Norwegian School of Economics
Bergen, Fall 2024

Hedging weather risk in agriculture

An empirical field study on the use of rainfall derivatives to hedge against crop yield losses due to adverse weather conditions

Jakob Grøner & Mathias Unhjem Øren

Supervisor: Svein-Arne Persson

Master thesis, Economics and Business Administration,

Major: Economic Analysis (ECO)

NORWEGIAN SCHOOL OF ECONOMICS

Preface

This master's thesis is part of our master's degree in Economics and Business Administration, specializing in Economic Analysis, at the Norwegian School of Economics (NHH) in Bergen.

The process of working on this thesis has been both demanding and rewarding. During our studies we have both found derivatives to be an exciting topic, and for our thesis we wanted to take a closer look at weather derivatives, which is relatively new but has significant potential in several industries. We came across studies on wind and hydropower producers, but we wanted to explore an industry with a slightly more complex relationship with weather. Despite difficulties encountered along the way, this journey has provided us with valuable knowledge and insights, making it an intellectually enriching experience.

We would like to express our gratitude to our supervisor, Svein-Arne Persson of the Department of Finance at Norwegian School of Economics, for his support and guidance throughout the process. His expertise and encouragement have been significant in shaping our work. Additionally, we would like to thank Mari Vengnes from Landbruksdirektoratet, for providing delivery data on grain crops between 2005 and 2023.

Abstract

Hedging revenues has been practiced for decades. Through the use of futures, forwards, swaps and options, companies have been able to hedge their associated risks. Historically, the process of hedging has mainly been related to price, while volumetric risks have not been considered to an equal extent. As the prices in Norwegian agriculture are set at a national level annually, the main source of risk relates to volume. The purpose of this thesis is to analyse the effect of hedging volumetric risk on the income of Norwegian farmers, by using weather derivatives. The effects are measured in both change of volatility and change in a simplified Sharpe ratio.

Weather derivatives are relatively underexplored compared to more traditional financial instruments and provide a new way of managing risks in industries that are strongly dependent on weather patterns. Agriculture relies heavily on weather conditions, and temperature, precipitation and sunlight all directly influence crop yields. This dependency makes the sector vulnerable to weather-related risks. Weather derivatives provide solutions to hedging against volumetric risks, such as fluctuating crop yields due to unpredictable weather, offering farmers a way to mitigate some of these uncertainties. After consulting with an expert in the field, we received clear indications that, in Norway, the primary weather risk is related to precipitation. While temperature and sunlight hours undeniably affect crops, their primary effect in the Norwegian climate is limited to influencing the time required for the grain to be ready for harvest. With this knowledge in mind, the thesis aims to construct precipitation derivatives that can help farmers stabilize their income. Agriculture is generally characterized by thin margins, and large fluctuations may result in substantial financial concerns.

As weather derivatives do not have a tradeable asset as the underlying, traditional methods to price financial derivatives, such as the Black-Scholes model, are not applicable. The majority of weather derivatives are traded over-the-counter (OTC) and consequently there is little data available, especially in Norway, on how to price them. This thesis will use some of the most common pricing methods that occur in previous literature – the historical burn analysis and Monte Carlo simulation. Through regression analysis, we aim to quantify the effect of precipitation on crop yields and use this as a basis for designing the derivative contracts. The findings of the analysis show that applying the proposed hedging strategy on an out-of-sample dataset through backtesting yielded an improvement in the majority of organization numbers.

Table of contents

List of tables

List of figures

Introduction

Agriculture is not only a central component to Norway's economy, it also provides jobs for nearly 40 000 people and plays a vital role in securing a certain degree of food security for the Norwegian population (Landbruksdirektoratet, 2024). Grains such as wheat, barley and oats are staples in the diets of millions of people, and the cultivating of these grains provide the primary source of income for thousands of farmers. By heavily subsidising the agriculture sector, the Norwegian government clearly indicates that it is desirable to maintain production, but it also faces major challenges with regards to climatic conditions. Norway's climate varies significantly across different regions and throughout the seasons, posing high risk to the crop yields. In fact, years with disadvantageous weather conditions can lead to a significant decrease in the amount of grain produced, and between 2022 and 2023 it was as much as 45% (SSB, 2024). Large fluctuations in the size of grain yields pose a significant threat for farmers and can make it very challenging to make ends meet.

Precipitation patterns greatly impact the quality and quantity of the crops. Periods of excessive rainfall or floods can lead to waterlogging, oxygen deprivation and an increase in the occurrence of diseases (Langan, et al., 2022). Similarly, periods of little rain or droughts leads to crops drying out and thus diminishing crop yields. In Norway, where the growth period for grain is mainly during the summer months, the weather becomes even more impactful as there is only one harvest per year for most grain types, as opposed to countries with more stable year-round climate. Volatile weather patterns can lead to significant volatility in crop yields, and farmers in countries with unpredictable weather thus face higher uncertainty in their incomes. Weather related risks represent one of the most significant challenges to grain farming, and understanding how to mitigate this risk is vital for maintaining more stable incomes. Both flooding and drought pose a threat to the Norwegian agricultural sector, and to address these challenges, the government has allocated increased subsidies for improving the drainage capabilities of farmers' land. Efficient draining systems can significantly reduce the losses in cases of excess rain, thereby stabilizing yields and making agriculture more resistant to extreme weather events. Drought on the other hand, is more difficult to handle. Irrigation systems represent big expenditures, both in terms of investment costs and operating costs, typically making them unprofitable.

Norwegian farmers do however not take on the full risk of weather conditions all by themselves. The government offers crop loss compensation programs, designed to support farmers financially during periods of severe weather conditions that ruin grain crops. Although this reduces risk for Norwegian farmers, there are indications that these reimbursement programs are not sufficient on their own. In fact, many farmers buy additional insurance to hedge themselves against poor weather conditions. In 2018, around every 4th farmer had purchased crop insurance on their own, despite the government providing a "safety net" (Finborud Børresen, Wold Haagensen, & Meskau, 2018). The compensation program states that "the yield must be less than 70 percent of the average yield for the crop group to trigger the basis for subsidies. This means that there is a 30 percent individual risk involved" (Landbruksdirektoratet, 2024). In the event of farmers facing liquidity problems subject to crop losses, it is also possible to apply for an advance payment of the compensation program. The decision regarding advance payment is made by the county governor and requires farmers to provide sufficient information and documentation (Landbruksdirektoratet, 2024). Should this application not be accepted, farmers might also face the issue of liquidity risk, as grain farming in general is characterized by thin margins and low profitability.

A possible and more modern way of reducing risk is the use of weather derivatives. Such derivatives can be based on parameters like rainfall or temperature, and thus offer farmers a way of hedging against adverse weather conditions. While the government's crop loss compensation programs and normal insurance only cover more extreme cases, weather derivatives can also help protect against losses from less severe, yet still impactful weather conditions. Weather derivatives can be used to resolve smaller, more frequent, events that would not trigger an insurance payout or government crop loss subsidies. For example, they could be structured to provide payouts in periods of prolonged dry periods that do not result in complete failure, but still reduce yields, subsequently leading to financial uncertainty. The combination of the government subsidies and use of weather derivatives can thus complement each other to significantly reduce the weather risk that Norwegian farmers face in their grain production.

Parts of our thesis are inspired by "The Use of Wind Speed Derivatives for a Norwegian Energy Producer" by Røkkum and Myrvang (2019). More specifically, the practical use of weather derivatives to manage weather risk for a real-world business has served as an inspiration. Our thesis does however differ in both industries and methodologies. Firstly, agriculture has a significantly more complex relationship with weather patterns compared to a wind energy producer. This additional complexity leads to additional challenges in our analysis, as agricultural production is influenced by numerous weather elements, but also more industry-specific factors such as crop rotation, growth stages and different sensitivities of different crop types. Consequently, our analysis requires more nuanced perspectives and is harder to interpret due to interrelated variables. Some of our pricing methods also differ, as we aim to make as few restrictive assumptions as possible. After finding coefficients from regression analysis, we add a step of optimizing the tick sizes, and thereby the corresponding prices. Lastly, we also perform a backtest on a different dataset to validate the effectiveness of the proposed hedging strategy.

Our contributions to the field include several aspects that have not been widely explored in the literature we have reviewed. This includes the method of price-loading, to account for the risk that the counterparty faces. Additionally, we model actual rainfall data to estimate physical probabilities which is then used in a Monte Carlo simulation to price the derivatives. Whereas most rainfall derivatives are based on accumulated rainfall over a certain period, we develop monthly option contracts based on rainless days.

Weather derivatives

Background

Weather derivatives are financial contracts that are dependent, much like commodity or stock derivatives, on some underlying factor that determines the payoff of the contract. The underlying factors of weather derivative contracts can be temperature, rainfall, wind, hours of sun and other weather conditions. The majority of weather derivative contracts are traded with temperature as the underlying weather index (Castano, 2024).

Weather derivatives were first developed in the United States during the 1990s. The product was initially invented as a method for utility companies to manage the risk associated with fluctuating temperatures. The first contracts taken into use were HDDs (heating degree days) and CDDs (cooling degree days) which today remain the most common type of weather derivative. These contracts provide payout based on how much the average daily temperature deviates from a benchmark, often set at 65 degrees Fahrenheit. The idea behind these derivatives was that once the temperature deviated from the benchmark, households tended to require either heating or cooling systems. The product thus served as a great way for utility companies to hedge the temperature risk that was inherent in their business (Castano, 2024).

Initially, weather derivatives were traded over-the-counter (OTC) and to this day it is estimated that 90% of weather derivative contracts are traded OTC (Robertson, 2023). However, there are also exchanges that offer weather products, the most prominent being the Chicago Mercantile Exchange (CME). They offered the first exchange traded temperature-based weather future in 1999 (Niedens, 2023). These contracts were based on a few American cities, since then the exchange has increased their offerings to include 13 American cities, four European and one in Asia. Even though the market did not experience rapid growth in the following decades, there has been a steady increase in contracts traded and in recent years there has been a significant rise in the market. According to the CME Group they saw an increase of 260% in the traded amount in 2023 compared to 2022 (Castano, 2024). This recent development in the weather products market could be due to more concern about climate change and market participants seeing more clearly how different industries are affected by weather.

Literature

Compared to stock or commodity derivatives, weather derivatives are less common in today's financial markets. However, their popularity has been growing in recent years as interest in alternative risk management solutions for climate-related challenges increases. This growing interest is accompanied by increasing support for the idea that weather derivatives can help mitigate risk in the agriculture sector. As concerns over climate change continue to grow, weather derivatives are considered as an alternative way for farmers to better control the unpredictability of weather.

In the research article "Effectiveness of weather derivatives as a hedge against weather risk in agriculture" from 2015, it is demonstrated that the application and effectiveness of such derivatives has been studied and proven in production of an extensive range of crops, such as wheat, barley, soybean, grapes and cotton (Štulec, Petljak, & Baković, 2015). The variance and standard deviation are considered the most common measure of volatility, while value-at-risk, certainty equivalent revenue, mean-variance criterion and utility function enhancement are used to a lesser extent. In other words, there are no universally accepted measure of the effectiveness of weather derivatives. Existing studies show that the effectiveness varies between geographical areas, crops and the time periods covered, and the volatility reduction ranges from 10.8% to 77.1% (Štulec, Petljak, & Baković, 2015).

The study refers to multiple other studies, of which the highest risk reducing performance was for corn, soybean and cotton production in two regions in USA. This study from 2004 utilized put options on both temperature and rainfall indices, and the risk reducing performance ranged from 16.6% to 77.1%. The authors highlighted that the weather derivatives design should be customized for each crop and each geographical area, which is something we also incorporate into our analysis. Marković and Jovanović (2011) studied the success of rain put options in the winter barley production in Germany. A 40.42% reduction in standard deviation was accomplished by utilizing put options.

Martina Bobriková (2022) highlights that the sector "is highly sensitive to meteorological elements that affect the yield of many crops", and points out temperature, rain, wind, snow, humidity and sunlight as important factors (Bobriková, 2022). By using weather options for the east of Slovakia, the study examines and evaluates hedging strategies for reducing the risk of excess rainfall. Whereas our analysis focuses on shortfall of rain, the study by Bobriková focuses on the opposite case. By using correlation analysis, rainfall is found to be the weather element with highest correlation with yields, which was in line with our expectations. The underlying of the weather derivatives is accumulated rainfall, and strike prices are based on historical data between 1980 and 2020.

The call and put options are priced using historical burn analysis which utilizes historical data to find hypothetical payouts for past scenarios, which is then used for pricing. The prices do however not take risk premiums into account. Nine different strategies are tested, including long calls, long straddles and long strangles with different strike values. By comparing the volatility of hedged yields with unhedged yields, in an ex-post analysis, the study finds that producers were able to significantly decrease their yield variation. The most effective strategy was found to be a long straddle strategy, reducing the variation coefficient by 15.66%. All nine strategies resulted in a decrease in the variation coefficient, and further emphasizes that weather derivatives can be an effective tool in the agriculture sector. While the primary focus of the study was to hedge against excess rainfall, the optimal strategy proved it was beneficial to hedge in both directions. We acknowledge that a similar strategy may have been beneficial in our case as well, especially in the past, but as the Norwegian government has increased subsidies for improving the draining abilities, we considered drought to be the largest threat to Norwegian crops in the future. The study does however not distinguish between different types of crops and areas, the strategies are not backtested, and also fails to add a risk premium, leading to overly positive results. Incorporating a risk premium into the option prices is essential, as it represents the market cost of acquiring the hedge. Without the inclusion of this added premium, which compensates the counterparty for the risks undertaken, there would not be sufficient incentive for them to participate in the transaction. Excluding a risk premium will lead to low prices that are not reflective of real-world conditions and could hence lead to unrealistically positive results.

Properties of the contract

Weather derivatives are typically structured as futures, options or swaps. In the analysis, we construct precipitation call options, and these contracts must specify the following information:

- Type of contract (call or put)

- Duration of the contract
- Underlying weather index
- Official station for where the index is measured
- Strike level
- Tick size

Options offer the possibility of large payouts, and the loss is limited to the initial option premium paid. A typical call option can be written on the following form:

$$Payout = Tick\ size * \max\ (Index - Strike, 0)$$

And the profit of the option may be defined as:

$$Profit = Payout - Compunded\ option\ premium$$

As the payout and the option premium occur at different points in time, any analysis should account for the timing discrepancy by adjusting one of the values so that both are expressed in terms of the same point in time. Once a contract has been agreed upon between two parties, the buyer of the option is obligated to pay an option premium to the seller. The option premium will vary with the strike level, the tick size, and the assumed probability distribution of the underlying index. If the measured index exceeds the strike level at maturity, the seller is obligated to pay the difference times the tick size to the buyer. The buyer of the option only makes a profit if the payout exceeds the compounded premium. Although a weather and a commodity option are similar in many aspects, there is one major difference. While a commodity option is based on an underlying tradeable asset like oil, the underlying of a weather option is an index which cannot be traded. While standard option pricing theory is largely based on creating a replicating portfolio to the option, this is not possible for weather derivatives since the underlying is not traded. This has major implications for the pricing of weather options which will be addressed later.

