Forecasting Cocoa Futures: Neural Network Using Weather Data

Keith Openshaw

9/7/24

Abstract

This paper explores the creation of a predictive model to forecast cocoa futures prices using a neural network. The model combines historical cocoa price data and real-time weather information to estimate the end-of-month price for cocoa futures. The goal is to develop a data-driven trading strategy that capitalizes on short-term price movements in the volatile cocoa market. Key inputs include weather data from the National Oceanic and Atmospheric Administration (NOAA) and cocoa futures prices from the Intercontinental Exchange (ICE). The model was tested and validated through a backtesting process to assess its effectiveness in predicting price trends and guiding trades. This research demonstrates the potential of using machine learning to improve trading strategies in commodity markets like cocoa.

1 Introduction

Cocoa futures have become increasingly volatile, with the 2024 price spike driven by extreme weather and diseases affecting production in West Africa. Given that the region supplies over 70\% of the world's cocoa, understanding whether such spikes can be predicted using weather data is crucial for traders.

This research develops a neural network model that integrates real-time weather data from NOAA with historical cocoa prices from the Intercontinental Exchange (ICE) to predict end-of-month cocoa prices. By focusing on weather's influence on cocoa yields, the model seeks to determine whether it could have predicted the 2024 price surge and improved trading decisions in the cocoa futures market.

Key question

Could this price surge have been predicted using weather data?

To answer this question, a model was developed that uses historical weather data to predict the value of Cocoa bean futures next month.

$$E(price) = f(x_1, x_2, \dots)$$

2 Methodology

2.1 Data and Model Inputs

To predict cocoa prices, the model uses 5 years of data relies on two primary sources:

NOAA Global Weather Data

This is a global dataset of satellite and weather station data, including atmospheric conditions like temperature and precipitation—available upon request from the National Centers for Environmental Information (NCEI) online portal.

The model is trained on weather data from two countries in Western Africa: **Ghana** and **The Ivory Coast**, which together make up ~60% of the worlds cocoa production. NOAA data consisted of daily precipitation and temperature readings, including averages and extremes.

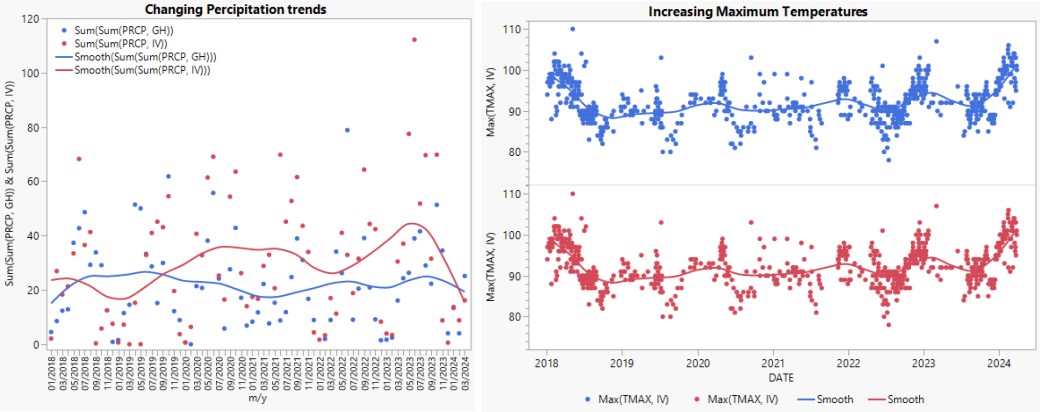

Cocoa Price Data

Cocoa Bean Futures are traded on the Intercontinental Exchange (ICE) in two locations: London (ICEU) and New York (NYCC), and there are five contracts a year.

This model is created using Cocoa **front-month** contracts (CC.1), which quotes the price of the contract closest to expiration. These offer a consistent quote for the price of cocoa and are typically the most actively traded contracts with the most liquidity available.

2.2 Model Creation

JMP 17 was used to prepare the data and perform the analysis. The predictive model was built using a neural network to capture complex, non-linear relationships between weather conditions and cocoa prices, aiming to forecast prices one month ahead.

Preprocessing

Data preprocessing included normalization of both historical cocoa prices (CC.1) and NOAA weather data to ensure consistency in scale. Missing weather data points were imputed using surrounding averages, and all inputs were synchronized on a monthly basis to align with the model's prediction target.

Inputs and Outputs

- Inputs: Historical Cocoa Prices (CC.1 front-month contracts), NOAA Weather Data (daily temperature and precipitation from Ghana and Ivory Coast).

- Output: Predicted end-of-month CC.1 price.

