# Enhancing Futures Trading Signals with Kalman Filtering

Nicholas Farr

AlgoGators Capstone Project

## Purpose of Research:

To determine whether Kalman filtering can improve the clarity and predictive power of short-term trading signals, enabling AlgoGators to execute more precise and profitable futures trades.

## Introduction:

Originally developed by Rudolf E. Kalman in 1960, the Kalman Filter has found applications in numerous fields, including engineering, navigation, robotics, and finance. In finance, it is particularly useful for predicting time series data, such as stock prices or futures contracts, which are influenced by a combination of trends and random fluctuations. By applying the Kalman Filter, traders can estimate the true value of a financial instrument, smooth out short-term volatility, and make more informed trading decisions.

This research explores the effectiveness of Kalman filtering in forecasting short-term directional movements and precise trade entry points in futures markets. Precision, not speed, is the primary focus—aligning with AlgoGators strategy to trade algorithmically without competing directly with high-frequency traders. The hypothesis is that applying these algorithms will improve trading accuracy, enabling AlgoGators to identify profitable market moves with greater confidence and precision in comparison to the current EWMA signal generation used by AlgoGators.

## Methodology: Building a Kalman Filter

***Note:*** *All Model Implementation was built and tested in python.*

Kalman Filtering:

The Kalman Filter is a recursive state-space estimation algorithm used to infer the underlying state of a system from noisy observation. In the context of financial time series, it is particularly effective for smoothing and predicting asset prices by filtering out market noise.

The Kalman Filter operates in two alternating steps:

**1. Prediction:** Uses a mathematical model to project the next state (e.g., future price) based on current estimates.

**2. Update (Correction):** Incorporates the actual observation (e.g., the true closing price) to correct the prediction, balancing the trust between the models forecast and new observations

## Implementation in Python:

**State Transition Matrix:**

```python
# State transition model — how we believe the system evolves
F = np.array([
    [1, dt],  # price = price + velocity * dt
    [0, 1]    # velocity stays constant
])
```

Represents how we expect the hidden state to evolve over time. The first row represents that tomorrow's price = today's price + today's velocity*(dt) - where DT is 1 due to having daily ticks. If this model were to have intraday ticks, this filter could update throughout the day.

**Measurement Model:**

```
# Measurement model — we only observe price, not velocity
H = np.array([[1, 0]])
```

This array represents what we can actually observe from the ES futures Market Data. I calculated them based on closing price, and there are no other observations such as velocity. Therefore since velocity isn't directly measurable we extract only the price (internal state) to what we observe.

**Initial Conditions:**

```
# Initial state estimate: [initial_price, initial_velocity]
x = np.array([df["Close_ES=F"].iloc[0], 0.0])

# Initial uncertainty in our estimate (very uncertain at start)
P = np.diag([1000.0, 1000.0])
```

At initialization, the Kalman Filter has access to only a single data point—the first observed closing price of the ES futures contract. With no information on movement yet, we assume the initial velocity (i.e., price change) to be zero. To reflect our lack of confidence at this early stage, the initial error covariance matrix P is set to a large value. This high uncertainty allows the filter to quickly adapt as more data becomes available and refine its internal state over time.

**Model Uncertainty:**

```
# How much we think price and velocity can randomly change
process_var_price = 1.0
process_var_velocity = 0.1
Q = np.diag([process_var_price, process_var_velocity])
```

These parameters are set to filter some of the market noise out. A value of 1 for process can be interpreted as the closing price can fluctuate $1 each day with little meaning to it ~ random fluctuations in price without signifying any trend or meaning behind the data. The same logic applies for the variance in velocity set as 0.1, suggesting that momentum is relatively stable unless observations force an adjustment.

**Measurement Noise:**

```
# How noisy our observations are — estimated from real market data
measurement_var = np.var(df["Close_ES=F"].diff().dropna())
R = np.array([[measurement_var]])
```

R measures the noise covariance, which can be interpreted as a quantification of how much the filter trusts the incoming data point. It is calculated as the variance of how much the ES futures close price fluctuates day to day. In periods of higher volatility this value is higher, becoming more cautious in overreacting to new prices. In periods of low volatility, this value will increase the weight of the new observation.