Advantages and challenges with weather derivatives

The main advantage of weather derivative contracts is to hedge weather risk. They are one of the few traded products that give businesses a way of hedging this specific risk. While other products might have correlation with certain weather events, like orange juice futures (Roll, 1984), very few are written directly on weather indices. The main alternative for weather derivatives is insurance or catastrophe bonds (CAT bonds), which have gained significant interest in recent years (Cueto, 2024). Despite insurance and CAT bonds being reasonable alternatives, there are differences in what type of weather they hedge the user against. These types of products are typically best suited to hedge against extreme weather events like hurricanes, floods and blizzards, while weather options on the other hand are better suited to hedge against significant, but not extreme weather fluctuations. For instance, HDD and CDD contracts are designed to pay out when temperatures deviate from 65 degrees Fahrenheit, making them useful even in the absence of extreme temperatures. Additionally, while consumers of insurance need to prove their losses to receive the payout from the insurance company, weather contracts are written on an objective index, making the payout process simpler and more predictable.

Another advantage of using weather derivatives as opposed to insurance, is the absence of moral hazard and adverse selection issues. Moral hazard occurs when the insured party, after obtaining coverage, increases the risk of their operations (Nickolas, 2024). This is because more risk typically implies higher expected return. Although more risk also increases the chance of big losses, this is no longer as detrimental since the worst-case scenarios are covered by insurance. Adverse selection arises when the insured party has more information about their business than the insurer does (Nickolas, 2024). The insured party thus has a better foundation for evaluating the terms of the contract. Even though the problems associated with moral hazard and adverse selection appear to solely concern the insurance company, they can have an indirect effect for consumers of insurance as well. Since insurance companies are aware of these effects, they can increase the price of the insurance to compensate. Weather derivatives eliminate these problems as the payout is simply based on some objective index, and there is thus an argument to be made that weather derivatives could be more cost effective than insurance.

Although weather derivatives do offer certain advantages, they remain relatively niche and are a newer product compared to insurance. They can only be traded on an exchange for a few places around the world, meaning that for most market participants trading OTC is the only alternative. OTC trading is not as liquid as trading on an exchange which is possible for other derivatives. This leads to some liquidity risk when entering contracts. Basis risk is also a challenge with weather derivatives, which is the risk that arises from the fact that the weather station normally is located some distance away from the location that is being hedged. When the distance increases so does the basis risk, as the likeliness of deviations in the weather increases.

Weather derivatives today

Utility companies remain some of the largest participants in the weather derivatives market (Robertson, 2023). However, there have been several other industries entering the market since the 1990s. Hedge funds and traders are increasingly using weather derivatives in their portfolios in chase for increased profits (Robertson, 2023). There is huge potential for businesses that have large volumetric weather risk to reduce the risk of their operations. The agricultural industry has also started using weather derivatives as their operations are highly affected by weather. Amongst agricultural companies that have started using these contracts are GrainCorp, a major global agricultural business (Evans, 2019). As crop yields are largely affected by the weather conditions like rainfall and temperature, weather derivatives provide a great way of hedging this vulnerability.

The contracts we will propose in this thesis will be precipitation-based options as this is the weather condition that has the largest impact on grain yields in Norway. There is no exchange for trading weather derivatives in Norway. The contracts we will propose would thus be traded OTC and a counterpart must be found. This could for example be banks, insurance companies or investors.

Pricing Methods

Historical burn analysis

Historical burn analysis is a pricing method commonly used in the valuation of weather derivatives. The approach is relatively straightforward, as it relies exclusively on historical weather data rather than modelling or forecasting techniques. By reviewing historical data, this method assesses the potential payout distribution of the proposed weather derivative by simulating its performance over past years. Jewson & Brix (2005) describes it as the "idea of evaluating how a contract would have performed in previous years". Although other methods may be more accurate or provide more information in some cases, historical burn analysis is a good first step in pricing almost any contract and might serve as a benchmark for other pricing methods.

The method uses historical data, which in our case will be precipitation data, to estimate the probability of certain weather events occurring. To ensure statistical reliability it is ideal to use as many years of data as possible, however this is under the assumption that the data does not contain any statistically significant trend. If there is uncertainty about the shape of the trend, it is beneficial to use fewer years (Jewson & Brix, 2005, p. 54). For each year in the dataset, the performance of the weather derivative is calculated based on actual weather data in the corresponding year. More specifically, this involves an assessment of whether the recorded weather data met the "strike" conditions of the contract, and consequently calculating the hypothetical payout that would have been issued. To find the expected payout of the derivative, the average payout across all years serves as an estimate, which is then discounted to its present value to set the price of the derivative.

The biggest advantage of the historical burn analysis is its simplicity. Other than detrending any potential trends, the method avoids complex modelling and other assumptions, relying solely on historical data to generate potential payouts. It does not require assumptions on probability distributions or advanced statistical models. Whenever assumptions are made in modelling data, "we may add something, but no assumptions are ever exactly correct, and so we also introduce errors" (Jewson & Brix, 2005, p. 66). Historical burn analysis is also quick to implement relative to many other methods. It requires less extensive computational resources than for example Monte Carlo simulation. It can also be mentioned that historical burn analysis offers transparency in pricing, as each year`s hypothetical performance is directly observable and interpretable.

This approach also has some disadvantages. The reliance on historical data assumes that future weather patterns will closely resemble those of the past, which may no longer be realistic in the context of climate change or other weather trends. Depending on what type of underlying the contract is based upon, it is important to assess whether historical data may still be representative for future events. For example, it is well established that global temperatures are rising, and for historical burn analysis to be applicable in such a context, it would be a prerequisite to take trend into account to accurately price a temperature derivative. Shifts in weather patterns could lead to both under- and overestimation of weather-related risks when relying only on historical data. Another challenge is that since historical burn analysis requires a relatively long period of data, it could be challenging to apply in regions where weather data is incomplete or of poor quality. These deficiencies could skew results. In some cases, this approach might also not capture extreme events if they are rare and fail to appear in historical data. This could in turn lead to mispriced derivatives that do not fully capture the weather risk. Additionally, extreme events could also be overestimated in their likelihood if the small sample of data contains unproportionate amounts of extreme outcomes.

Monte Carlo simulation

Monte Carlo simulation is a common tool used in option pricing, and it can be used for pricing weather options despite them not having a tradeable underlying asset. The method relies on simulations based on probability distributions to calculate an expected return for a contract, which then is used to estimate a price for the option. More specifically, once the payout structure of the contract is determined, all parameters are known in advance except for the price/level of the asset/index at a point in the future. To estimate the unknown future value of the price/index, Monte Carlo simulation generates possible values from a probability distribution. Once a distribution has been selected, a large number of observations from this distribution is drawn to remove randomness and thus reduce the variance of the estimate. For each observation drawn, the payout of the contract is calculated and discounted, resulting in a vector that contains a large number of payouts. The payouts are discounted to account for the time value of money, given the period between purchasing the contract and receiving the potential payout. In standard option pricing theory, risk-adjusted probabilities are normally utilized. This approach has the advantage that one can consequently discount payouts using the risk-free rate, however certain adjustments are required when discounting weather options, which we will discuss in more detail. After obtaining the discounted payout vector, the average payout is estimated to find the price of the contract.

There are many advantages with using Monte Carlo simulation to price option contracts and other derivatives. One of them is that it provides a great way of dealing with uncertainty in the estimation. Unlike historical burn analysis which is based on a limited number of observations, Monte Carlo simulation has unlimited observations available, effectively reducing uncertainty in the estimation of the option price. Another advantage is the simplicity of the method, which makes it easy to employ. This is an important concept for models that often can get overlooked. Less complex methods are more comprehensible for users and thus make it easier to understand the limitations of the method. While more complex models sometimes can yield better precision in estimates, they can often be misunderstood, which can lead to dramatic consequences.

Monte Carlo simulation is also a very flexible method and works for several different types of derivatives, including weather derivatives. Other models, such as the widely known Black & Scholes method, are limited in the types of derivatives they can be applied to, as it requires the underlying asset to be tradable. Lastly, Monte Carlo simulation has the advantage that it can utilize historical data in its estimation process. This will be further demonstrated when later pricing the precipitation options, by using real data to approximate distributions.

Despite Monte Carlo simulation being an effective and robust framework, it also has some disadvantages. One example is that a distribution of the underlying index must be assumed. Should the distribution be a poor approximation of the actual underlying, the results will be unreliable. However, there are ways of evaluating the chosen distribution and how well it aligns with historical data, which minimizes the risk of selecting an incorrect distribution. Another issue with the methodology is that it requires computational power to run thousands of simulations. If multiple derivatives are priced simultaneously, the computational power required can become an issue in certain cases. Developments in modern technology have lessened this challenge significantly, as most standard computers nowadays possess sufficient computational power.

When pricing standard financial options with Monte Carlo simulation it is common to discount payouts with the risk-free rate. The reason is that the presence of a liquid tradeable should in theory allow for market participants to replicate the payout structure of the option by combining a position in the underlying and a riskless bond. Since one can recreate the option's payout by investing in other products, the price of the replicating portfolio and the option must be the same for there to be no arbitrage opportunities. Since the replicating portfolio can be discounted at the risk-free rate, so can the option itself (Kenton, 2024). However, the pricing of weather derivatives requires a different approach, as the underlying is not a traded asset and consequently a replicating portfolio does not exist. Therefore, one must account for the weather risk when computing the price of a weather option. This represents the risk the counterpart of the contract faces, since their position cannot be hedged in the market. This can be resolved by treating the weather option as a standard financial option and then adding a premium to the calculated price. The appropriate magnitude of this adjustment will be discussed later. Even though a premium is added to account for weather risk, payouts will still be discounted at the risk-free rate in both pricing methods, to adjust for the time value of money.

Data

In this section, the data used in the analysis will be presented. High quality data is essential when trying to quantify the impact of precipitation on crops, but also for use of weather derivatives pricing. When collecting data, common issues include both missing data and data that seem unreasonable. An essential step in any analysis is therefore to first clean the data.

Delivery data

Landbruksdirektoratet has provided delivery data (kg of grain per type per year) for all organization numbers that have delivered grain to grain reception facilities during the period from 2005 to July 2023. Organization numbers represent the individual level of farming operations, rather than companies. The data is limited to the counties of Akershus and Østfold from 2005 up until the formation of Viken in 2020, when Akershus, Buskerud and Østfold were merged into one county. There is also a second dataset, containing the delivery data of Viken between 2020 and 2023. The subsequent analysis will aim to quantify the impact of precipitation on crop yields, utilizing the first part of the delivery data, while the second part will be used for backtesting. Going forward, "delivery data" will refer specifically to the data from 2005 to 2019, unless stated otherwise.

The delivery data contains information on organization number, year, county number, municipality number, type of grain and quantity (in kg). The types of grain are wheat, rye, barley, oats, triticale, feed peas and oilseeds, field beans and lupins. Figure 1 visualizes the total delivered amount (in kg) by grain type in the period 2005-2019 in the two counties.

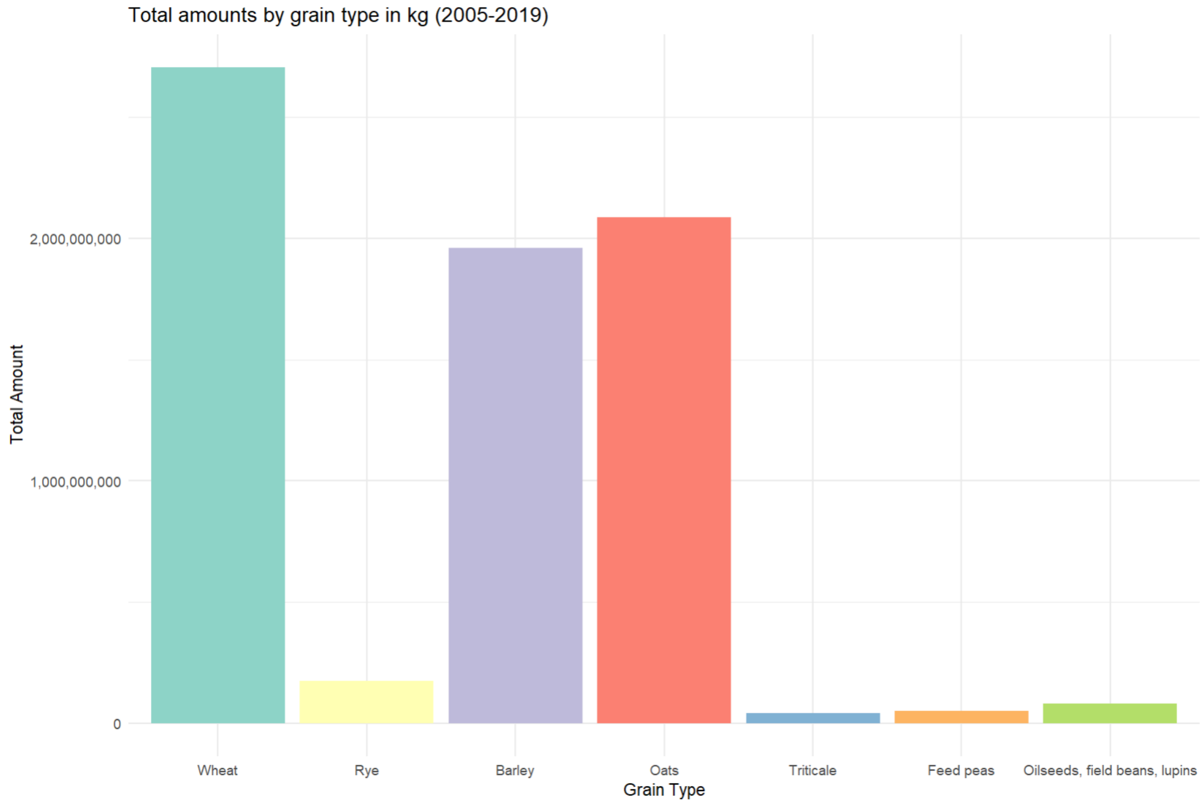

Figure 1: Grain quantities

The dataset contains of 118.597 observations, providing a substantial amount of data for the analysis. Each organization number usually occur many times in the dataset, as they typically produce various types of grain over many years, meaning the dataset consists of panel data. The original dataset did not contain information about the size of the areas related to the production, which posed a significant issue for the analysis. As one object of the analysis was to ultimately quantify the impact of precipitation on crop yields, areas would naturally play a significant part. The general idea is that a larger field would, on average, incur greater absolute losses or gains in quantity due to weather related incidents. To gather information on the size of areas, we were able to retrieve information about applications for production subsidies, by organization numbers. This subsidy is described as a collective term for several subsidy schemes that enterprises with livestock and/or crop production can apply for, and they constitute an important source of income for many farm owners (Landbruksdirektoratet, 2024)**.** These subsidy applications are made annually, based on total designated area per grain (in decares) hence providing the information we desire.

