultural and Resource Economics*, 36(1):28–47.

Buffett, W. (2002). Berkshire Hathaway Inc. 2002 annual report. Technical report.

Campbell, J. Y. and Hentschel, L. (1992). No news is good news: An asymmetric model of changing volatility in stock returns. *Journal of Financial Economics*, 31(3):281–318.

Carr, P. and Wu, L. (2006). A tale of two indices. *Journal of Derivatives*, 13(3):13–29.

Chakravarty, S., Gulen, H., and Mayhew, S. (2004). Informed trading in stock and option markets. *Journal of Finance*, 59(3):1235–1257.

Chen, Y.-L. and Tsai, W.-C. (2017). Determinants of price discovery in the VIX futures market. *Journal of Empirical Finance*, 43:59–73.

Chow, K. V. and Denning, K. C. (1993). A simple multiple variance ratio test. *Journal of Econometrics*, 58(3):385–401.

Christie, A. A. (1982). The stochastic behavior of common stock variances: Value, leverage and interest rate effects. *Journal of Financial Economics*, 10(4):407–432.

Clark, T. E. and West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1):291–311.

Collins, R. A. (1997). Toward a positive economic theory of hedging. *American Journal of Agricultural Economics*, 79(2):488–499.

Conlon, T. and Cotter, J. (2013). Downside risk and the energy hedger's horizon. *Energy Economics*, 36:371–379.

Cremers, M. and Weinbaum, D. (2010). Deviations from put-call parity and stock return predictability. *Journal of Financial and Quantitative Analysis*, 45(2):335–367.

Cross, F. (1973). The behavior of stock prices on Fridays and Mondays. *Financial Analysts Journal*, 29(6):67–69.

Daigler, R. T. and Rossi, L. (2006). A portfolio of stocks and volatility. *Journal of Investing*, 15(2):99–106.

DairyNZ (2017). Dairynz economic survey 2015-16. Technical report, DairyNZ, https://www.dairynz.co.nz/media/5787208/dairynz-economic-survey-2015-16.pdf.

Dash, S. and Moran, M. T. (2005). VIX as a companion for hedge fund portfolios. *Journal of Alternative Investments*, 8(3):75.

DeLisle, J., Doran, J., and Krieger, K. (2014). Volatility as an asset class: Holding VIX in a portfolio. Working paper.

Deng, G., McCann, C. J., and Wang, O. (2012). Are VIX futures ETPs effective hedges? *Journal of Index Investing*, 3(3):35–48.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263.

Egelkraut, T. M. and Garcia, P. (2006). Intermediate volatility forecasts using implied forward volatility: The performance of selected agricultural commodity options. *Journal of Agricultural and Resource Economics*, 31(3):508–528.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the Econometric Society*, 50(4):987–1007.

Engle, R. F. and Granger, C. W. (1987). Co-integration and error correction: Representation, estimation, and testing. *Econometrica: Journal of the Econometric Society*, 55(2):251–276.

Fleming, J. (1998). The quality of market volatility forecasts implied by S&P 100 index option prices. *Journal of Empirical Finance*, 5(4):317–345.

Fleming, J., Ostdiek, B., and Whaley, R. E. (1995). Predicting stock market volatility: A new measure. *Journal of Futures Markets*, 15(3):265–302.

Fleming, J., Ostdiek, B., and Whaley, R. E. (1996). Trading costs and the relative rates of price discovery in stock, futures, and option markets. *Journal of Futures Markets*, 16(4):353–387.

Fleming, J. and Whaley, R. E. (1994). The value of wildcard options. *The Journal of Finance*, 49(1):215–236.

Fonterra (2017). Farmgate milk price statement. Technical report, Fonterra Cooperative Group Limited, https://www.nzx.com/files/attachments/266472.pdf.

French, K. R., Schwert, G. W., and Stambaugh, R. F. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, 19(1):3–29.

Frijns, B., Tallau, C., and Tourani-Rad, A. (2010). The information content of implied volatility: evidence from Australia. *Journal of Futures Markets*, 30(2):134–155.

Frijns, B., Tourani-Rad, A., and Webb, R. I. (2016). On the intraday relation between the VIX and its futures. *Journal of Futures Markets*, 36(9):870–886.

Giot, P. (2003). The information content of implied volatility in agricultural commodity markets. *Journal of Futures Markets*, 23(5):441–454.

Gonzalez-Perez, M. T. (2015). Model-free volatility indexes in the financial literature: A review. *International Review of Economics & Finance*, 40:141–159.

Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *Journal of Finance*, 50(4):1175–1199.