**Kalman Loop Logic (Training and Building the Filter):**

```python
kf_prices = []
kf_velocity = []

for z in df["Close_ES=F"].values:
    # --- 1. Predict step ---
    x = F @ x                    # Predict next state (price + velocity)
    P = F @ P @ F.T + Q          # Predict uncertainty

    # --- 2. Update step ---
    y = z - (H @ x)              # Innovation = observed - predicted
    S = H @ P @ H.T + R          # Total uncertainty of the innovation
    K = P @ H.T @ np.linalg.inv(S)  # Kalman Gain (blending factor)

    x = x + (K.flatten() * y)    # Update the state estimate
    P = (np.eye(2) - K @ H) @ P  # Update the uncertainty

    # --- 3. Store results ---
    kf_prices.append(x[0])       # Smoothed price
    kf_velocity.append(x[1])     # Estimated velocity (momentum)
```

*Note: @ is the matrix multiplication operator (the dot product)*

The Kalman loop iterates over each ES Closing price and continues to update with every new closing price ~ 3 years of market data. The loop first performs a prediction step by projecting the next state estimate using the state transition matrix (F). The error covariance matrix (P) is simultaneously updated to reflect growing uncertainty using the process noise (Q). In the update step, the algorithm computes the innovation – defined as the difference between the actual observed price and the predicted price – then evaluates its uncertainty through the innovation covariance S. Using this, the Kalman Gain (K) is derived to determine how much the prediction should be corrected based on the new data. The state and uncertainty are then updated accordingly. Finally, the filtered (smoothed) price and estimated momentum are stored for each day, enabling downstream comparison against other forecasting models such as EWMA.

## Evaluation Strategy:

Currently, AlgoGators utilizes Estimated Weighted Moving Averages (EWMA) as a way to generate signals and find trends. This filter can be robust, allowing for a lot of optimization and control within the trend following strategy, as well as can be crossed with different spans to enhance signal accuracy. The Kalman filter is inherently less robust, however, can provide huge benefit when modeling latent momentum and updating adaptively giving a statistical advantage in non-static environments like futures markets.

To evaluate the Kalman Filter's forecasting performance, we used its Smooth Price generation and tested it against a variety of metrics. Calculating metrics such as Sharpe Ratio, or anything dealing with returns would require implementing a viable trading strategy that potentially longs, shorts, and holds positions based on signals generated from the Kalman Filter. However, the scope of this research is to identify how accurate the produced signals are, rather than how effective a strategy with them can be. To identify accuracy of modeling the ES Futures Time Series data a variety of metrics were calculated.

- **MAE (Mean Absolute Error): Measures average magnitude of forecast errors.**

$$MAE = \frac{1}{n} \sum_{i=1}^{n} \left| Y_i - \hat{Y}_i \right|$$

- **RMSE (Root Mean Squared Error): Penalizes larger errors more heavily.**

$$RMSE = \sqrt{\sum_{i=1}^{n} \frac{(\hat{y}_i - y_i)^2}{n}}$$

- **MAPE (Mean Absolute Percentage): Expresses error as a percentage of actual price.**

$$\text{MAPE} = \frac{1}{n} \sum_{i=1}^{n} \left| \frac{Y_i - \hat{Y}_i}{Y_i} \right|$$

These metrics provide a comparison of how accurate the Kalman prediction is at signaling and forecasting on futures markets (particularly CME E-mini S&P 500 Futures) in comparison to two EWMA spans (16,64) and (64, 256). For each EWMAC pair, the “slow” EWMA was used as a 1-day-ahead forecast by shifting the series forward, simulating a real-time predictive environment.