We acknowledge that acquiring area data through subsidy applications could already be a potential source of error in our analysis. In any case where an organization number fails to apply for this subsidy, the corresponding delivery data observation would be excluded from our analysis, as there would be missing data on area. These exclusions may result in loss of valuable information but are based on the understanding that the size of the area is a prerequisite for accurately evaluating each observation. An area of zero could also be subject to data retrieval issues or data omissions, and the corresponding delivery data would similarly be excluded. Another source of error arises when organization numbers incorrectly report the area in their applications. This can lead to inaccurate productivity figures, as underreporting the area will inflate productivity estimates, while overreporting will deflate them. Since no intuitive method for adjusting or modifying the area data itself could be identified, we have chosen to utilize the data as it is. We will, however, at a later stage, exclude certain observations based on their production per area, and hence address some of these uncertainties.

The delivery data contains seven different types of grain. These types all have different characteristics and responds differently to climate factors, such as precipitation, humidity and temperature. With this in mind, the idea was to limit the analysis to only one type of grain. After speaking with people involved in agriculture, we were left with the impression that wheat was the grain type most sensitive to weather conditions and therefore most suitable for the analysis. However, the structure of the delivery data made this choice more difficult. Wheat is a collective term for both spring wheat and winter wheat, and the delivery data did not separate between these two types. Spring wheat are sown and harvested in the same year, typically over a 4 month period, while winter wheat are sown in the fall and harvested in the fall of the following year. Thus, winter wheat remains in the fields throughout the winter. For food production, spring wheat is used primarily, as winter wheat is mostly used for animal feed because its quality often deteriorates (Langerud, 2023).

As the delivery data did not separate between the two types, and winter wheat seems to be significantly more dependent on climate factors over the winters we were certain that the analysis could not consider these two types as equivalent. When subsequently doing regression analysis, we suspected that this could lead to biased estimators. For this reason, we decided to focus elsewhere. In Norway, barley is grown on about half of the cereal area and has historically been the grain type with the highest production (Holtet, 2024). Additionally, our agricultural contact indicated that oats, which makes up a large share of Norway's yearly grain harvest, is not as sensitive to precipitation as barley. Studies also show that drought stress could reduce grain yield by 49-87% in barley (Samarah, 2005; Samarah et al., 2009). As barley is both the most widely cultivated grain in Norway and particularly sensitive to precipitation, we decided to focus the analysis on this grain type.

Limiting the dataset to delivery data on barley, we extracted the top 500 largest producers in Akershus and Østfold and filtered out only the observations that had corresponding data on area. This left us with 6283 observations. To address the issue of inaccurate productivity numbers, due to uncertainty around area data, we defined a new variable called PPA – production per area (kg per decare). In the period 2005-2019 the average kg per decare of barely was just shy of 400. In our dataset, there were some outliers, with the most extreme observation having a PPA of around 5000. As there were some uncertainty around both the delivery data but especially the area data, we wanted to filter out all observations with a PPA higher than 800, roughly double of the 15-year mean. This led to roughly 1.1% of observations being excluded.

As the goal of the analysis is to quantify a relationship between precipitation and crop yields and subsequently a hedging strategy, we also wanted to exclude outliers that were likely to have been caused by other factors than climate. Such factors may include weed, insects, viruses, bacteria and fungi. To do this, we used an approach that was based on filtering out observations where PPA was below the commonly recognized "worst year" in the time interval. 2018 was affected by an extreme drought that resulted in very poor grain yields and has since been referred to as a "disaster year". This situation highlighted the vulnerabilities in agriculture and the significant impact of climate change on farming practices. By filtering out all observations where an organization numbers PPA was worse than in 2018, we were able to exclude another 10% of the observations. We find it highly likely that these observations, were caused by factors that do not directly relate to precipitation. Ultimately, we were left with 5579 observations, distributed over 494 different organization numbers. These observations are illustrated in Figure 2. The following figure, Figure 3 displays how PPA changes for 20 random organization numbers over the time period.

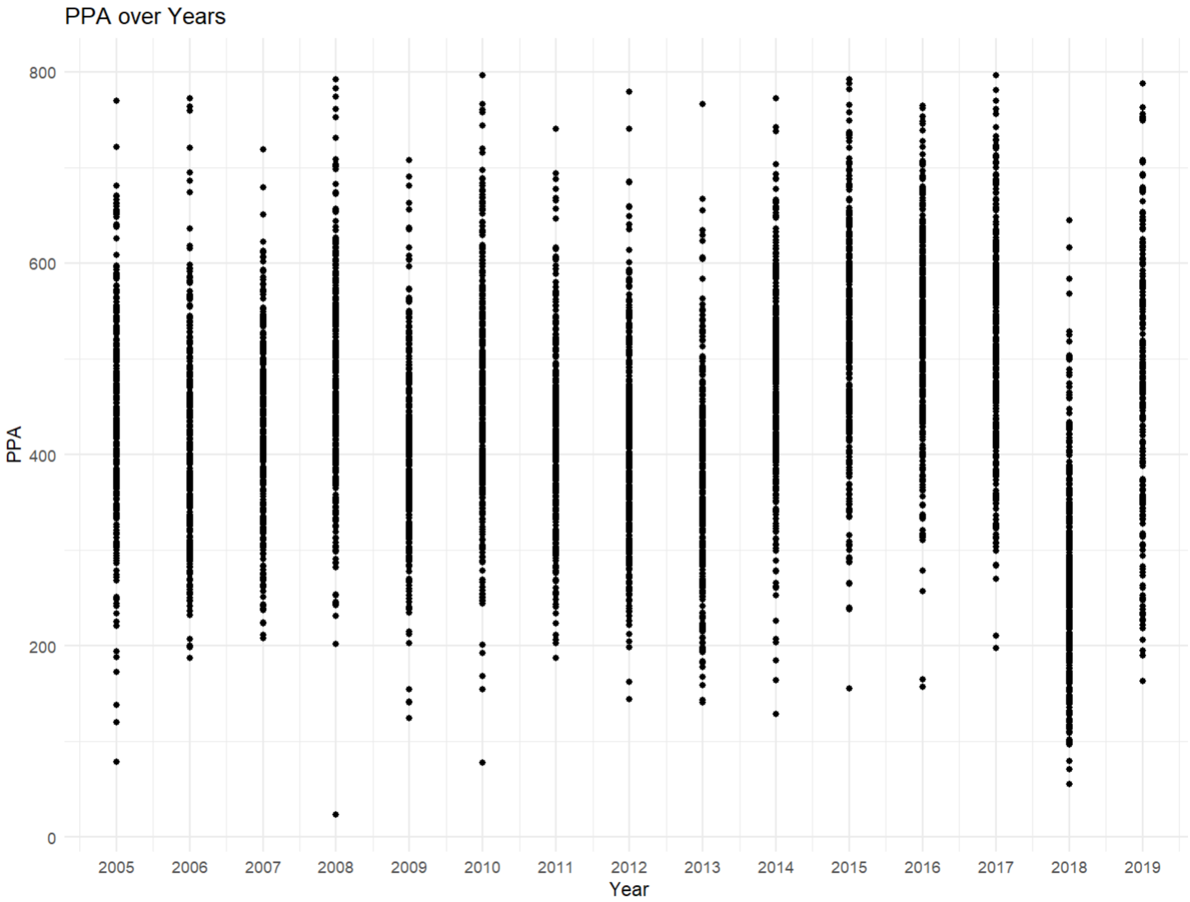

Figure 2:Yearly PPA values

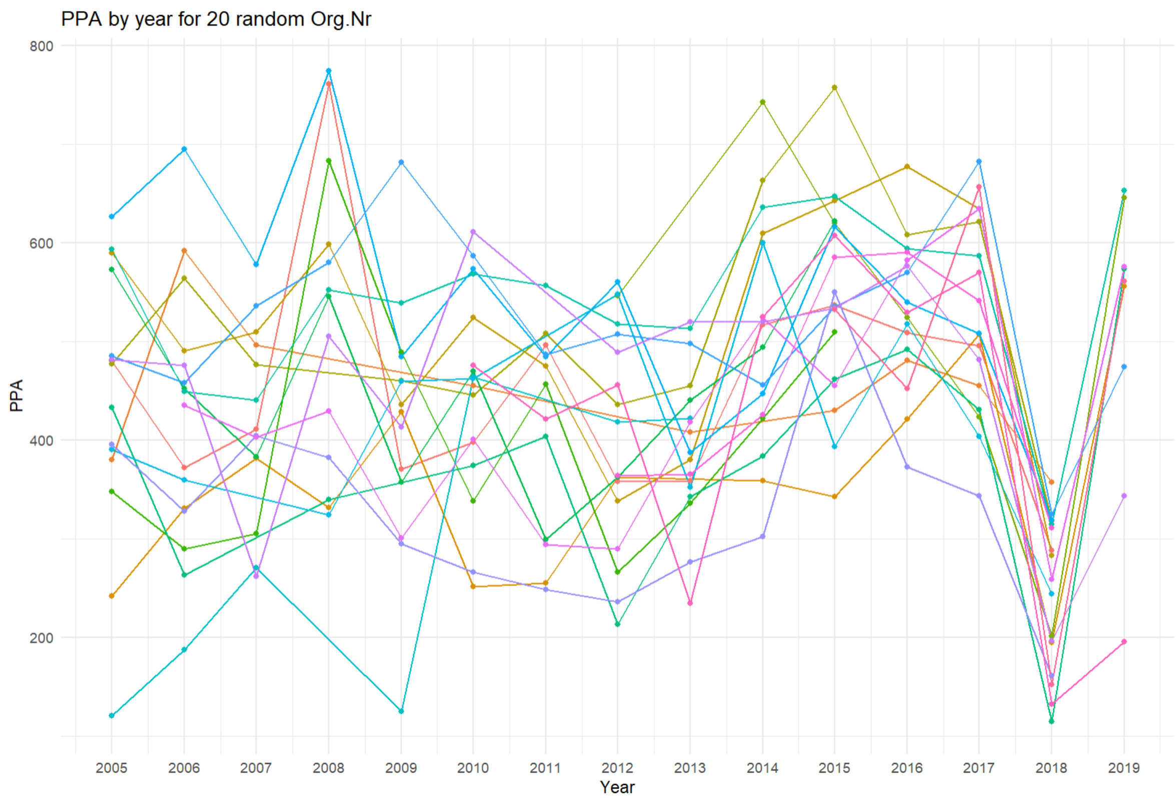

Figure 3: PPA for 20 random organization numbers

A central reason to why we limited the dataset to the 500 largest producers has to do with the concept of “crop rotation”, which is essential to organic grain farming. This concept refers to growing different crops in a set sequence, alternating between plant families over time. The goal is to reduce the risks of pests and diseases specific to certain plants. For example, plants with deeper roots bring nutrients from deeper soil layers, which subsequent crops benefit from. Rotation promotes soil biodiversity, contributing to better soil structure. Crop rotation also improves weed management by using different weed control strategies for different types of crops (Frøseth, 2011). Limiting the dataset to the 500 largest producers were made on the assumption that these farms are better at following this practice.

Weather data

In the following, the data on precipitation from the two counties, Østfold and Akershus, will be presented. The filtering and cleaning of the delivery data left us with 494 different organization numbers. When pricing the weather derivatives, it is essential to collect weather data from independent weather stations, ideally with very close proximity to the farm associated with the organization number. The idea of independent weather stations, in the form of a third party, is vital in pricing of weather derivatives. As farmers would be the buyers of the contracts, their own precipitation measurements would not be considered independent, and thus not reliable to any counterparty. As the payout depends on weather data, it would open up the possibility of tampering the measurements and hence committing fraud. For a counterparty to participate, it would require an independent third party to measure the precipitation, which would act as the underlying in the contracts. This leads to basis risk, which represents the risk due to the distance between the area that the buyer wants to cover and the actual location the contract will be written on. Although both counties have numerous independent weather stations, it would be very difficult to link up each farm to a weather station without knowing its exact location. It would also require a considerable time investment.

Data on daily, monthly and yearly precipitation in the two counties can be retrieved from Norsk Klimaservicesenter. A general issue with this data is that most weather stations have a lot of missing data both in the main period of interest (2005-2019), but also in the years prior. An extensive historical dataset greatly improves the accuracy of rainfall modelling when trying to identify appropriate distributions. With insufficient data on most weather stations, we chose one weather station for each county, based on both the number of observations available, but also a discretionary assessment of where the concentration of farms was highest in the two counties. As we only use precipitation data from two weather stations, it also means that there will only be two “sets” of precipitation data each year. This means that all organization numbers belonging to the same county also have the same precipitation data, when later merging the delivery data table and rainfall table by year. For Akershus, we chose to retrieve data from the weather station of Gardermoen, and for Østfold, we chose Halden, which is one of the oldest weather stations in the country. The former weather station has full data on daily precipitation all the way back to 1957, whereas the latter goes back to 1883. This provides a good dataset for modelling rainfall.

Using data from only one weather station in each county incurs a greater basis risk than individually assigning the closest weather stations. Counties are large, and two farms on the opposite ends of the periphery might experience different weather conditions. Assigning both of these farms to the same weather station, might provide substantial basis risk. Although using more than one weather station would be preferable, this approach would simply not be achievable without knowing exact locations and having access to more weather stations with complete weather data. As most weather stations had several days or periods of missing data, it also raises doubt about the accuracy of the data when it was available. Whenever possible, we tried to compare measurements of multiple weather stations within a county, and the overall agreement was that large differences were rare. Overall, the two selected weather stations have an extensive dataset of historical observations, and the accuracy is believed to be high.

Table 1 summarizes the average daily precipitation (in mm) for the two weather stations. In general, the weather station at Gardermoen generally record higher levels of precipitation. Both weather stations experience increasing precipitation over the five most relevant months.