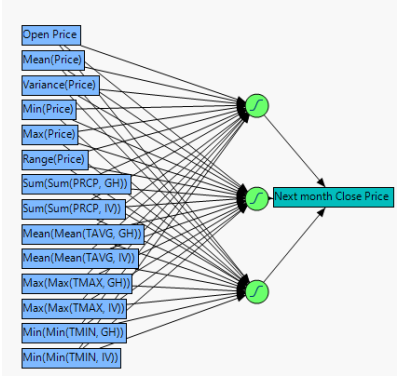

Neural Network Architecture:

The model consists of three hidden layers, designed to capture both simple and complex patterns in the data. The dataset was split into 66% for training and 33% for validation to ensure model robustness and prevent overfitting.

2.3 Backtesting Framework

The neural network programming in JMP provides built in metrics on model accuracy and significance. However, it is necessary to implement a backtesting strategy to assess whether the model's predictions actually have any trading potential.

The backtesting uses simple conditional logic to compare the expected future price of Cocoa (model's prediction) with the current price. It then chooses to "buy" or "short" depending on where it predicts the price to move.

$$\begin{cases} buy, & p_{expected} > p_{actual} \\ short, & p_{expected} < p_{actual} \end{cases}$$

Returns are then aggregated using the positive return on "buy" months and the negative of returns on "short' months.

$$profit = \begin{cases} p_{next\ close} - p_{open}, & buy \\ p_{open} - p_{next\ close}, & short \end{cases}$$

Special Considerations

Futures contracts often require special considerations when tied to physical assets. When Cocoa futures contracts expire, the owner of Cocoa contracts at expiration must take physical delivery of the underlying asset. Cocoa contracts expire on the 16th or 17th of the corresponding month, therefore, to avoid taking delivery: this strategy closes all positions on the 15th.

Backtesting Summary:

- Trades were opened on the 18th of each month and closed on the 15th of the following month to avoid physical delivery
- Trades were done at market close
- Only front-month contracts were traded
- Returns were calculated based on the long or short return over contract holding period

3 Results

3.1 Model Predictions

The model showed strong accuracy in predicting price movements, including the surge seen by 2024.

It performs well on both Training and Validation sets, with high R^2 values and lower error metrics on the validation set (RASE and MAD). It preforms especially well at recognizing the rapid price increase in 2023 and adapting to higher valuations. However, it underestimated the nearly 100% price surge in 2024.

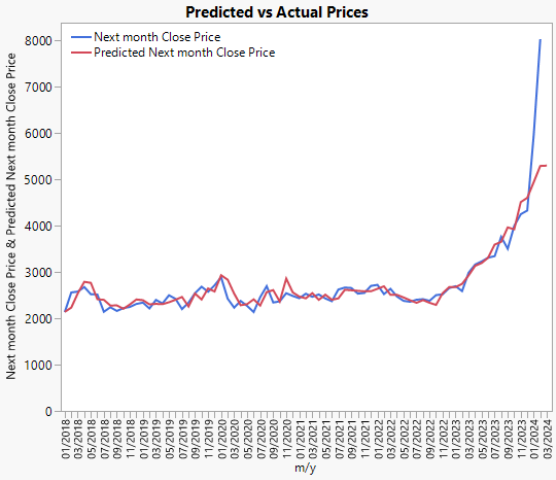

$$f(\dots) = p_{expected}$$

Model NTanH(3)

Training		Validation	
Next month Close Price		**Next month Close Price**	
Measures	**Value**	**Measures**	**Value**
RSquare	0.8056907	RSquare	0.8195397
RASE	450.88533	RASE	170.63101
Mean Abs Dev	192.21802	Mean Abs Dev	120.26734
-LogLikelihood	353.91712	-LogLikelihood	157.40261
SSE	95549864	SSE	698758.62
Sum Freq	47	Sum Freq	24

It's predicted accuracy is impressive and the decrease in RASE and MAD from training to validation suggests the model generalizes well and with a low chance of overfitting.

3.2 Backtesting Results

The model also showed impressive performance when applied to the backtesting framework.