## Results:

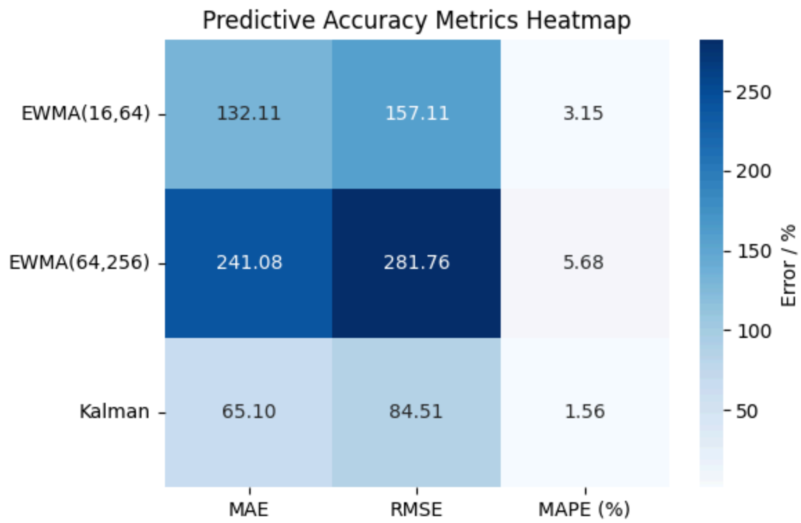

*Figure 1: Heatmap Accuracy Metrics*

Visualized in the Figure Above, it is clear that a Kalman Filter outperforms both of the EWMA spans in all three standard error metrics. It achieved a Mean Absolute Error (MAE) of 65.10, compared to 132.11 and 241.08 for the shorter and longer EWMA spans, respectively. The Root Mean Squared Error (RMSE) shows a similar trend, with the Kalman Filter outperforming both EWMAs by more than 50%. Finally, in terms of Mean Absolute Percentage Error (MAPE)—which normalizes error relative to price—the Kalman Filter achieved just 1.56%, compared to 3.15% and 5.68% for the EWMAs. These results confirm that the Kalman Filter not only produces more stable forecasts but also more accurate directional and magnitude predictions for ES futures prices.