	Average daily precipitation (in mm)	
	Østfold (Halden)	Akershus (Gardermoen)
April	43,27	48,63
May	49,83	62,66
June	65,99	74,82
July	77,11	84,39
August	88,71	97,55

Table 1: Average daily precipitation by county

Figure 4 and Figure 5 below shows the histograms of frequencies of accumulated precipitation by month, in Østfold (Halden) and Akershus (Gardermoen) respectively. It is worth noting that the former weather station depicts 100 years of data, while the latter only depicts 67 years. As indicated by Table 1, the accumulated precipitation levels generally seem to increase by month. The magnitude of the most extreme observations also seems to increase by month. In both counties there are observations of more than 200mm of precipitation in the later summer months.

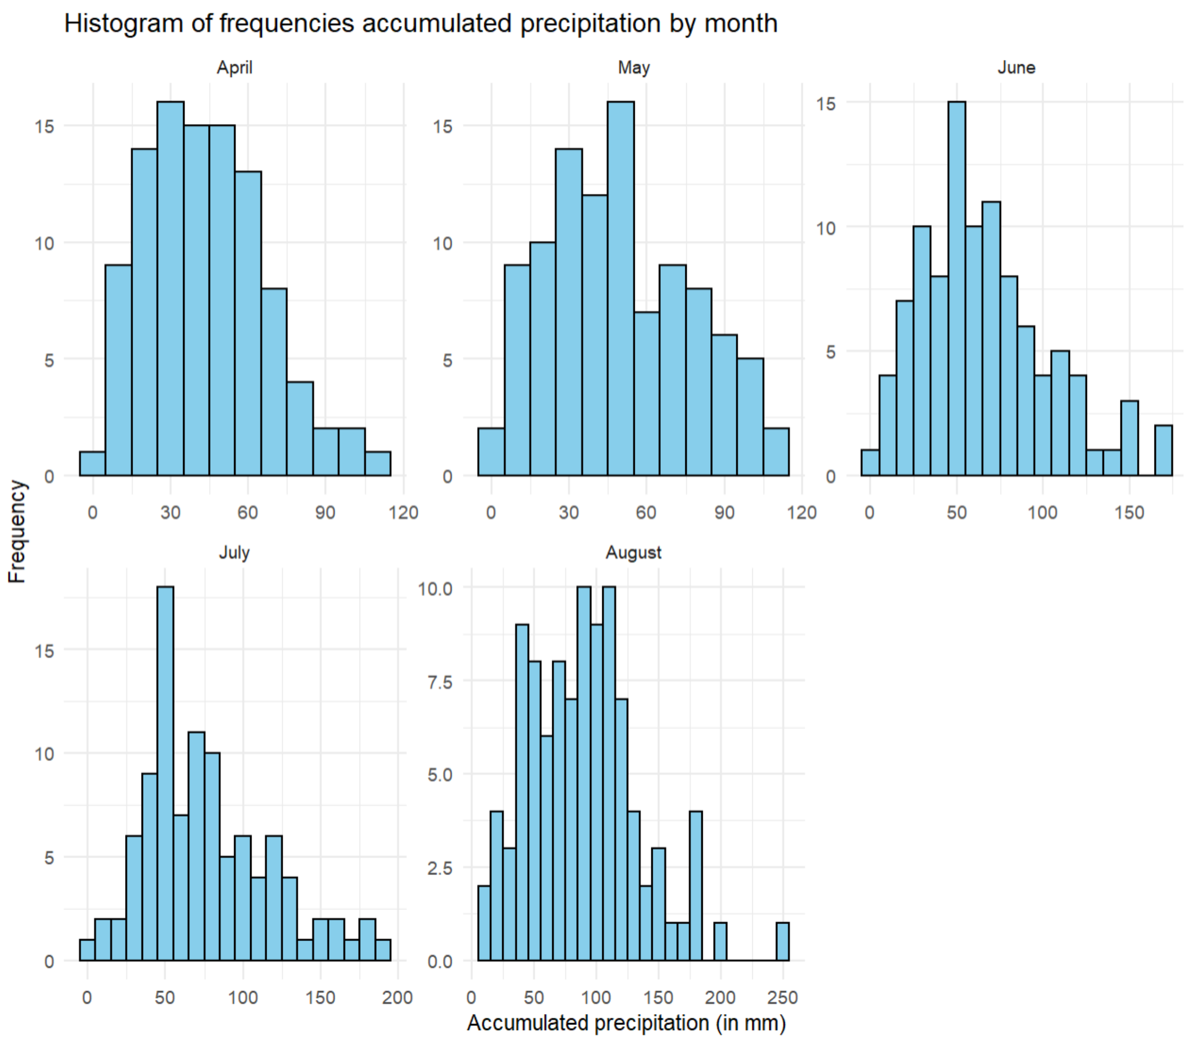

Figure 4: Monthly precipitation in Østfold

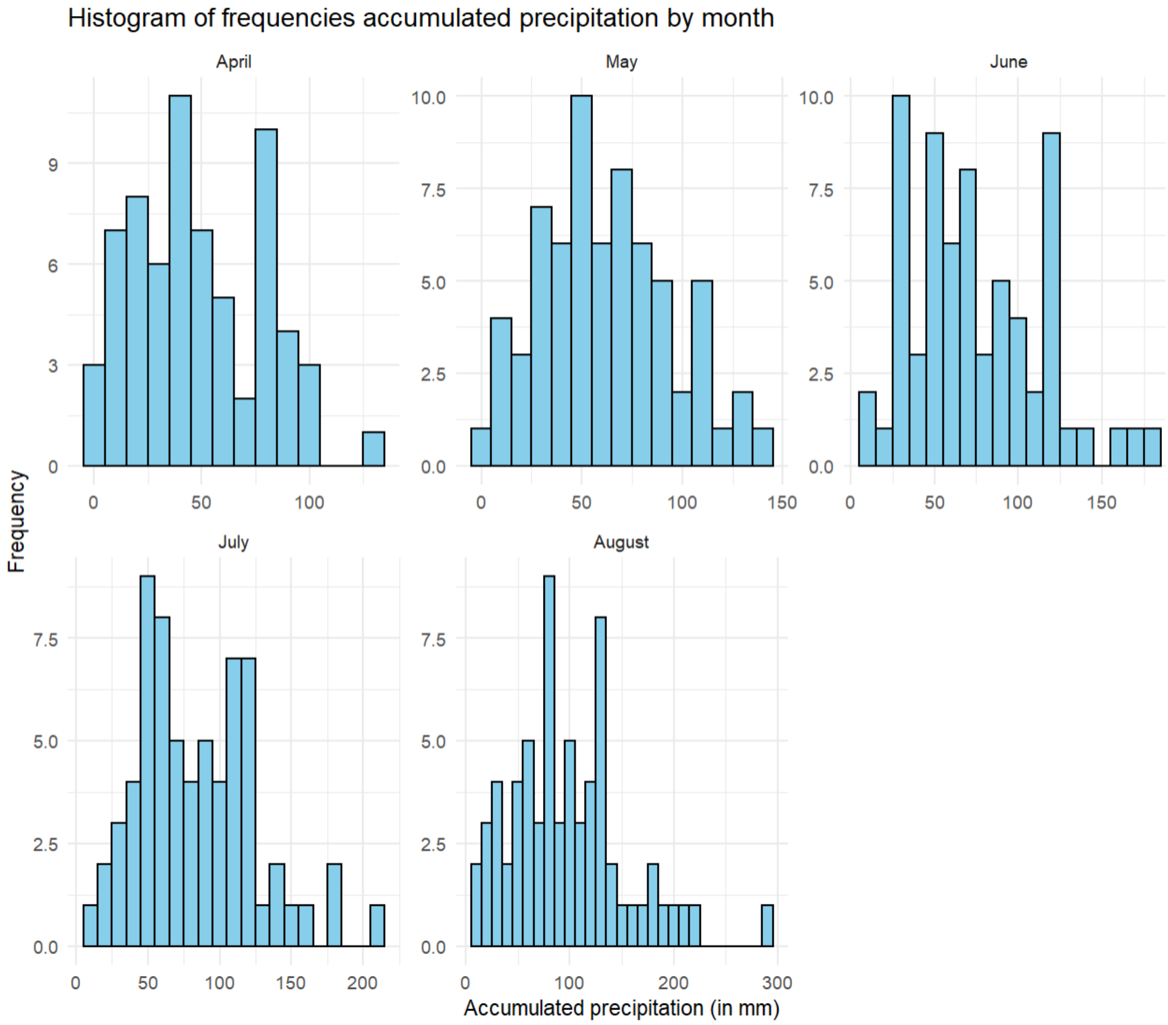

Figure 5: Monthly precipitation in Akershus

When trying to identify and understand the relationship between precipitation and crop yields, we found it sensible to participate in a discussion with a farmer with extensive expertise in the field. An initial idea was to regress the quantity per area on accumulated rainfall over the entire period where the grain is in the soil. However, we were informed this would be a bad approach, as large monthly fluctuations can be very important. A summer with alternating flood and drought will typically have very different effects on a crop, compared to a summer where precipitation is more stable, even if the total accumulated precipitation is the same between the two scenarios. An example of a year with great fluctuations is 2023. The harvest was referred to as one of the worst ever, and the primary reason was weather conditions. May and June were characterized by drought, while the occurrence of the extreme weather event “Hans” in August flooded large areas and destroyed crops for many farmers (Børresen, Mo, Stenberg, & Sveen, 2023). Both drought and flood, if the drainage systems are not efficient, pose significant threats to crop yields. For this reason, a regression model would therefore need to somehow consider these fluctuations, and not just aggregate the precipitation of the entire period. A sensible approach would therefore be to use accumulated precipitation by months as explanatory variables. The idea would then be to find suitable contracts for each month to hedge against weather risk.

As weather forecasts typically only predict weather for about a week ahead, there is very little benefit in postponing the design of the precipitation options while waiting for weather forecasts. It would also add a complex factor into the pricing of the derivatives. It could however be interesting to see if there is a relationship between accumulated rainfall in the months prior and the month of the contract itself. Should such a relationship proven to be present, a reasonable approach would be to adjust the hedging strategy based on precipitation levels in the preceding month(s). To assess whether this would be a viable approach, a Pearson correlation test was performed.

The Pearson correlation test is a statistical test that assesses both the strength and the direction of the linear relationship between two continuous variables. The purpose of this test is to examine if there is a statistically significant relationship between precipitation in a month and precipitation in previous months. The null hypothesis is that there is no correlation between the two variables, while the alternative hypothesis is that there is a correlation between the two variables. $r = 1$ indicates a perfect positive linear relationship, while $r = -1$ indicates a perfect negative linear relationship. $r = 0$ indicates no linear relationship. For there to be sufficient evidence to reject the null hypothesis and conclude that a significant correlation exists, the p-value must be smaller than the significance level, typically 0.05. The results of the Pearson correlation test are displayed in the table below:

Pearson correlation test: precipitation in each month vs precipitation in preceding months

	April		May		June		July		August	
	r	P-value	r	P-value	r	P-value	r	P-value	r	P-value
3 months	0,12	0,24	-0,01	0,91	0,01	0,88	0,06	0,58	-0,03	0,80
2 months	0,05	0,61	-0,16	0,70	-0,09	0,35	0,06	0,56	-0,04	0,73
1 month	0,12	0,23	0,01	0,87	-0,11	0,26	0,09	0,38	0,02	0,85

Table 2: Pearson correlation test

Table 2 displays both the strength and p-values of the correlation between accumulated precipitation in a month and of the preceding three, two and one months respectively. The absolute values of the correlation coefficients are small, with -0,16 being the largest value. However, the p-value is in all cases significantly higher than the significance level of 0.05, which indicates that there is not enough evidence to reject the null hypothesis. Overall, the Pearson correlation test concludes that there is no significant relationship between precipitation levels in any of the individual months and its preceding months. This has great implications on the hedging strategy presented in the analysis; it appears unwarranted to update or adjust the hedging strategy along the way.

Analysis

The analysis that follows can be divided into three separate sections that are built upon each other. The first part will aim to estimate the effect rain has on yields of barley in Norway. In this section we will use different regression models in order to give a precise measure on how rainfall affects crop yields. Several models will be presented and both strengths and weaknesses with be discussed. The next section will focus on designing and pricing rainfall derivatives for the months deemed as impactful, based on the results from our regression analysis. Initially, results from the historical burn analysis will be presented, and later serve as a benchmark when using Monte Carlo simulation to price the derivatives. Lastly, the third part will focus on evaluating the effectiveness of the options designed in the previous section. We seek to quantify the effect a product like rainfall options can have on the volumetric risk that Norwegian farmers face due to shifting weather conditions.

Part 1: Regression Analysis

County-wise analysis

Initially, the idea was to create a regression where the aggregated quantity of an entire county was regressed on the accumulated rainfall for April through August, as these are the months when the grain usually experiences its growth cycle. The general expectation of rainfall was that shortfall of rain would be harmful for the crops, while excess rainfall would similarly damage the crops. Thus, it was expected that there would be an optimal amount of rain, positioned somewhere between the extreme cases of drought and flood, and any deviation, in any direction, would lead to a decrease in yield. These expectations would lead to a non-linear relationship between rainfall and production. To account for this, squared terms of rainfall of each month were added to the model. The dynamics of yield against rainfall, meant that we expected to see a negative coefficient for the squared terms. Figure 6 describes roughly how we expected the relationship to look like.

