Cross-Market Arbitrage and Market Segmentation: Evidence from High-Frequency Interest Rate Data

Stefano Clerici

Dissertation written under the supervision of Professor Geraldo Cerqueiro.

Dissertation submitted in partial fulfilment of requirements for the MSc in Finance, at the Universidade Católica Portuguesa, 5/1/2026.

Abstract

This thesis investigates the efficiency of financial markets by analyzing the segmentation between institutional Short-Term Interest Rate (STIR) futures listed on the CME and the emerging landscape of regulated prediction markets (Kalshi). While traditional finance theory posits that arbitrage forces should enforce the Law of One Price across functionally equivalent assets, we hypothesize that structural barriers and distinct participant compositions create persistent pricing discrepancies.

Methodologically, we utilize a proprietary high-frequency dataset comprising tick-by-tick data for Fed Funds and SOFR futures, alongside limit order book data for binary event contracts. We first construct a rigorous "fair value" forward curve for the Effective Federal Funds Rate (EFFR), accounting for the term structure of calendar spreads and meeting probabilities. Subsequently, we develop a "hybrid hedging strategy" that utilizes a static replication approach to map the linear risk exposure of traditional futures onto the binary payoff structure of prediction markets.

Our empirical results strongly validate the "Limits to Arbitrage" hypothesis. We document that while the institutional market operates with near-perfect efficiency, the prediction market exhibits systematic inefficiencies driven by retail behavioral biases (favorite-longshot bias) and one-sided liquidity constraints. By implementing a cross-market arbitrage strategy, we demonstrate the ability to generate risk-free alpha, confirming that market segmentation effectively insulates retail mispricing from institutional correction.

Title: Cross-Market Arbitrage and Market Segmentation: Evidence from High-Frequency Interest Rate Data

Autor: Stefano Clerici

Keywords: Market Segmentation, Prediction Markets, Interest Rate Futures, Limits to Arbitrage, High-Frequency Trading.

Resumo

Esta tese investiga a eficiência dos mercados financeiros através da análise da segmentação entre os futuros institucionais de Taxas de Juro de Curto Prazo (STIR) listados na CME e o cenário emergente dos mercados de previsão regulamentados (Kalshi). Embora a teoria financeira tradicional postule que as forças de arbitragem devam impor a Lei do Preço Único entre ativos funcionalmente equivalentes, hipotetizamos que barreiras estruturais e composições distintas de participantes criam discrepâncias de preços persistentes.

Metodologicamente, utilizamos um conjunto de dados proprietário de alta frequência, composto por dados tick-by-tick para futuros de Fed Funds e SOFR, juntamente com dados do livro de ordens para contratos de eventos binários. Primeiramente, construímos uma curva a termo de "valor justo" rigorosa para a Taxa Efetiva dos Fed Funds (EFFR), contabilizando a estrutura a termo dos calendar spreads e as probabilidades das reuniões do FOMC. Subsequentemente, desenvolvemos uma "estratégia de cobertura híbrida" que utiliza uma abordagem de replicação estática para mapear a exposição ao risco linear dos futuros tradicionais na estrutura de payoff binário dos mercados de previsão.

Os nossos resultados empíricos validam fortemente a hipótese dos "Limites à Arbitragem". Documentamos que, enquanto o mercado institucional opera com eficiência quase perfeita, o mercado de previsão exibe ineficiências sistemáticas impulsionadas por vieses comportamentais de retalho e restrições de liquidez unilaterais. Ao implementar uma estratégia de arbitragem entre mercados, demonstramos a capacidade de gerar alfa livre de risco, confirmando que a segmentação do mercado isola efetivamente o mispricing do retalho da correção institucional.

Título: Arbitragem entre Mercados e Segmentação de Mercado: Evidência a partir de Dados de Taxas de Juro de Alta Frequência

Autor: Stefano Clerici

Palavras-chave: Segmentação de Mercado, Mercados de Previsão, Futuros de Taxas de Juro, Limites à Arbitragem, Trading de Alta Frequência.

Table of Contents:

Chapter 1: Introduction

Financial markets have long served as the primary mechanism for aggregating dispersed information and pricing future economic outcomes. In the domain of monetary policy, Short-Term Interest Rate (STIR) futures listed on the Chicago Mercantile Exchange (CME) represent the global standard for institutional hedging. These markets, dominated by sophisticated participants such as banks and high-frequency trading firms, are characterized by deep liquidity and high informational efficiency.

However, the recent financial landscape has witnessed the emergence of a structurally distinct venue: regulated prediction markets. Platforms like Kalshi, designated as a Designated Contract Market (DCM) by the Commodity Futures Trading Commission (CFTC) in 2020, have democratized access to financial risk transfer, allowing retail participants to trade "event contracts" on economic indicators. Unlike traditional futures, which offer linear payoffs suitable for balance sheet management, prediction markets utilize binary options designed for direct speculation on specific outcomes.

This thesis investigates the interaction between these two parallel ecosystems. Specifically, it addresses a central research question: In an era of high-frequency efficiency, does market segmentation between institutional futures and retail prediction markets allow for the persistence of risk-free arbitrage opportunities?

Theoretical literature suggests that in a unified market, the "Law of One Price", which state that assets generating identical cash flows and carrying equivalent risk profiles must trade at the same price, irrespective of the venue in which they are exchanged. In an efficient market characterized by rational agents and the absence of friction, any deviation from this equilibrium creates a risk-free arbitrage opportunity. Market participants are expected to exploit such discrepancies, buying the undervalued asset while simultaneously selling the overvalued one, thereby enforcing price convergence and restoring the singular valuation of the underlying risk, should prevail; any discrepancy between the implied probabilities of the CME and Kalshi should be rapidly eroded by arbitrageurs. However, the theory of "Limits to Arbitrage" (Shleifer and Vishny, 1997) argues that structural barriers, such as regulatory segmentation, capital constraints, and distinct clientele, can prevent this correction, allowing mispricings to endure.