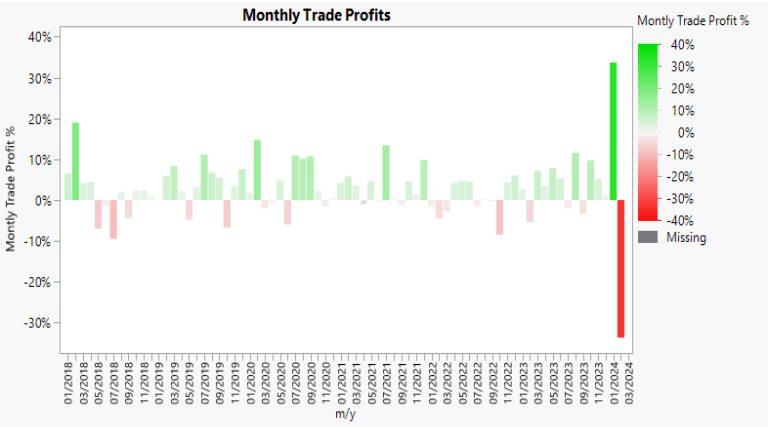

Expected Return (Monthly)	**2.74%**
Annualized	**27.39%**

Correct Trades:	**50**
Incorrect:	**24**
Ratio	**67.6%**

Var	0.00614
Standard Deviation	0.078
Skew	-1.122
Kurtosis	3.148
Min Return	-33.86%
Max Return	18.83%

It correctly predicted the directions of monthly price movements **67.6%** of the time and had a monthly expected return of **2.74%**.

The "Monthly Trade Profits" chart shows how overwhelmingly positive these trades are and demonstrates the consequences of the 2024 price underestimate—which resulted in a 34% loss. The implications and analysis of this will be discussed later.

3.3 Comparisons

The model's performance was compared against the S&P 500 and a 'buy only' cocoa bean strategy to assess competitive opportunity.

S&P 500

The model significantly outperformed the S&P 500 over the same period, generating an Annualized Excess Return of **17.05%**. The profits from the model are represented in the following graph and table as "Monthly Trade Profit%".

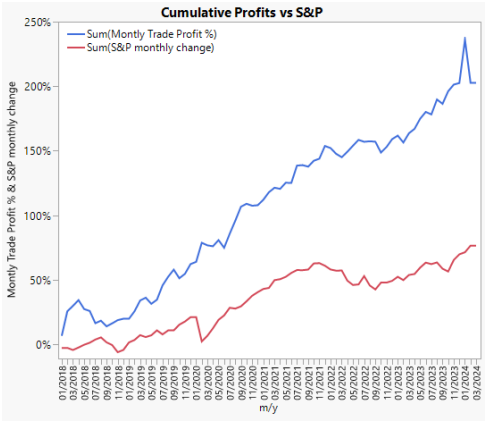

	Montly Trade Profit %	S&P monthly change	Excess Return
Mean	2.739%	1.034%	1.705%
Std Dev	0.0784174475	0.040789945	0.0941664672
Min	-33.66%	-18.92%	-38.83%
Max	33.71%	8.92%	33.71%

Unsurprisingly, the model has a standard deviation roughly 2x as large as the S&P and has significantly larger maximum and minimum returns.

Buy Only Cocoa Strategy

CC1 contracts have been an incredibly profitable investment due to the discussed price surge, so it is important to analyze whether our strategy is better then simply buying and holding these contracts for the same monthly period. This also addresses whether the "short" feature of our backtesting model contributes to performance.

A 'Buy Only' strategy that purchases CC1 contracts every month and holds for the same period as the model generated an annualized return of 18.87%. While better than the S&P, this is still 8.52% less than the return of our 'Buy and Short' strategy.

4 Discussion

The model predicted monthly cocoa prices extremely well during periods of high volatility, and successfully predicted the 2023/2024 surge—the goal for which it was originally designed. It boasts an impressive 27% annualized return, significantly outperforming buy-only contracts and the S&P 500 over the same period. All metrics discussed clearly demonstrate strong potential for profitable trading using this trading strategy.

The model's impressive performance highlights a strong correlation between weather fluctuations in West Africa and cocoa bean futures prices, confirming the initial hypothesis behind its development. This finding suggests that changes in weather patterns have a direct influence on cocoa supply, driving futures prices.

4.1 Limitations

Our model was able to draw a direct predictive link between weather data and the price of cocoa contracts and performed very well in the backtesting framework. However, it is extremely difficult to know how this strategy would perform in a real trading environment. This section illustrates the key risks and limitations that could challenge the profitability of this strategy.

- **New Valuations** - Perhaps the greatest limitation of this model relates to the price surge it intended to predict. The model was successful in predicting this surge, however cocoa price valuations have

never been the same since 2022. Prices have remained higher, with no sign of coming down. This suggests that investors and commodities traders have changed their long-term assumptions and predictions for the future of the cocoa market. The model was trained on cocoa prices before this price surge, and it's possible that valuations have changed enough to make the model obsolete. This concern is reflected by the last month of the model's predictions—prices were so high that the model expected a decline regardless of the weather data. It is impossible to know whether the price surge has made the model obsolete; only further testing and time will be able to answer this question.