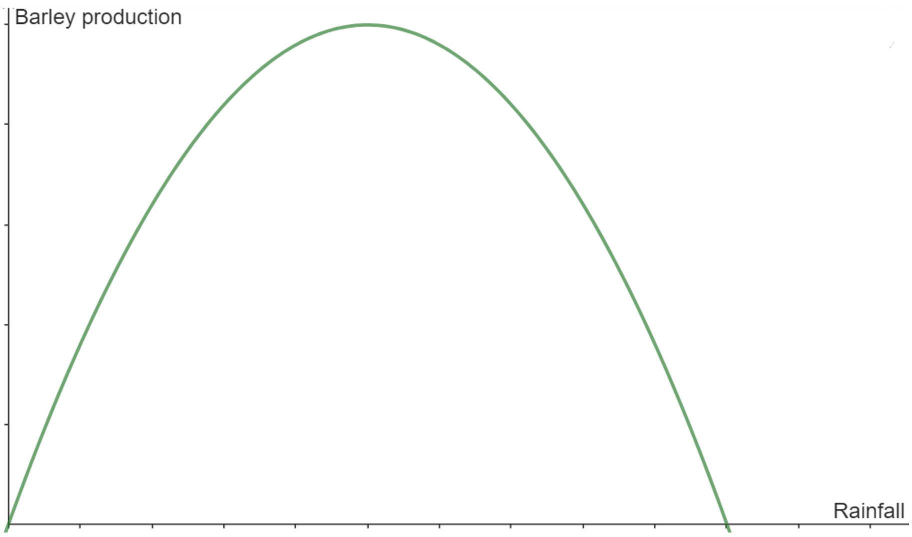

Figure 6: Expectation of relationship between yield and rainfall

The first model includes rainfall by month and their squared terms. As the data originates from a period of 15 years, a time variable, called Index, was included. The variable is expected to pick up the general effects of better technology and more effective production processes that are developed during the time interval. These explanatory variables are regressed against the dependent variable PPA (production per decare). The results of this regression is shown below:

	Dependent variable		
	PPA		
Index	12,372 ***		
	(2,763)		
Pre_Apr	2,658	Pre_Apr_sq	-0,037
	(2,243)		(0,026)
Pre_May	-0,991	Pre_May_sq	0,009
	(1,357)		(0,009)
Pre_Jun	-4,708 **	Pre_Jun_sq	0,024 *
	(2,149)		(0,014)
Pre_Jul	5,142 ***	Pre_Jul_sq	-0,024 ***
	(0,749)		(0,004)
Pre_Aug	3,590 ***	Pre_Aug_sq	-0,012 ***
	(0,836)		(0,003)
Observations	30		
R^2	0,836		
Adjusted R^2	0,719		
F-statistic	7,853 *** (df=11;17)		
Note	*p<0.1; **p<0.05; ***p<0.01		

Table 3: County-wise regression model

With an R-squared of 83.6% and an adjusted R-squared of 71.9%, the explanatory variables explain a lot of the variation in PPA. The Index coefficient shows a positive, statistically significant trend which confirms the expectation that newer technology has led to an increase in PPA throughout the years. In terms of the rainfall variables, the non-squared terms are significant for June, July and August on a 5% level, while the squared terms are significant only for July and August on a 5% level, and June on a 10% level. This suggests that it is rainfall in the second half of the growing season that has the biggest impact on the yield of crops. Furthermore, the months that were significant on a 5% level had negative coefficients for the squared terms, which aligned with the expectations of the relationship between rain and yield. The F-statistic in a regression analysis tests the overall significance of the model. It compares the model with the null model, and measures how likely it is that at least one of the independent variables is significantly related to the dependent variable. The critical value on a 5% significance level in this case is 2.41, and as the F-statistic exceeds this value, it is highly improbable that all the coefficients collectively have zero effect.

Despite the results from this regression being in line with the expectations, a key issue remained: we only had a limited number of observations of the dependent variable, more specifically there were only 30 observations, one for each county over 15 years.
The lack of observations in PPA was concerning and raised doubt about the reliability of the results. As the delivery data had information on an individual level, both in quantities and size, the number of observations would significantly increase by rather looking at organization numbers. This would introduce more variance in the data, which again would improve the model and increase the reliability of the estimated effects. By making a regression on a farm-specific level it would be possible to consider the individual differences between farms and thus improve the precision of our estimates. Individual differences that could affect the results include the quality of the soils, better machinery, crop rotation and a different understanding of cultivating crops in general.

Fixed effects and random effects

The next two models that follow will utilize panel data regression. It is worth noting that the model that was just presented is also a panel data regression, where the two counties act as the individuals. As the data consists of only two individuals, a random effects model is not applicable, and a fixed effects model must be used. Panel data refers to data on a set of individuals over time, and in our case the individuals are organization numbers. An important aspect of panel data regression is handling the unobserved individual specific effects. There are several ways of controlling for these, but the most common are to either use a fixed effects model or a random effects model. The random effects method assumes no correlation between unobserved specific effects and the explanatory variables, while the fixed effects method does not require such an assumption. However, if there are no signs of correlation between the unobserved individual specific effects and the explanatory variables it is preferable to use random effects because it utilizes both within and between-entity variation while fixed effects only use within variation. Using a random effects model when a fixed effects model is appropriate, might lead to biased estimates (Mustafa, 2023).

A common method of determining whether a fixed effects or random effects model is more suitable is the Hausman-test. This test examines if there is a correlation between the time-invariant individual specific effects and the explanatory variables. If the test concludes there is a significant relationship, then a fixed effects model is preferred. If the Hausman test finds no significant correlation, the random effects model is preferred as it requires fewer assumptions and is more efficient. The choice to use random or fixed effects in the following models will depend on the results from the Hausman test. In our regressions there could be several unobserved individual factors that are correlated with rainfall. One example is soil quality, as the soil varies between each individual farm due to factors such as crop rotation and weed management. There could also be differences between counties, as different magnitudes of rainfall may reform the soil differently, leading to different characteristics. Should this be the case, which is also our prediction, a fixed effects model is preferred.

Organization-based analysis

As the previous model used the aggregated data on a county level, it did not take individual variation into account, and the two following models will relate to organization number. Additionally, it was found sensible to add temperature as a control variable. From discussions with people in the Norwegian grain industry, it was indicated that temperature did influence the crop yield. However, temperature in Norway was not considered as a main deciding factor of crop yields. The effect is primarily limited to impacting the length of the growth period.
Given this understanding of the interplay between temperature and yield, temperature appeared to be a good control variable, which could improve the efficiency of precipitation estimates. We

anticipated that temperature would have a significant effect on the yield, and that its interaction with rainfall could provide valuable insights. The results from this regression can be seen in Table 4.

	Dependent variable		
	PPA		
Index	9,247 ***		
	(0,668)		
Ned_Apr	3,112 ***	Ned_Apr_sq	-0,035 ***
	(0,444)		(0,005)
Ned_Mai	-0,320	Ned_Mai_sq	0,004 ***
	(0,252)		(0,002)
Ned_Jun	-3,803 ***	Ned_Jun_sq	0,012 ***
	(0,553)		(0,004)
Ned_Jul	3,561 ***	Ned_Jul_sq	-0,018 ***
	(0,265)		(0,001)
Ned_Aug	2,755 ***	Ned_Aug_sq	-0,009 ***
	(0,209)		(0,001)
Temp_Apr	4,305 ***		
	(1,631)		
Temp_Mai	-3,603		
	(2.435)		
Temp_Jun	-1,329		
	(2,130)		
Temp_Jul	-16,068 ***		
	(1,681)		
Temp_Aug	10,006 ***		
	(2,256)		
Observations	5579		
R^2	0,403		
Adjusted R^2	0,401		
F-statistic	3222,243 ***		
Note	* $p<0.1$; ** $p<0.05$; *** $p<0.01$		

Table 4: Organization-based regression model

Compared to the first regression model, the adjusted R-squared has dropped dramatically from around 72% to 40%. This was aligned with expectations, as introducing farm specific observations into the regression induces a big increase in the variation of the dependent variable. Although introducing monthly temperatures as control variables increases the explanatory power, the effect of increased observations outweighs this effect. Even though the explanatory power decreases, this does not indicate that one model is better than the other, as the number of observations is very different. In fact, the standard errors of the rainfall coefficients have dropped dramatically, supporting the expectation that the model would be more precise. The index coefficient still proves to be highly significant, providing additional confirmation that improved technology and methods has had an impact on yields. The temperature coefficients prove to be significant for the months of May, July and August, suggesting that crops benefit from higher temperatures in April and August and lower temperatures in July, relative to what has been observed in the time period.

For the rainfall coefficients, not only is July and August significant, April and June are now also significant, while May only has a significant squared term. With respect to the squared terms, July and August still have negative coefficients, so does April. However, the coefficients of the squared terms of May and June are positive, indicating the opposite curvature properties of what was expected. It is hard to find a logical reason for this, as it was expected that both drought and flood would severely damage the crops. We believe there could be several reasons for these results, one possibility is that the data may contain a significant amount of randomness. Factors such as incorrect delivery data or incorrect area data may have skewed the results, while weeds and fungi could also have had an impact. Another argument is that, due to only using one weather station for each county, there may simply not be enough variation in the rainfall data to accurately describe the relationship between yields and rainfall.

If the precipitation measurements for May and June during the time interval are less representative of typical rainfall patterns, this could lead to unexpected results in the regression analysis. Another reasoning has to do with the interplay between months. Large amounts of rain in June for example, which has been the case in some years, could be outweighed by the effect of shortfall of rain in the preceding or following months, and thus leading to less efficient estimates. In other words, the interplay between months may have resulted in some unexpected findings. It would be beneficial to go back further in time and include more observations of both rainfall and production per decare, however, we do not possess data on farm sizes and production before 2005. Even though the model shows an unexpected relationship between barley crop yields and rainfall in May and June, there are several insights from this model in terms of interpreting the overall relationship between rain and crop yields. The large increase in observations give more reason to put faith in the last model than the previous one, as the data utilized contains more variation. Lastly, the F-statistic suggest that it is highly unlikely that all the coefficients collectively have zero effect. This regression model used a fixed effects model, based on results obtained from the Hausman test.

Organization-based analysis: rainless days

Although the previous regression models prove a significant relationship between precipitation and crop yields, we found it sensible to use another regression model as the basis for the derivatives. As mentioned earlier, there has been a development over the last couple years where the subsidies for drainage systems have increased. Additionally, conversations with our agricultural contact indicated that a well-drained field would encounter the problem with excess rainfall very differently than a poorly drained field, as a properly implemented drainage system would, to a great extent, prevent any issues from arising. We acknowledge that hedging against flood conditions in the past might have been beneficial for farmers, but as of today, and moving forward with recent developments, we consider the major obstacle to crop yields to be drought. For the analysis, the primary focus will thus be on drought conditions, and we introduce number of rainless days in a month as an explanatory variable.

The aim of the new regression model is to quantify the impact of shortfall of rain on yields, that consequently would be used as tick sizes in the option contracts. As we more or less set aside the issues of flooding in the regression analysis, the explanatory power of the model would also decrease significantly, but this was not considered as a big issue. The new regression model can be found below:

$$PPA = Index + \sum_{m=4}^{8} \alpha_m * DA_m + \sum_{m=4}^{8} T_m$$

where DA_m (days above), is the number of rainless days per month above its historical mean. Consequently, α_m quantifies the impact of one additional rainless day in a month on the PPA. If the number of rainless days in a month falls below the historical mean, the DA variable takes the value zero. T is the temperature and is included to improve the efficiency of α_m. The regression is a PLM model, which allows for controlling individual effects. In the county of Østfold, the historical means of rainless days in the period April-August is 19, 19, 19, 18, 17 respectively. This model provides a very intuitive and simple basis for the derivative contracts. Should there for example be 23 rainless days in April, the payout of the weather derivative for April would be $4 * a_4$. As this regression model is based on additional rainless days, relative to the historical mean, we would expect coefficients to be negative. In other words, another rainless day worsens the growing conditions, leading to poorer yields. The coefficients of the regression model can be found below:

	Dependent variable
	PPA
Index	0,588
	-0,722
DaysAbove_Apr	-9,345***
	(1,451)
DaysAbove_May	-46,153***
	(6,062)
DaysAbove_Jun	75,452***
	(7,567)
DaysAbove_Jul	59,116***
	(8,414)
DaysAbove_Aug	-15,253***
	(2,554)
Temp_Apr	77,376***
	(10,535)
Temp_May	-9,041
	(6,006)
Temp_Jun	-32,982***
	(5,410)
Temp_Jul	-74,812***
	(9,845)
Temp_Aug	-117,729***
	(13,240)
Observations	2124
R^2	0,405
Adjusted R^2	0,340
F-statistic	118,503*** (df=11; 1915)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 5: Organization-based regression model: rainless days

All a_m coefficients are highly statistically significant, p-values are infinitely small. Some of the coefficients are, however, positive. This is not in line with the general expectations, and the results suggest that drought in June and July is beneficial. These positive coefficients may reflect complex interactions between rainfall, temperature and growth stages. A possible explanation is that crops often have different water requirements at different stages, and some crops may be able to withstand a period of dryness without having a big impact on yield. Another possible explanation is that a period of dry weather might be associated with fewer diseases or fungal infections, which could have a positive effect on yields. Fungal growth thrives in wetter conditions, and the model might capture some of the trade-off between drought and disease. The effect of temperature may also interact in complex ways with rainfall.

Another explanation is that, since the regression model is only based on 15 years of data, and the historical means are based on a significantly longer period. There might not have been enough dry June and July observations to push the coefficients into the expected negative direction. If April and May are wet, crops could suffer from waterlogging, experiencing oxygen deprivation and restricting root growth. If drier conditions in June and July follow, it can help drying out the soil, possibly leading to better crops. Although the signs of some coefficients are surprising, it is important to remember that there are many factors other than just rainless days that plays a part in crop yields. With an adjusted R-squared of 34%, the explanatory power is lower than the previous regression model, but this was expected, as we only aim to quantify the impact of drought over the sample period, and not flood conditions. The explanatory variables are also different; instead of accumulated rainfall by month, and their squared terms, this model uses rainless days. This is one of the reasons why the signs of each month may differ across the three models.

The coefficient of DaysAbove_Apr tells us that an additional rainless day in April above the mean, is associated with a reduction of 9.34kg barley per decare, on average. The coefficient of May is almost five times larger, indicating that an extra rainless day in May has much greater impact on crop yields. For the hedging strategy, a sensible approach would therefore be to purchase calls on months April, May and August, in which an extra rainless day has a proven negative impact on yields. The coefficients of these months would act as the tick size in the derivatives contracts.

Part 2: Pricing

Historical burn analysis

The historical burn approach utilizes historical data to find hypothetical payouts of the proposed derivative contract. By taking the average of the payouts, discounting them and adding a risk premium, a price for the derivative can be found. This will serve as a benchmark for the Monte Carlo approach that is set to follow.

The regression model on rainless days quantified the impact of an additional rainless day on PPA. These will serve as tick sizes for the derivative contracts. A key consideration in this context is that the tick size is based on quantity per area, and not actual monetary values. The reason for this is that prices of barley, and other types of grain, is set on a nation-wide scale annually. Every year, an agricultural settlement, *Landbruksoppgjøret,* takes place through negotiations between the state, *Norges Bondelag* and *Norsk Bonde- og Småbrukarlag* (Regjeringen, 2024). During these negotiations the parties agree on prices for each type of grain. With this information in mind, the farmer can only hedge the volumetric risk, which is why the tick sizes are based on quantities. In practical terms, this means that every year the tick size would simply be multiplied by the annual agreed upon price, which for 2024 was 4.03 kr/kg for barley (Felleskjøpet, 2024). Instead of adding the step of multiplying all observations of PPA by the price in the corresponding year, only to adjust it back to, for example 2005 or 2019 prices, we found it more sensible and far simpler to only base the tick sizes on quantities in kg per decare. The payout of the derivatives based on rainless days can then be written as:

$$Payout = TS_m * \max\left(RD_m - S_m, 0\right)$$

where RD_m is the number of rainless days in month m, while S_m is the strike level, set to the historical mean (also rounded). TS_m is the tick size of each month and has the property $TS_m = \alpha_m$ from the last regression model.