Hatchett, R. B., Brorsen, B. W., and Anderson, K. B. (2010). Optimal length of moving average to forecast futures basis. *Journal of Agricultural and Resource Economics*, 35(1):18–33.

Hibbert, A. M., Daigler, R. T., and Dupoyet, B. (2008). A behavioral explanation for the negative asymmetric return–volatility relation. *Journal of Banking & Finance*, 32(10):2254–2266.

Hudson, M. A., Leuthold, R. M., and Sarassoro, G. F. (1987). Commodity futures price changes: Recent evidence for wheat, soybeans and live cattle. *Journal of Futures Markets*, 7(3):287–301.

Hull, J. C. (2008). *Options, Futures, and Other Derivatives*. Upper Saddle River, N.J.: Pearson/Prentice Hall.

Husson, T. and McCann, C. J. (2011). The VXX ETN and volatility exposure. *PIABA Bar Journal*, 18(2).

Jiang, G. J. and Tian, Y. S. (2005). The model-free implied volatility and its information content. *The Review of Financial Studies*, 18(4):1305–1342.

Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica: Journal of the Econometric Society*, 59(6):1551–1580.

Johnson, L. L. (1960). The theory of hedging and speculation in commodity futures. *The Review of Economic Studies*, 27(3):139–151.

Jorion, P. (1995). Predicting volatility in the foreign exchange market. *The Journal of Finance*, 50(2):507–528.

Kenyon, D. and Clay, J. (1987). Analysis of profit margin hedging strategies for hog producers. *Journal of Futures Markets*, 7(2):183–202.

Khalifa, A. A., Miao, H., and Ramchander, S. (2011). Return distributions and volatility forecasting in metal futures markets: Evidence from gold, silver, and copper. *Journal of Futures Markets*, 31(1):55–80.

Kim, H. S., Brorsen, B. W., and Anderson, K. B. (2010). Profit margin hedging. *American Journal of Agricultural Economics*, 92(3):638–653.

Kim, M., Garcia, P., and Leuthold, R. M. (2009). Managing price risks using and local polynomial kernel forecasts. *Applied Economics*, 41(23):3015–3026.

Koeman, J. and Białkowski, J. (2015). Efficiency of hedging against fluctuating prices of dairy products. *Applied Finance Letters*, 4(1 & 2):6–11.

Kroner, K. F., Kneafsey, K. P., and Claessens, S. (1995). Forecasting volatility in commodity markets. *Journal of Forecasting*, 14(2):77–95.

Lee, C. I., Gleason, K. C., and Mathur, I. (2000). Efficiency tests in the French derivatives market. *Journal of Banking & Finance*, 24(5):787–807.

Lien, D. and Shrestha, K. (2014). Price discovery in interrelated markets. *Journal of Futures Markets*, 34(3):203–219.

Lo, A. W. and MacKinlay, A. C. (1988). Stock market prices do not follow random walks: Evidence from a simple specification test. *The Review of Financial Studies*, 1(1):41–66.

Loughrey, J., Thorne, F., Donnellan, T., Hennessy, T., and O'Connor, D. (2018). An evaluation of suitable tools to manage price/income volatility at dairy farm level in ireland. Technical report, Dairy Research Ireland, `https://www.teagasc.ie/publications/2018/an-evaluation-of-suitable-tools-to-manage-priceincome-volatility-at-dairy-farm-level-in-ireland.php`.

Manfredo, M. R. and Richards, T. J. (2007). Cooperative risk management, rationale, and effectiveness: The case of dairy cooperatives. *Agricultural Finance Review*, 67(2):311–339.

Manfredo, M. R. and Sanders, D. R. (2004). The forecasting performance of implied volatility from live cattle options contracts: Implications for agribusiness risk management. *Agribusiness: An International Journal*, 20(2):217–230.

Martin, L. J. and Hope, D. (1984). Risk and returns from alternative marketing strategies for corn producers. *Journal of Futures Markets*, 4(4):513–530.

Maynard, L. J., Wolf, C., and Gearhardt, M. (2005). Can futures and options markets hold the milk price safety net? Policy conflicts and market failures in dairy hedging. *Review of Agricultural Economics*, 27(2):273–286.

Neyhard, J., Tauer, L., and Gloy, B. (2013). Analysis of price risk management strategies in dairy farming using whole-farm simulations. *Journal of Agricultural and Applied Economics*, 45(2):313–327.

Padungsaksawasdi, C. and Daigler, R. T. (2014). The return-implied volatility relation for commodity ETFs. *Journal of Futures Markets*, 34(3):261–281.

Pan, J. and Poteshman, A. M. (2006). The information in option volume for future stock prices. *The Review of Financial Studies*, 19(3):871–908.