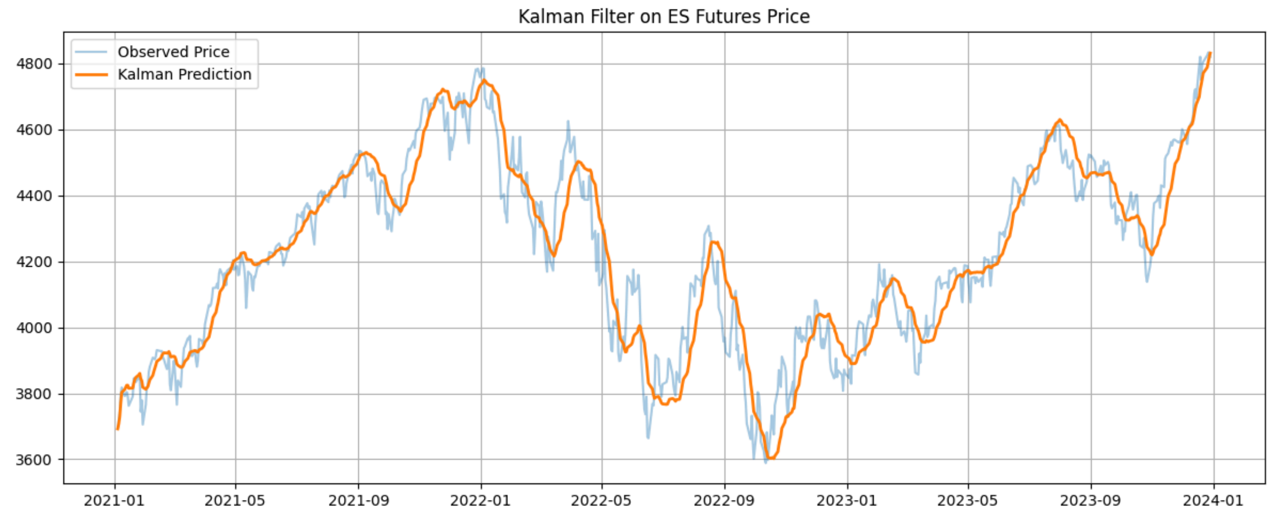


*Figure 2: Kalman Prediction on ES Futures Price*

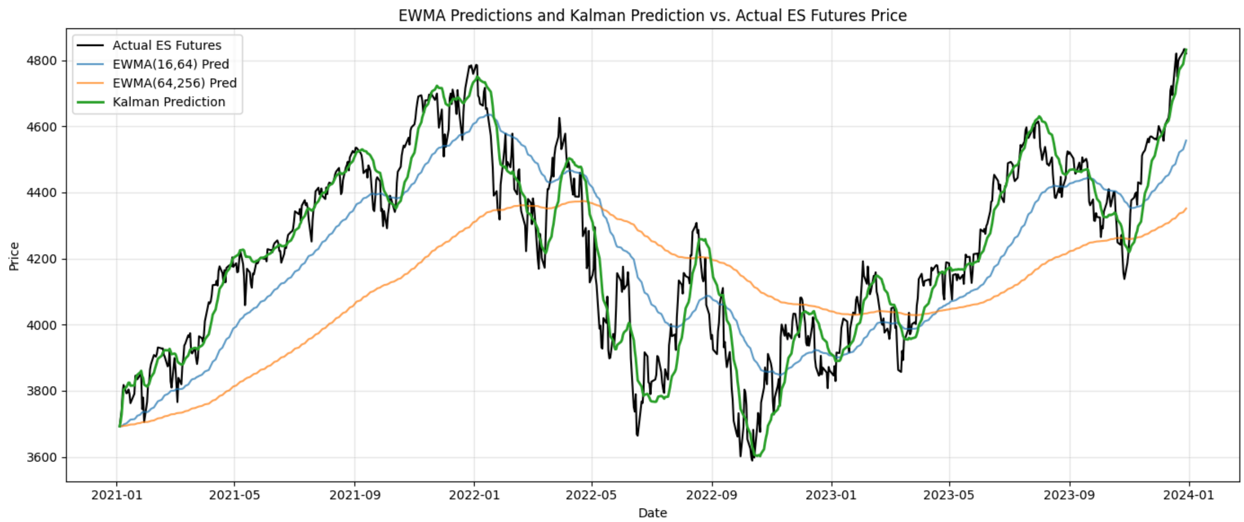


*Figure 3: Overlay of All 3 Signal Forecasting Methods vs ES Futures*

Figures 2 and 3 illustrate how each forecasting model tracks the ES futures prices over the 3 year span. The Kalman filter outputs the prediction with minimal lag, in comparison both of the EWMA spans have a significant amount of lag in their price prediction. The filter is especially effective during periods of volatility—smoothing noise without overshooting major turning points—demonstrating its value as a dynamic state estimator for time series data. These visualizations show how effective the Kalman filter can be when attempting to model time series data

## Further Testing and Experimentation:

This research proved that the Kalman Filter is more dynamic and fits non static market data better than a traditional EWMA filter. However it is important to note that the EWMA models are not traditionally used to predict precise price levels. Rather they are strong at identifying overall direction and trend rather than daily price values.

In addition, future research could explore how this model performs under regime switching conditions, as well as how its velocity or change in market price can be leveraged for dynamic position sizing. Kalman filtering can be further optimized in various ways potentially resulting in further performance and financial gains.

## Trading System Implementation:

Another limitation of this experiment was the lack of implementing a trading system. As stated above, EWMA filters are not particularly known for predicting the next day's price, rather they are crossed against different time spans and then used to generate a predicted overall trend. Setting parameters and thresholds to when you should enter a short/long position given the strength of this filter's signal can still result in profitable trading despite their limited accuracy when predicting exact price levels. Valuable future research would be to implement a standardized trading strategy using both of these filters and compare metrics such as Sharpe Ratio.

## Conclusion:

Regardless of the theoretical difference between the Kalman Filtering and EWMA span filters, the Kalman smoothed price predicting seems to perform highly and is well suited to predict ES futures time series data. The Kalman Filter's ability to dynamically update both price and momentum states when provided new information gives the filter an edge in volatile and non-static data sets.