For Østfold we use 100 years of data, which means that the historical burn analysis is based on 100 observations per month. Each of the hypothetical payouts are discounted at the risk-free rate for one, two and five months respectively, to find the day 0 value of the payout. In the Norwegian market, the most commonly used benchmark for the risk-free rate is the 10-year government bond yield (PWC & FFN, 2023). Between 2005 and 2016, the average rate was

3.2% (SSB, 2017), and this will be used in the pricing of the derivative. As the discounting period is very short, any small change to the risk-free rate has very small impact on the discounted payouts.

During the introduction of the Monte Carlo simulation it was mentioned that when discounting returns from weather derivatives, it is not sufficient to only use the risk-free rate like one would do with financial derivatives. The reason being that the underlying of weather derivatives is not tradeable and thus the position is not possible to fully hedge. Thus, a risk premium equal to the market price of weather risk must be added. Cao and Wei argue that in theory the risk premium should be zero if weather risk could be hedged, however, empirical findings suggest that the risk premium is in fact positive (Cao & Wei, 2004). Härdle and Cabrera also find that the risk premium for temperature derivatives is in fact non-zero and positive, suggesting that the use of only the risk-free rate is in fact incorrect (Härdle & Cabrera, 2009). Cao and Wei provide a framework for how to estimate the risk premium for weather risk by using historical temperature data and assumptions about risk aversion, and certain economic parameters. The limitation of this method in our case is its inapplicability to rainfall risk. The approach is designed to assess weather risk for temperature derivatives, and it cannot be assured that the risks associated with temperature and rainfall will align.

Another more unconventional way of estimating the market price of weather risk is using online platforms that provide weather derivatives traded OTC. These platforms provide prices for option premiums and hence implied risk premium can be extracted. This method is used by Blom in his thesis from 2009 where he uses a platform called Weatherbill to estimate prices for specific contracts that are almost 100% guaranteed to provide a payout. For a price of 110 dollars, the contract pays out 100 dollars if the minimum temperature during a summer day is below 50 degrees Celsius in Oslo. Given that the payout is more or less guaranteed, the price of the contract can then be used to find the risk premium. This method results in a risk premium of 10% (Blom, 2009). This method is unconventional, yet creative. However, the same issue still remains as the contract is based on temperature and not rainfall. We will thus provide a third and final way of estimating the risk premium for rainfall options.

A typical counterparty of the proposed weather derivative contracts could be insurance companies. A common measure of profitability in the insurance industry is the “combined ratio". The ratio can be found by dividing the sum of incurred losses and expenses by the earned premium. A ratio below 100 percent indicates that the company is making an underwriting profit, while a ratio of more than 100 percent indicates that it pays out more in claims than it receives from premiums. The combined ratio does not take investment income into account, which is typically a substantial share of total income for these companies. Many insurance companies still consider the combined ratio as the best way to measure success as it only includes profits earned through efficient management (Hayes, 2020).

To account for risk in the price of the derivatives, we therefore find it sensible to multiply the average discounted payouts by the inverse of the combined ratio. Taking the inverse approximates the surplus or profit margin relative to premiums, which can approximate the risk premium. This margin adjusts for the insurer`s risk exposure, and without some sort of margin any counterpart would have no incentive to enter a contract. Norway's largest insurance company, Gjensidige, aim for a combined ratio of 84% in 2024 and 2025 (Vosgraff, 2024). Taking the inverse of this number, we get an option price multiplier of 1.19, implying an expected margin of 19%. It is important to emphasize that this does not constitute a more traditional risk premium and should more accurately be described as a form of margin-loading the price. We believe this method to be the best way of risk-adjusting the price of the derivatives for two reasons. Firstly, the method is realistic and provides a real-world view of what a counterpart like an insurance company would require as a premium. Additionally, it is a conservative measure compared to that we have seen from other papers. This would underestimate the effectiveness of the derivatives in the consequent analysis, and it is advisable to take a cautious approach when evaluating effectiveness.

The price of the call options using the historical burn analysis can be found by the following formula:

$$C_m = e^{-rt} * \frac{1}{n}\sum_{n=1}^{100} \max\left(RD_{m,n} - S_m, 0\right) * TS_m * \frac{1}{CR}$$

where r is the risk-free rate, t is time, TS_m is tick size, and CR is the combined ratio. The expectation of rainless days above the historical mean is the average of the annual differences

between rainless days per month and the strike level, where strikes also are set to the historical means. The prices of each derivative contract are summarized in
Table 6 below.

	Tick size	Maturity (in months)	Strike (rainless days)	Risk-free rate	CR	**Price**
April	9,345	1	19	3,2 %	1,19	**18,189**
May	46,153	2	19	3,2 %	1,19	**99,972**
August	15,253	5	17	3,2 %	1,19	**38,508**

Table 6: Historical burn analysis prices

By the historical burn analysis, the prices of the weather derivatives for April, May and August, with different tick sizes, is 18.189, 99.972 and 38.508 respectively. These prices are based on only 100 observations per month, which is far less than a typical number of simulations in a Monte Carlo pricing approach. With fewer observations, the probability of skewed prices increases, but the historical burn analysis still provides a benchmark for the Monte Carlo simulation that follows.

Monte Carlo simulation

The main assumption of using Monte Carlo simulation is that a distribution for the underlying must be selected. Therefore, conducting a thorough analysis of the data to identify an appropriate distribution is critical. Before selecting a distribution, it may be useful visualize the data to gain insight into which approximation might be suitable. In Figure 7 histograms of frequencies of rainless days per month in the last 100 years is illustrated.

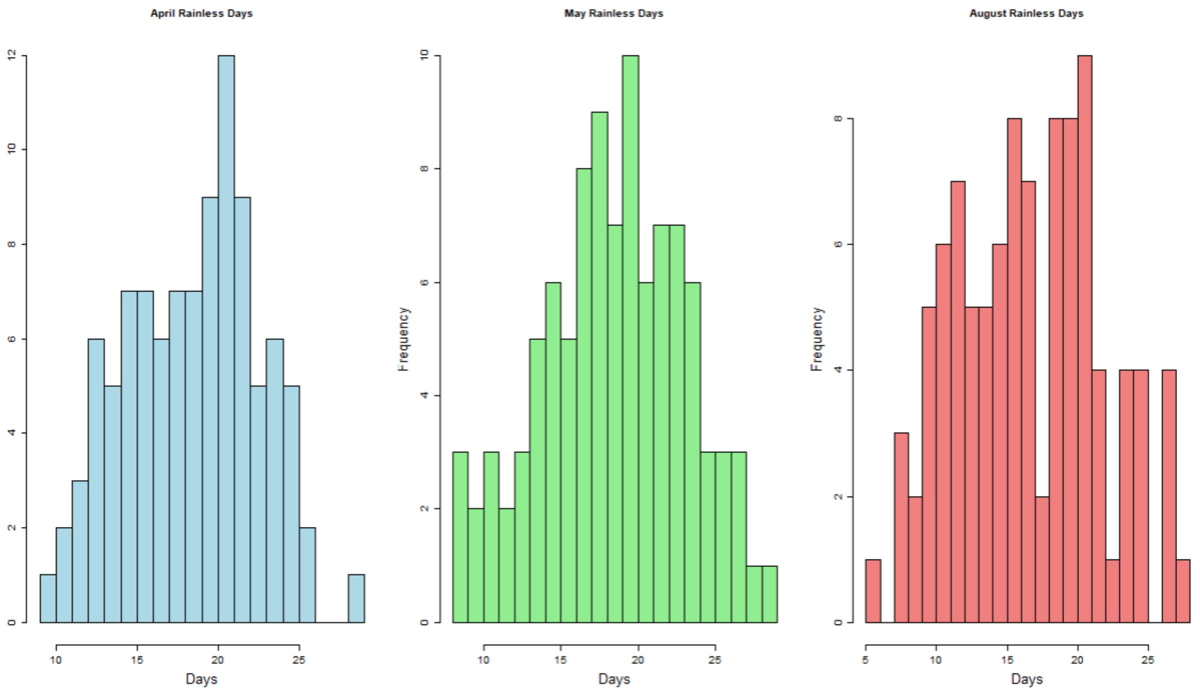

Figure 7: Rainless days distribution Østfold

Month	Mean	Standard deviation
April	18,88	4,08
May	18,89	4,66
August	16,99	5,16

Table 7: Means and standard deviations - rainless days

From Figure 7 it is evident that the distribution for each month has some resemblance to a normal distribution. However, the mean is not centrally located, it is skewed towards the right side of the x-axis, as confirmed by the means presented in Table 7. Even though each histogram shares some similarities, there are still some differences that will influence the choice of distribution and thus the price of the contract for that specific month. The most observed number of rainless days for each month is similar at around 18-22 days, but looking at frequencies, the August distribution is somewhat flatter than the two other months, implying a larger standard deviation, as confirmed in Table 7. Lastly, examining extreme observations is vital, as they can significantly impact option prices. Notably, there are more extreme observations in August compared to April. While April only has a single observation below 10, August has several. In the other end of the scale, the figure confirms that August also has more observations of 25 and over, than April.

The next step after visualizing the data is to fit a suitable distribution for each month. In order to find the most appropriate distribution, Maximum Likelihood Estimation (MLE) will be used. After reviewing the data, we decided to try four different distributions; normal, Weibull, lognormal and gamma. To determine which distribution has the best fit for each dataset, the Akaike Information Criterion (AIC) is utilized. This measurement is commonly used when deciding which distribution has the best fit to a given dataset and tests the goodness of fit while penalizing excessive use of parameters and model complexity (Kenton, 2023). The value of the measure is calculated as:

$$AIC = 2k - 2\ln(L)$$

where k represents the number of parameters and L is the goodness of fit. Another commonly used measurement is the Bayesian Information Criterion. The formula shares some similarities with AIC and can be found below:

$$BIC = \ln(n)\, k - 2\ln\,(L)$$

As the second terms of the two respective formulas are identical, the only true difference is that AIC multiplies *k* by 2, while BIC multiplies *k* by *ln(n)* where the latter usually is larger. In our case, this indicates that the BIC will be greater in absolute terms, but the ranking between the different distributions will be unchanged. For this reason, when selecting distributions, only AIC values will be compared. AIC is a relative measure used to compare multiple distributions to determine which one is better. A lower AIC indicates a better fit. The tables below present the AIC scores of each distribution across the three months in question.

Distribution	AIC
Normal	569,1224
Weibull	567,9033
Lognormal	576,8155
Gamma	572,8557

Table 8: AIC scores for April

Distribution	AIC
Normal	595,6934
Weibull	593,9114
Lognormal	608,1893
Gamma	601,9199

Table 9: AIC scores for May

Distribution	AIC
Normal	616,1413
Weibull	614,6389
Lognormal	623,8379
Gamma	618,2048

Table 10: AIC scores for August

From Tables 8-10, it is evident that the Weibull distribution is the best fit for all months, as it has the lowest AIC values. However, the differences between Weibull and normal distribution are small, and normal distribution may serve as a suitable alternative choice. The gamma distribution provides similar AIC values for certain months but is overall inferior to the Weibull and normal distributions. The lognormal distribution, on the other hand, is the least suitable fit across all months

Even though the Weibull and normal distribution proved to be the best approximations, they are not entirely optimal. An important property of the underlying historical datasets is that there are no values below zero, and no values above 30 or 31. In other words, it is impossible to have less than 0 rainless days in a month, nor more than the number of days in a month. The distributions above do not fulfil these requirements.

When examining the probability of the normal distribution being less than 0, a small, but positive number is obtained for all months. The probability of exceeding 30 days in April is 0.32%. The Weibull distribution is only defined for positive numbers and thus have 0% chance of being below 0, however, the probability of exceeding 31 days for the month of August is 0.20%. Even though these probabilities are generally very small, they may still have an impact on the prices of the options, as the occurrence of extreme values could result in a few simulations with unattainable payoffs. Using these distributions could thus lead to inaccuracies and bias in the estimates of the prices of the option contracts.

To take these properties into account, it is sensible to introduce truncated distributions, which limits the range of obtainable values. As the Weibull and normal distributions proved to be better fits than gamma and lognormal distributions for all months, only the two former distributions will be truncated. After fitting the best truncated normal distributions, we are left with the following results:

Distribution	Mean April	Mean May	Mean August
Normal	18,88	18,89	16,99
Truncated normal	18,93	18,97	17,04

Table 11: Means of truncated and non-truncated normal distributions

Distribution	SD April	SD May	SD August
Normal	4,08	4,66	5,16
Truncated normal	4,15	4,77	5,27

Table 12: Standard deviations of truncated and non-truncated normal distributions

The means and the standard deviations of the normal distributions have increased for all months when utilizing a truncated normal distribution. This implies that the distributions have become flatter, as well as slightly skewed to the right. To compare the fit of normal versus truncated normal, the AIC values can be examined.

Distribution	AIC April	AIC May	AIC August
Normal	569,1224	595,6934	616,1413
Truncated normal	568,4239	594,6467	615,2993

Table 13: AIC scores normal vs truncated normal distribution

Based on the AIC scores for each distribution, it is evident that the truncated normal distribution provides a superior fit for all months when compared to the non-truncated normal distribution. Below, the same comparison is conducted for the Weibull distribution.

Distribution	Shape April	Shape May	Shape August
Weibull	5,283715	4,637500	3,669002
Truncated Weibull	5,283376	4,637382	3,669167

Table 14: Shape parameters

Distribution	Scale April	Scale May	Scale August
Weibull	20,51014	20,68170	18,86126
Truncated Weibull	20,50961	20,68222	18,86034

Table 15: Scale parameters

Distribution	AIC April	AIC May	AIC August
Weibull	567,9033	593,9114	614,6389
Truncated Weibull	567,9033	593,9114	614,6389

Table 16: AIC scores Weibull vs truncated Weibull distribution

Similar to how the mean and standard deviation establish the appearance of a normal distribution, the shape and scale parameters establish the appearance of the Weibull distribution. While the shape parameter governs the form of the distribution, the scale parameter stretches or compresses the distribution along the x-axis. From Table 14 and Table 15, it is apparent that the differences in shape and scale parameters between the Weibull and truncated Weibull distribution are negligible, and have no effect on the AIC. It was expected that the Weibull distribution would be less affected by truncating it compared to the normal distribution, as the former do not allow for negative values and the area that exceeding 30 or 31 were smaller than for the normal curve. The subsequent analysis will utilize the truncated Weibull distribution as the preferred option, but results from using the truncated normal curve will also be presented.

After selecting the distribution of the underlying, the next step is to price the option contracts through a Monte Carlo simulation. The risk-free rate and the risk premium are set to 3.2% and 1.19, respectively, similar to the historical burn analysis.