- **Competing Traders.** Sophisticated commodities traders are incredibly intelligent, and it is likely they have similar predictive models—or at least the capability to make them. If competing traders make similar or more accurate models, the proprietary advantage of this model would disappear. Historical prices do not clearly indicate that such traders already exist, therefore it is possible that this application is niche enough to remain untested.
- **Weather Data Availability.** This problem would become prominent when trying to operate this strategy at scale competitively. In order to obtain the training weather data for this model, a request had to be submitted with the NOAA discussing the exact data requested and its purposes, and it took multiple days to confirm and receive the data. This would be impractical for more sophisticated models as new weather data would have to be integrated for predictions constantly. To operate this strategy optimally, traders would need a new source of data or special relationship with NOAA data distribution.

None of these concerns were evident enough from backtesting to doubt the profitability of the strategy. However, it is important to acknowledge these challenges and account for them in any real-world application of the cocoa trading strategy.

4.2 Optimization and Improvements

The model showed impressive predictive power despite its relative simplicity. It predicts prices mainly as a function of precipitation and temperature, but there are a number of factors that influence cocoa prices that the model is not designed to consider.

The following discusses improvements and additions to the model that could increase predictive power and the feasibility of this trading strategy.

- **Climate change models.** The model doesn't account for long-term climate change and its impact on cocoa production. Incorporating long-term climate projections into the model would allow it to anticipate long term shifts in production and provide significant opportunities for trading. It's possible this data would help explain the fundamental shift in valuations that has occurred in the last few years.
- **Supply factors other than weather.** Beyond weather, diseases and political instability can disrupt cocoa supply. In the last few years flooding has caused black pod disease and swollen-shoot virus to limit the cocoa yield, and constant labor disputes with farmers have driven the production costs higher independent of weather. Adding variables such as these would make the model more responsive to sudden supply shocks and long-term geopolitical trends.
- **Demand side factors.** The model focuses heavily on weather-driven supply but ignores demand-side influences, such as consumer trends and cocoa grindings data. Demand can shift significantly based on market trends, affecting prices independently of supply. By integrating these factors, the model would offer a more comprehensive prediction and better capture market dynamics.

- **Optimizing holding periods.** The current strategy always opens trades on the 18th of the month and exits on the next 15th to avoid physical delivery. This was chosen to generalize the strategy for any month or year; however, it likely misses profitable opportunities and trades at inopportune times. Future models should optimize trades, so they occur on the best day possible, and held for the best amount of time depending on the month and time to expiry.
- **Optimizing Contract Selection.** This model relies solely on front-month contracts (CC.1), which are the most liquid and closest to expiration. By incorporating longer-dated contracts such as second- or third-month futures the model could capture more seasonal patterns and potential offer more trading opportunities. This would be especially promising combined with longer term price determinants such as climate change or geopolitical turmoil.

5 Conclusion

This study presents a neural network model that successfully predicts cocoa futures prices using weather data from West Africa. The model demonstrated strong predictive power, notably during periods of volatility, including the 2023 and 2024 price surge. It achieved an annualized return of 27%, outperforming both a "buy-only" cocoa strategy and the S&P 500, while accurately predicting price movements 67.6% of the time.

The results underscore the value of weather data in predicting commodity prices, particularly for crops like cocoa, where production is deeply affected by environmental conditions. The model's ability to capitalize on weather-driven price fluctuations highlights its potential as a tool for traders in volatile markets.

Key Limitations and Opportunities for Future Work

While the model performed impressively, it faces several limitations. One of the most significant challenges is the structural shift in cocoa prices following the 2024 surge, raising concerns about the model's ongoing relevance. Additionally, the logistical challenges of acquiring real-time weather data present operational hurdles, and there is a risk that competing models developed by sophisticated traders could erode the model's competitive advantage.

To enhance the model's robustness and adaptability, future improvements should focus on:

- Adapting to Market Changes: Incorporating long-term climate change projections and geopolitical factors could help address new pricing dynamics in the cocoa market, making the model more responsive to shifts in supply and demand.
- Expanding Data Inputs: Introducing demand-side variables like consumer trends and cocoa grindings, as well as integrating other supply-side disruptions such as disease outbreaks, could provide a more comprehensive prediction of price movements.
- Strategic Enhancements: Optimizing contract selection and holding periods could increase profitability by capitalizing on seasonal patterns and better aligning trades with market conditions.

The current model offers a strong foundation for trading cocoa futures using the predictive model; however, adapting to market evolution and expanding the scope of input data will be key to sustaining its performance in an increasingly competitive environment.