Peterson, P. E. and Leuthold, R. M. (1987). A portfolio approach to optimal hedging for a commercial cattle feedlot. *Journal of Futures Markets*, 7(4):443–457.

Poon, S.-H. and Granger, C. W. (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature*, 41(2):478–539.

Rapach, D. E., Strauss, J. K., and Zhou, G. (2010). Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. *The Review of Financial Studies*, 23(2):821–862.

RBNZ (2018). Financial stability report. Technical report, Reserve Bank of New Zealand, `https://www.rbnz.govt.nz/financial-stability/financial-stability-report`.

Schlusche, B. (2009). Price formation in spot and futures markets: exchange traded funds vs. index futures. *Journal of Derivatives*, 17(2):26–40.

Schroeder, T. C. and Hayenga, M. L. (1988). Comparison of selective hedging and options strategies in cattle feedlot risk management. *Journal of Futures Markets*, 8(2):141–156.

Schwarz, T. V. and Szakmary, A. C. (1994). Price discovery in petroleum markets: Arbitrage, cointegration, and the time interval of analysis. *Journal of Futures Markets*, 14(2):147–167.

Shadbolt, N. M. and Apparao, D. (2016). Factors influencing the dairy trade from New Zealand. *International Food and Agribusiness Management Review*, 19(B):241–255.

Shapiro, B. and Brorsen, B. W. (1988). Factors affecting farmers' hedging decisions. *North Central Journal of Agricultural Economics*, 10(2):145–153.

Shu, J. and Zhang, J. E. (2012). Causality in the VIX futures market. *Journal of Futures Markets*, 32(1):24–46.

Signori, O., Briere, M., and Burgues, A. (2010). Volatility exposure for strategic asset allocation. *Journal of Portfolio Management*, 36(3):105–116.

Smith, G. and Rogers, G. (2006). Variance ratio tests of the random walk hypothesis for South African stock futures. *South African Journal of Economics*, 74(3):410–421.

Stein, J. L. (1961). The simultaneous determination of spot and futures prices. *The American Economic Review*, 51(5):1012–1025.

Szakmary, A., Ors, E., Kim, J. K., and Davidson, W. N. (2003). The predictive power of implied volatility: Evidence from 35 futures markets. *Journal of Banking & Finance*, 27(11):2151–2175.

Triantafyllou, A., Dotsis, G., and Sarris, A. H. (2015). Volatility forecasting and time-varying variance risk premiums in grains commodity markets. *Journal of Agricultural Economics*, 66(2):329–357.

Turvey, C. G. and Baker, T. G. (1989). Optimal hedging under alternative capital structures and risk aversion. *Canadian Journal of Agricultural Economics/Revue Canadienne D'Agroeconomie*, 37(1):135–143.

Wang, Z., Fausti, S. W., and Qasmi, B. A. (2012). Variance risk premiums and predictive power of alternative forward variances in the corn market. *Journal of Futures Markets*, 32(6):587–608.

Weber, E. J. (2009). A short history of derivative security markets. In *Vinzenz Bronzin's Option Pricing Models*, pages 431–466. Springer, Berlin, Heidelberg.

West, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica*, 64(5):1067–1084.

Whaley, R. E. (1993). Derivatives on market volatility: Hedging tools long overdue. *The Journal of Derivatives*, 1(1):71–84.

Whaley, R. E. (2006). *Derivatives: markets, valuation, and risk management*. Hoboken, New Jersey: John Wiley & Sons, first edition.

Whaley, R. E. (2009). Understanding VIX. *Journal of Portfolio Management*, 26(3):12–17.

Whaley, R. E. (2013). Trading volatility: At what cost? *Journal of Portfolio Management*, 40(1):95–108.

White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica: Journal of the Econometric Society*, 50(1):1–25.

Wolf, C. A. and Widmar, N. J. O. (2014). Adoption of milk and feed forward pricing methods by dairy farmers. *Journal of Agricultural and Applied Economics*, 46(4):527.

Wood, W. C., Shafer, C. E., and Anderson, C. G. (1989). Frequency and duration of profitable hedging margins for texas cotton producers, 1980–1986. *Journal of Futures Markets*, 9(6):519–528.

Yang, S.-R. and Brorsen, B. W. (1994). Daily futures price changes and non-linear dynamics. *Structural Change and Economic Dynamics*, 5(1):111–132.

Yoon, B.-S. and Brorsen, B. W. (2005). Can multiyear rollover hedging increase mean returns? *Journal of Agricultural and Applied Economics*, 37(1):65–78.