	Tick size	Maturity (in months)	Strike (rainless days)	**Price HBA**	**Price MCS (truncated Weibull)**	**Price MCS (truncated normal)**
April	9.345	1	19	**18.189**	**17.593**	**17.518**
May	46.153	2	19	**99.972**	**98.705**	**99.663**
August	15.253	5	17	**38.508**	**37.299**	**37.484**

Table 17: Monte Carlo simulation prices

By utilizing the Monte Carlo simulation approach, with a truncated Weibull distribution, the prices of the options are 17.593, 98.705 and 37.299 for April, May and August, respectively. The results are derived from 100 000 simulations. First and foremost, it is evident that these prices differ very little from the historical burn analysis, which is a great sanity test. The difference in prices is even smaller when comparing the two distributions in the Monte Carlo approach. This indicates that the choice of distribution does not greatly impact the prices. This also lessens the limitation of the model by assuming a specific distribution for the underlying index.

Part 3: Results

Optimizing

In the previous part of the analysis, the tick sizes of the contracts were set equal to the coefficients from the regression analysis. While these coefficients have an intuitive value, in the sense that it quantifies the impact of an additional rainless day on the PPA, these tick sizes are not necessarily optimal for reducing the variation in crop yields. To find the impact of purchasing derivative contracts, a comparison between the old PPA and new PPA can be conducted. While the old PPA represents the production per area without hedging, the new PPA incorporates both the payouts of the option contracts and the option premiums paid, of which the latter will remain constant across all years. Once again, all these values are based on quantities (kg per decare), not monetary values, and would in practical terms have to be multiplied by the annual predetermined price (in kr/kg). The new PPA can then be written as:

$$New\ PPA_{y,i} = Old\ PPA_{y,i} + \sum_{m=4}^{8} q_m (DA_m * TS_m - C_m) \tag{1}$$

where y is year, i is the organization number, DA_m is how many rainless days above the strike in a month, TS_m is the tick size of the month and C_m is the price of the call option purchased. q_m is defined as the quantity of an option purchased in a month, where the tick size is set to the coefficients from the regression, as a reference point. As the formula includes both tick size and the call option premium, which has a linear relationship with the tick size, q_m allows for scaling the optimal tick size and the corresponding price. As the use of weather derivatives is aimed to reduce risk, it can be sensible to compare the change both in standard deviation, but also in a simplified version of the Sharpe ratio. While the former metric is more common in the existing literature, the latter also aim to limit the reduction in mean PPA values. A "naïve" approach would be to purchase contracts with tick sizes found from the regression, which is equivalent to setting $q_m = 1$ for all m. The results of this approach can be found in the table below.

Mean		SD	
Old PPA	New PPA	Old PPA	New PPA
458,88	357,25	105,53	96,52

Table 18: Changes in PPA - naïve approach

If all farmers had purchased calls for April, May and August every year they were producing barley between 2005 and 2019, the average annual production per decare (and thereby also net income) would decrease by roughly 22.1%. Since the derivatives essentially function as a form of insurance, it is expected that net income would decrease in the long run. In other words, the results align with our general expectations about direction, while magnitude remains more difficult to predict. During single years of additional rainless days, net income may rise from purchasing rainfall derivatives, but only if the number of additional rainless days is sufficient to offset the cost of call option premiums. In the long term, however, the mean is expected to decrease, as the option prices include a general margin loading to account for risk. The standard deviation of the PPA falls with roughly 9.5%, which is also in line with our expectations: by hedging against adverse weather conditions, the variability in production should subsequently decrease.

The Sharpe ratio is among the most commonly used financial metrics, which compares the return of an investment with its risk (Fernando, 2024). The Sharpe ratio can be written as:

$$Sharpe\ ratio = \frac{R_p - R_f}{\sigma_p}$$

where R_p is the expected return of the portfolio, R_f is the risk-free rate and σ_p is the standard deviation of the portfolio`s excess return. As we, in this scenario, do not have a risk-free alternative to PPA, we can introduce a simplified Sharpe ratio as:

$$SSR_i = \frac{Mean\ PPA_i}{SD\ PPA_i}$$

As the naïve approach resulted in PPA decreasing by roughly 22.1%, and the standard deviation only by 9.5%, the SSR consequently decreases, and the strategy is not economically sensible. Although the standard deviation decreases, it is outweighed by a large reduction in mean PPA, which is undesirable. An alternative hedging approach is to set different tick sizes, and thereby also different option prices, by changing the weights q_m.

The first optimization approach is based on minimizing the mean standard deviation across organization numbers. More specifically, each organization number`s standard deviation of

PPA, after hedging, is estimated and thereafter the weights q_m are changed to minimize the mean of standard deviations. To ensure a certain degree of robustness and relevance of the tick sizes, only organization numbers with observations in at least half of the interval`s years are included in the optimization process. This restriction intends to avoid biasing estimates toward organization numbers with sporadic data, potentially distorting the estimates. By focusing on entities with more consistent data, the parameters better reflect tendencies relevant to the full timeframe. There are, however, very few organization numbers with sporadic data, as the data already was limited to the 500 largest producers, and this restriction therefore has very little impact on the optimized values. The tick sizes found from the regression was 9.344, 46.154 and 15.253 for April, May and August respectively. Optimization shows that the optimal tick sizes, after applying the optimal weights, are 1.689, 31.246 and 9.455. In other words, it is optimal to set all tick sizes lower than the regression results suggest. This leads to the following results:

Mean		**SD**	
Old PPA	New PPA	Old PPA	New PPA
458,88	392,76	105,53	92,33

Table 19: Changes in PPA - first optimization approach

By applying this strategy, the average annual PPA (and thereby net income) decreases only by 14.4%. For comparison, the naïve approach led to a decrease of 22.1%. Additionally, the standard deviation falls by 12.5%, roughly 3% greater than the naïve approach. In other words, both metrics indicate better results by using optimized weights. However, as the decrease in mean PPA is larger than the decrease in standard deviation of PPA, the simplified Sharpe ratio decreases, and the hedging-strategy is still not profitable. These results clearly illustrate the weakness of optimizing the weights to minimize standard deviations. By only minimizing the variation, this approach fails to achieve the goal of keeping the change in mean PPA as small as possible. Further analyses show that even when setting the option prices unrealistically high, in our case multiplied by a million, and holding the tick sizes constant, the optimized weights remain almost unchanged. This indicates that it still would be preferable to purchase these overpriced options, solely because the standard deviation decreases. The PPA however becomes negative for all organization numbers, which clearly is sub-optimal. This optimization strategy clearly shows that focusing solely on reducing the total variation, overlooks the fact that the risk-adjusted performance decreases, and the aim of minimizing the reduction between PPA and new PPA is not taken into account.

With this in mind, another possible approach is to find the weights that maximize the mean of SSRs after hedging. For each organization number, mean of new PPA and standard deviation of new PPA is estimated. These values are found by calculating the mean and standard deviation of all the annual observations of an organization number. The simplified Sharpe ratio, for organization number i, is then given by:

$$New\ SSR_i = \frac{Mean\ New\ PPA_i}{SD\ New\ PPA_i}$$

The objective function of the optimization is:

$$\max \frac{1}{i} \sum_{i=1}^{I} New\ SSR_i$$

with the constraint $q_m \geq 0$ and where new PPA is defined by (1)

The optimization process finds the optimal weights q_m that maximizes the mean of SSRs. Results show that the optimal tick sizes go from 9.344, 46.154 and 15.253, to 1.465, 17.944 and 0 respectively. It is, in other words, no longer optimal to purchase call options for August. For April and May, it is beneficial to set the tick sizes lower than what the regressions results suggest. The results of this hedging strategy can be found below:

Mean		SD	
Old PPA	New PPA	Old PPA	New PPA
458,88	430,93	105,53	95,07

Table 20: Changes in PPA - second optimization approach

The mean PPA falls by 6,1% relative to no hedging. The standard deviation also falls by 9.9%. While the simplified Sharpe ratio decreased for both the naïve and the minimized SD approach, this method leads to an increase in the simplified Sharpe ratio by 4.2%, which is a desirable result. The hedging strategy ensured that around 78% of organization numbers had their simplified Sharpe ratio improved. In the best case, an improvement of 33% could be observed. The hedging strategy has been optimized in-sample to achieve the best possible results, and with these improvements in place, the next step is to test how this strategy performs on a new, out-of-sample dataset to examine its robustness and effectiveness beyond the initial sample.

Backtesting

The dataset used for backtesting contains the delivery data of Viken between 2020 and 2023. The structure is the same as the first dataset, containing organization number, year, county number, municipality number, type of grain and quantity. Daily precipitation data was retrieved from Norsk Klimaservicesenter. To incorporate sizes of area, we had to use the same approach of extracting information from applications of production subsidies. As a result, there is still some uncertainty around the organization numbers` PPA. There could be issues of omitted data, some farmers might have failed to apply in time, or the reported number could be incorrect. As the formation of Viken also included Buskerud, we also ensured that all organization numbers used for backtesting had appeared at least once in the original dataset. The effectiveness of the hedging strategy is examined only on organization numbers that produced barely in all four years of the backtesting interval. We found it desirable to include organization numbers with as many observations as possible, as there already were few observations, and any reduction could lead to great changes in standard deviation.

The optimal strategy from the optimization step was to adjust the monthly option weights to maximize the mean of simplified Sharpe ratios. Applying this approach, it was found optimal to purchase call options for April and May only, with tick sizes of 1.465 and 17.944 respectively. The corresponding prices are 2.756 and 38.351 ("PPA prices").

Implementing the optimized hedging strategy on the backtesting dataset, results in a mean decrease in standard deviations of 16.4%. In fact, as many as 98.9% of the organization numbers had a reduction in standard deviation. In the best case, an improvement of around 35% was observed. The annual results for this specific organization number can be shown below:

Year	PPA	Option payout	Option premium	New PPA	Change in PPA
2020	**386,41**	25,27	41,11	**370,57**	-4,10 %
2021	**426,14**	4,40	41,11	**389,43**	-8,61 %
2022	**486,41**	7,33	41,11	**452,63**	-6,94 %
2023	**297,36**	71,77	41,11	**328,03**	10,31 %

Table 21: Results from hedging strategy - best case

For reference, an organization number which experienced an improvement close to the mean, had the following results:

Year	PPA	Option payout	Option premium	New PPA	Change in PPA
2020	**579,98**	25,27	41,11	**564,15**	-2,73 %
2021	**467,43**	4,40	41,11	**430,72**	-7,85 %
2022	**482,23**	7,33	41,11	**448,45**	-7,00 %
2023	**280,79**	71,77	41,11	**311,46**	10,92 %

Table 22: Results from hedging strategy - random organization number

The two tables illustrate how farmers, to some extent, can stabilize their yields by purchasing these rainfall derivatives. As these two organization numbers are linked to the same weather station, and should consequently experience very similar precipitation patterns, the results also show that PPA is dependent on other things than just rainfall. Between 2020 and 2021 the first of the two organization numbers see their old PPA increase by around 10%, while for the randomly selected organization number it decreases by 19.4%. This great difference in differences could be a result of for example weed growth, fungi, lack of crop rotation, different equipment or just inaccurate PPA numbers. It demonstrates that stabilizing yields can be challenging, as numerous other factors not directly linked to weather may cause large fluctuations.

The option premium paid is constant across all years (in PPA prices), as the optimization process found a single common optimal hedging strategy, independent of year. In the first three of four years, there were relatively few drought days, and as a result, the use of weather derivatives was a net negative. In 2023 however, May was particularly dry, which led to a positive payoff from the derivatives. An important note here is that the 2023 crops were greatly affected by the extreme weather "Hans" which also caused large areas to be flooded, leading to unusually poor yields. Fields with inefficient drainage systems experienced the greatest losses, as crops were deprived of oxygen and drowned. In other words, 2023 experienced both drought and flood, and as the proposed derivatives are only based on drought, many organization numbers would struggle to fully stabilize their yields between 2022 and 2023, as flood management is unaccounted for. Consequently, it may be meaningful to examine the top performer of 2023 in PPA, which, unless there are randomness in the data, is more likely to have an efficient drainage system and be less affected by flood.

Year	PPA	Option payout	Option premium	New PPA	Change in PPA
2020	**539,37**	25,27	41,11	**523,54**	-2,94 %
2021	**537,85**	4,40	41,11	**501,14**	-6,83 %
2022	**653,68**	7,33	41,11	**619,90**	-5,17 %
2023	**537,04**	71,77	41,11	**567,71**	5,71 %

Table 23: Results from hedging strategy - top performer of 2023

Compared to the two other farmers, the yields are much more stable, and the floods of 2023 seems to have smaller impact. The difference between 2022 and 2023 was 116.64 PPA units, and after the use of derivatives it was reduced to 52.19, a decrease of roughly 55%. As the government has allocated increasing funds for farmers to improve their draining abilities, one might argue that the candidate in Table 23 is more representative of the proposed derivatives` impact in the future.

As discussed earlier, even though reduced standard deviation is a desirable property when purchasing derivatives, it does not tell the full story. It is vital to consider how the average values of PPA move in relation to the movement of the standard deviation. After estimating the simplified Sharpe ratios in the backtesting dataset, the hedging strategy increases the SSR for as much as 95.5% of the organization numbers. The mean is a 16.9% increase in SSR, and the top decile ranges from 31.6% to 47.8%. At the other end of the scale, the greatest decrease in the simplified Sharpe ratio is just 5.2%. The results demonstrates that the strategy developed based on the original dataset, proves to be highly successful during backtesting.

The option hedging strategy proposed above was based on setting the strikes equal to the historical means of rainless days per month. This provides a straightforward basis for the option contracts and may serve as a suitable approach for individuals who are moderately risk tolerant. As farmers will have different attitudes towards risk, it could be of interest to examine the effectiveness of hedging when setting different strike levels. This will naturally change the corresponding prices, and for call options the prices will be higher the lower strike, as the expected payout increases. The attitude towards risk depends on multiple factors such as risk tolerance, financial goals and investment experience. Varying the strike values involves a trade-off: lower strike values come with a higher premium but offers a greater probability of the option being in the money, which aligns with the preference of low risk and more insurance against loss. Higher strike values result in lower premiums, but also lower probability of being in the money, and hence less insurance against loss. Ultimately, the choice of strike value depends on the investor`s preferred balance between cost, potential reward and their willingness to accept risk. The table below shows the prices of the call options in three scenarios: higher risk aversion with strike values two days below the means, moderate risk aversion with strike equal to the means (already presented in the results above), and lower risk aversion with strike values two days above the mean.