## Works Cited:


Iarslan, S. (2024, May 25). Implementing a Kalman filter-based trading strategy. Medium. https://medium.com/@serdarilarslan/implementing-a-kalman-filter-based-trading-strategy-8dec764d738e

Qpy. (2024, May 25). Implementing a Kalman filter-based trading strategy. Implementing a Kalman Filter-Based Trading Strategy. https://quantitativepy.substack.com/p/implementing-a-kalman-filter-based

Research/analysis/02 kalman filter based pairs trading.ipynb at ... (n.d.-a). https://github.com/QuantConnect/Research/blob/master/Analysis/02 Kalman Filter Based Pairs Trading.ipynb

ZoroDareVohnCoder. (2024, June 30). Using the unscented Kalman filter on stock prices. https://www.signalpop.com/2024/06/30/using-the-unscented-kalman-filter-on-stock-prices/

# Disclaimer

The information, trading strategies, and materials presented in this report are provided strictly for educational and informational purposes and are offered without any guarantees or warranties regarding their accuracy, completeness, reliability, or timeliness. These materials are not intended to serve as financial, legal, tax, investment, or other professional advice, and no content herein should be interpreted as a recommendation to buy, sell, or hold any security, financial product, or instrument, nor as an endorsement of any specific strategy, practice, or course of action. Users of this material are strongly encouraged to conduct their own independent research, analysis, and due diligence or seek advice from qualified professionals before making any financial or investment decisions.

The authors, contributors, and distributors of this material are not acting as fiduciaries and do not assume any legal or ethical duty to the user. No advisory or client relationship is created through the use of this material. Trading and investing, particularly in derivatives, options, futures, and other leveraged products, involve substantial risks. These risks may include, but are not limited to, the potential for significant financial losses, the loss of principal, and exposure to unpredictable market conditions, volatility, or adverse economic factors. Users should be aware that past performance is not indicative of future results, and reliance on historical data or strategies described herein may not lead to favorable outcomes. Additionally, users are solely responsible for ensuring their trading or investment activities comply with applicable laws, regulations, and tax obligations in their jurisdiction.

By accessing and using this material, users acknowledge and agree to the following terms:

1. Assumption of Risk: Users accept all responsibility for the risks inherent in any trading or investment activity based on the information or strategies presented in this material.
2. No Liability: To the fullest extent permitted by applicable law, the creators, contributors, and distributors of this content disclaim all liability for any losses, damages, claims, or expenses—whether direct, indirect, incidental, or consequential—arising from reliance on or use of this material. This includes, but is not limited to, lost profits, trading losses, or disruptions to financial plans.
3. Indemnification: Users agree to indemnify, defend, and hold harmless the creators, contributors, and distributors of this material from and against any claims, damages, liabilities, or expenses (including attorney's fees) arising out of the user's reliance on or use of the information provided.
4. Personal Responsibility: Users affirm that they have independently evaluated the risks associated with any trading or investment decision and agree to proceed at their own discretion and liability.

Furthermore, users should be aware that the creators, contributors, and distributors of this content retain ownership of all intellectual property rights associated with this material. Unauthorized reproduction, distribution, or modification of this content is strictly prohibited and may violate applicable intellectual property laws. By accessing or downloading this content, users agree to abide by these terms and conditions and refrain from sharing, reselling, or otherwise redistributing this material without express written permission from its authors.

This material is provided on an "as is" basis and may include inaccuracies, typographical errors, or incomplete data. The authors reserve the right to make corrections or changes to the content at any time without notice. However, they are under no obligation to update, supplement, or revise this material to reflect changes in market conditions, new information, or other developments.

Users are strongly advised to consult with a licensed attorney, financial advisor, tax professional, or other qualified specialist to evaluate the implications of using the strategies or materials provided. Any reliance on the content of this material is entirely at the user's own risk. This disclaimer serves as a legally binding agreement between the user and the creators of this content, and users indicate their acceptance of these terms by accessing or utilizing this material in any form.

Lastly, trading and investing involve inherently speculative activities that may not be suitable for all individuals. The suitability of any particular investment strategy, asset, or trading approach depends on the user's specific financial situation, risk tolerance, objectives, and experience. Users should exercise caution and consider seeking professional guidance when engaging in financial activities that expose them to significant risks.