	Tick size	Maturity (in months)	**Price (higher risk aversion)**	**Price (moderate risk aversion)**	**Price (lower risk aversion)**
April	9,345	1	**31,097**	**17,593**	**8,345**
May	46,153	2	**163,981**	**98,705**	**52,135**
August	15,253	5	**58,356**	**37,299**	**21,706**

Table 24: Option prices for three different risk profiles

In the scenario of higher risk aversion, we find that the optimal strategy to the original dataset is to only purchase a call option for May with a tick size of 7.941. Utilizing this strategy on the backtesting data, there is an average decrease in standard deviation by 11.72%. As much as 97.8% of the organization numbers experienced a decrease in standard deviation. In terms of the simplified Sharpe ratio, the average increase of implementing the hedging strategy is 11.08%, with 96.6% experiencing an improvement. Even though the option prices are significantly higher than for the moderate risk scenario, the strategy still proves to be highly successful, as the payouts both increase and are more frequent.

In the lower risk aversion scenario, the optimal hedging strategy is to purchase call options for both April and May, with tick sizes of 5.135 and 33.749 respectively. Utilizing this strategy in the backtesting sample, there is an average decrease in standard deviation by 14.38%. As much as 96.6% of organization numbers experienced a decrease in standard deviation. And finally, in terms of the simplified Sharpe ratio, the average increase of implementing the hedging strategy is 13.27%, with 92.1% experiencing an improvement. Compared to the other two scenarios, the option prices are considerably cheaper, but on the other hand, payouts are more infrequent as the strike values are two and four days higher respectively. For May the tick size is roughly 4.25 times larger in the lower risk aversion scenario compared to the higher risk aversion scenario, which indicates larger payouts in extreme cases, but smaller payouts in less extreme cases due to differences in strike values. The results shows that the hedging strategy for all three risk profiles proves to be highly successful in the backtesting, further confirming the benefits of using weather derivatives for effective risk management.

In the literature section, it was found in a study by Marković and Jovanović (2011) that the winter barley production in Germany experienced a 40.42% reduction in standard deviation by utilizing rain put options. The effect is thus greater than in our analysis, but several factors may explain this outcome. For instance, our analysis focused on rainless days, while the study from Germany relied on accumulated rainfall. Additionally, Norway and Germany have very different climatic conditions, which leads to both different growing conditions and different soils, but also differences in pricing. Despite differences in approach and geographical context, the results from the German study strengthen and validate our own findings. Their analysis emphasizes and proves the advantage of utilizing weather derivatives, demonstrating that they can be an effective tool to reduce risk in the agriculture sector.

Conclusion

The objective of this thesis was to examine whether the use of precipitation derivatives could successfully mitigate some of the weather risk that are present in the agriculture sector. Through regression analysis, we were able to estimate how much an additional rainless day over the mean impacts crop yields. These estimates were then used to design the weather derivatives. The prices of the proposed derivatives differ very little between the historical burn approach and the Monte Carlo simulation, both when assuming a normal distribution and a Weibull distribution, which suggests that the prices are relatively robust. In the following step, three different hedging strategies were presented; a "naïve" approach setting tick sizes equal to the coefficients of the regression, a second approach that minimized the standard deviation, and a third approach that maximized the simplified Sharpe ratio. To balance the objective of minimizing the reduction in the mean PPA value while simultaneously lowering the standard deviation, we opted for the third approach. This optimization method resulted in 78% of organization numbers, in-sample, having an improvement in their simplified Sharpe ratio.

To check the robustness of the proposed hedging strategy, a backtest was conducted on a new dataset, containing the same organization numbers but a different time interval. The strategy resulted in a mean decrease in standard deviations by 16.4% for the scenario of moderate risk. Of all the organization numbers, 98.9% had their standard deviation reduced. For the simplified Sharpe ratio, similar results were found; an improvement was observed for 95.5% of organization numbers. The mean was a 16.9% increase in SSR, and the top decile ranges from 31.6% to 47.8%. In other words, the hedging strategy was very successful also out-of-sample. The results could be of interest to both farmers on an individual level, but also potentially the government. During years of adverse weather conditions, the government disburses hundreds of millions of kroner, and the proposed precipitation derivatives may help to stabilize these payouts which also fluctuate greatly.

For further studies, there are a few aspects that could have been improved or approached differently. As the weather derivatives require an independent third party to measure the precipitation, and most weather stations had insufficient or incomplete data, the basis risk is greater than desired. Two farms on opposite sides of the periphery of a county might experience slightly different weather conditions, which this analysis fails to account for. Even though we selected weather stations where the concentration of farms appeared to be the highest, it would be beneficial to include more weather stations with complete data. When subsequently modelling and thereby predicting rainless days by county, it might lead to somewhat inaccurate price estimates. The price of a weather derivative may, in practice, slightly differ within a county.

Another aspect has to do with climate change. As temperatures are rising all around the world, the optimal hedging strategy might also change in the future. Although our agricultural contact suggested that temperatures on its own has not destroyed a crop in Norway over the past 20 years, we still believe that temperature can be of importance in the future. The interplay of temperature and precipitation is likely to shift to some extent in the future. This interplay introduces an additional layer of high complexity to the optimal hedging strategy. It may be beneficial to adjust the precipitation derivative to account for temperature fluctuations to increase the accuracy of the hedge, but on another side, it would also involve substantial costs in terms of data collection, analysis and model development. With these challenges in mind, it is worth questioning whether the additional benefits of taking temperature into account could justify the increased time and financial investment required.

Declaration on the use of AI tools in the work on this master's thesis

Name of the AI tool:

ChatGPT 3.5/4.0

Purpose of using the tool:

- Used as a tool for generating R-code in some contexts and assisting with debugging error messages
- Also applied for language improvements through providing synonyms/word replacements and ensuring better text flow

We are aware that we are responsible for all content of this master's thesis, including the parts where AI tools are used. We are responsible for ensuring that the thesis complies with ethical rules for privacy and publication.

Bibliography

Bobriková, M. (2022). *Weather risk management in agriculture using weather derivatives*. Italian Review of Agricultural Economics, 77(2), 15–26. https://www.researchgate.net/publication/307526869_Weather_Risk_Management_in_Agriculture

Børresen, M. F., Mo, A.-K., Stenberg, M., & Sveen, J. (2023, September 18). *Elendig kornavling – begynte med tørke og sluttet med flom.* Retrieved from NRK: https://www.nrk.no/innlandet/arets-kornavling-blir-liten-og-av-darlig-kvalitet-_-forst-torke-og-sa-mye-nedbor-og-flom-1.16558941

Cao, M., & Wei, J. (2004). *Weather Derivatives Valuation and Market Price of Weather Risk.* Journal of Futures Markets, 24(11), 1065-1089. https://doi.org/10.1002/fut.20122

Castano, I. (2024, September 2). *Weather Derivatives Grow as Risks Intensify*. Retrieved from OpenMarkets: https://www.cmegroup.com/openmarkets/energy/2024/Weather-Derivatives-Grow-as-Risks-Intensify.html

Cueto, J. (2024, June 25). *Catastrophe bonds gain popularity among retail investors*. Retrieved from ReInsurance Business: https://www.insurancebusinessmag.com/asia/news/reinsurance/catastrophe-bonds-gain-popularity-among-retail-investors-494460.aspx

Evans, S. (2019, May 22). *Nephila worked with Allianz on GrainCorp weather risk hedging*. Retrieved from Artemis: https://www.artemis.bm/news/nephila-worked-with-allianz-on-graincorp-weather-risk-hedging/

Felleskjøpet. (2024). *Pris og prognoser*. Retrieved from Felleskjøpet: https://www.felleskjopet.no/korn/pris-og-prognoser/

Fernando, J. (2024, August 4). *Sharpe ratio: Definition, Formula, and Examples*. Retrieved from Investopedia: https://www.investopedia.com/terms/s/sharperatio.asp

Finborud Børresen, M., Wold Haagensen, V., & Meskau, F. (2018, August 18). *Tre av fire bønder er ikke forsikret mot avlingssvikt*. Retrieved from NRK: https://www.nrk.no/innlandet/tre-av-fire-bonder-er-ikke-forsikret-mot-avlingssvikt-1.14169786

Frøseth, R. B. (2011, July 14). *Vekstskifte i økologisk korndyrking*. Retrieved from Agropub: https://www.agropub.no/fagartikler/vekstskifte-i-okologisk-korndyrking

Hayes, A. (2020, July 31). *Combined Ratio: Definition, What It Measures, Formula, Examples.* Retrieved from Investopedia: https://www.investopedia.com/terms/c/combinedratio.asp

Holtet, E. K. (2024, October 21). *Bygg.* Retrieved from Store Norske Leksikon: https://snl.no/bygg

Härdle, W. K., & Cabrera, B. L. (2009). *The Implied Market Price of Weather Risk*. Applied Mathematical Finance, 19(1), 59-95. https://doi.org/10.1080/1350486x.2011.591170

Jewson, S., & Brix, A. (2005). *Weather Derivative Valuation: The Meteorological, Statistical, Financial and Mathematical Foundations.* New York: Cambridge University Press.

Kenton, W. (2024, June 27). *Monte Carlo Simulation: What It Is, How It Works, History, 4 Key Steps*. Retrieved from Investopedia: https://www.investopedia.com/terms/m/montecarlosimulation.asp

Kenton, W. (2023, June 28). *Goodness-Of-Fit*. Retrieved from Investopedia: https://www.investopedia.com/terms/g/goodness-of-fit.asp

Landbruksdirektoratet. (2024). *Norsk landbruk - tall og fakta.* Retrieved from Landbruksdirektoratet: https://www.landbruksdirektoratet.no/nb/norsk-landbruk-tall-og-fakta

Landbruksdirektoratet. (2024). *Produksjonstilskudd og avløsertilskudd i jordbruket.* Retrieved from Landbruksdirektoratet: https://www.landbruksdirektoratet.no/nb/jordbruk/ordninger-for-jordbruk/produksjonstilskudd-og-avlosertilskudd-i-jordbruket

Landbruksdirektoratet. (2024). *Tilskudd ved produksjonssvikt.* Retrieved from Landbruksdirektoratet: https://www.landbruksdirektoratet.no/nb/jordbruk/ordninger-for-jordbruk/tilskudd-ved-produksjonssvikt?openStep=87af4add-5b11-408f-8268-cb8708683470-3

Langerud, A. (2023, March 20). *Vår- og høstkorn.* Retrieved from NDLA: https://ndla.no/nb/subject:57f73fac-aef5-4ab9-9ce0-55f031c24768/topic:d59c1ae7-d9a3-4c0f-b3ce-d8aac53d5385/topic:ae907e7d-89d7-401e-894d-3a7e832ff8cd/resource:d2e541ad-0c6d-45d5-842a-5b32a581d77e

Langan, P., Bernád, V., Walsh, J., Henchy, J., Khodaeiaminjan, M., Mangina, E., & Negrão, S. (2022). *Phenotyping for waterlogging tolerance in crops: current trends and future prospects*. Journal of Experimental Botany, 73(15), 5150-5153. https://doi.org/10.1093/jxb/erac243

Mustafa, A. (2023, July 8). *Understanding Random Effects and Fixed Effects in Statistical Analysis*. Retrieved from Medium: https://medium.com/@akif.iips/understanding-random-effect-and-fixed-effect-in-statistical-analysis-db4983cdf8b1

Nickolas, S. (2024, August 19). *Understanding the Difference Between Moral Hazard and Adverse Selection*. Retrieved from Investopedia: https://www.investopedia.com/ask/answers/042415/what-difference-between-moral-hazard-and-adverse-selection.asp

Niedens, L. (2023, October 18). *Listed weather deriviatives grow in popularity but have limitations.* Retrieved from Investopedia: https://www.investopedia.com/listed-weather-derivatives-grow-in-popularity-but-have-limitations-8362112

PWC & FFN. (2023). *Risikopremien i det norske markedet.* Retrieved from PWC: https://www.pwc.no/no/publikasjoner/risikopremien-2023.pdf

Regjeringen. (2024). *Jordbruksoppgjøret.* Retrieved from Regjeringen: https://www.regjeringen.no/no/tema/mat-fiske-og-landbruk/jordbruk/innsikt/jordbruksoppgjoret/id2354584/

Robertson, H. (2023, October 12). *Use of weather derivatives surges as extreme climate events rock the globe*. Retrieved from Reuters: https://www.reuters.com/markets/global-markets-weather-derivatives-analysis-pix-2023-10-11/

Roll, R. (1984). *Orange Juice and Weather.* American Economic Association, 74(5). https://econpapers.repec.org/article/aeaaecrev/v_3a74_3ay_3a1984_3ai_3a5_3ap_3a861-80.htm

Røkkum, K. S., Myrvang, L. B. (2019, May 23). *The Use of Wind Speed Derivatives for a Norwegian Energy Producer.* Retrieved from NTNU Open: https://ntnuopen.ntnu.no/ntnu-xmlui/bitstream/handle/11250/2617919/no.ntnu%3Ainspera%3A2359025.pdf?sequence=1&isAllowed=y

Samarah, H. N. (2005). *Effects of drought stress on growth and yield of barley*. Agronomy for Sustainable Development, 25(1), 145–149. https://doi.org/10.1051/agro:2004064

Samarah, N. H., Alqudah, A. M., Amayreh, J. A., & McAndrews, G. M. (2009). *The effect of late-terminal drought stress on yield components of four barley cultivars*. Journal of Agronomy and Crop Science, 195(6), 427–441. https://doi.org/10.1111/j.1439-037X.2009.00387.x

SSB. (2017). *Gjennomsnittlig utlåns- og innskuddsrente i bankene, prisstigning, beregnet realrente, effektiv rente for innenlandske obligasjoner (10 års) og Nibor (3 md.) per 31. desember. Prosent (SÅ 454).* Retrieved from Statistisk sentralbyrå: https://www.ssb.no/300692/gjennomsnittlig-utlans-og-innskuddsrente-i-bankene.per-31.desember.prosent-sa-454

SSB. (2024, February 19). *Korn og oljevekster, areal og avlinger.* Retrieved from Statistisk sentralbyrå: https://www.ssb.no/jord-skog-jakt-og-fiskeri/jordbruk/statistikk/korn-og-oljevekster-areal-og-avlinger

Štulec, I., Petljak, K., & Baković, T. (2015). *Effectiveness of weather derivatives as a hedge against weather risk in agriculture*. Agricultural Economics (AGRICECON), 188(8). https://doi.org/10.17221/188/2015-AGRICECON

Vosgraff, S. K. (2024, October 22). *Gjensidige doblet bunnlinjen*. Retrieved from Finansavisen: https://www.finansavisen.no/finans/2024/10/22/8194513/gjensidige-doblet-bunnlinjen