To test this hypothesis empirically, we focus on the pricing of Federal Reserve monetary policy decisions. We utilize a high-frequency dataset comprising tick-by-tick data from CME Fed Funds Futures and Secured Overnight Financing Rate (SOFR) futures, alongside order book data from Kalshi event contracts. The methodological contribution of this work is twofold. First, we construct a rigorous pricing model that extracts a "fair value" forward curve for the Effective Federal Funds Rate (EFFR) and SOFR from the institutional market, accounting for the complex term structure of calendar spreads and meeting probabilities. Second, we develop a novel "hybrid hedging strategy" that maps the linear risk of traditional futures onto the binary payoff structure of prediction markets.

Our results provide strong empirical evidence of market segmentation. While we confirm that the CME market is highly efficient, with internal anomalies persisting for only seconds, we find that the prediction market exhibits persistent structural inefficiencies. By simulating a strategy that arbitrages the institutional consensus against retail sentiment, we demonstrate the ability to generate a risk-free profit, alpha, effectively immune to directional market movements. Specifically, we document that retail liquidity constraints and behavioral biases create a "favorite-longshot bias" in prediction markets, which can be systematically exploited by a sophisticated agent bridging the two venues.

Chapter 2: Theoretical and Institutional Background

2.1 Short-term U.S. interest rates and benchmark rates

Short-term U.S. money markets are summarized by a small set of benchmark overnight interest rates. The Effective Federal Funds Rate (EFFR) is the volume-weighted average rate at which depository institutions lend balances held at the Federal Reserve to one another on an unsecured overnight basis. As such, the EFFR reflects both the stance of monetary policy and prevailing credit and liquidity conditions in the unsecured interbank market.

SOFR rate, by contrast, is a broad measure of the cost of overnight borrowing collateralized by U.S. Treasury securities in the repurchase agreement (repo) market. Because transactions underlying SOFR are secured by high-quality collateral, SOFR embeds substantially less credit risk than unsecured funding rates such as the EFFR. The spread between unsecured and secured overnight rates can therefore be interpreted, with appropriate caution, as a proxy for the market-implied credit and liquidity premium in unsecured overnight funding.

Following the global benchmark reform, SOFR has replaced USD LIBOR as the main overnight reference rate for derivatives and cash instruments. In this new framework, EFFR remains central for monetary policy implementation, while SOFR has become the primary benchmark for pricing and hedging in derivatives markets. The interaction between these two rates, one unsecured and policy-driven, the other secured and market-driven, is at the core of the empirical analysis in this thesis.

2.2 Futures on short-term interest rates

Futures contracts written on short-term interest rates serve as the primary mechanism for market participants to trade and hedge expectations regarding future overnight funding conditions. These instruments represent the global standard for institutional hedging. Designed for professional participants, contracts listed on the Chicago Mercantile Exchange (CME) require significant margin outlays and a sophisticated understanding of convexities and basis risks, effectively creating a high barrier to entry for non-professional actors. Consequently, these markets are dominated by banks, asset managers, and high-frequency trading firms, ensuring high liquidity and pricing efficiency. The contracts considered in this study, Fed Funds Futures (FFF), one-month SOFR futures, and three-month SOFR futures, are the foundational tools utilized by these institutional players to manage interest rate exposure.

All these contracts are quoted according to the standard convention of

$$P = 100 - R \tag{1}$$

where P denotes the futures price and R is an interest rate expressed in percent per annum. Under standard no-arbitrage arguments and abstracting from small risk premia, the rate R can be interpreted as the market-implied expectation of a suitable average of the underlying overnight rate over the contract's reference period. A one-basis-point increase in the implied rate therefore corresponds to a 0.01 decrease in the futures price.

In a Fed Funds Futures (FFF) contract, the underlying is the arithmetic average of the daily EFFR over a given calendar month, including weekends. Similarly, a one-month SOFR futures contract references the arithmetic average of daily SOFR over its delivery month. In both cases, the contract can be viewed as settling on the simple average overnight funding cost over that month.

A three-month SOFR futures contract is instead linked to a compounded SOFR rate over a predefined reference quarter. The payoff is structured to replicate the return from continuously rolling an overnight secured position at SOFR throughout the quarter, with reinvestment of principal and interest on each business day. As a result, the three-month SOFR futures price reflects the market-implied compounded secured funding cost over the quarter.

These instruments provide complementary information: FFF embody expectations about the future path of the unsecured overnight policy rate (EFFR), while SOFR futures summarize expectations about secured funding conditions. Their joint behavior allows one to infer both the expected policy path and the expected credit component in unsecured funding. The distinction between secured and unsecured overnight rates is central to the interpretation of the information embedded in futures prices. The EFFR, being an unsecured interbank rate, reflects not only expectations about the stance of monetary policy, but also credit and liquidity premia in the market for reserves. SOFR, in contrast, is a secured rate based on collateralized transactions backed by U.S. Treasuries and is therefore much less sensitive to changes in the perceived credit risk of borrowers.

The basis spread between instruments referencing EFFR and those referencing SOFR provides a natural measure of this differential. In the context of this thesis, the basis can be defined either as the difference between the implied EFFR and implied SOFR rates over a given month, or, equivalently, as the appropriately scaled difference between Fed Funds Futures prices and one-month SOFR futures prices. Under the maintained assumption that other distortions are small, this basis can be interpreted as a market-implied credit component in unsecured overnight funding.

2.3 Day-count conventions, reference periods, and IMM dates

The mapping from futures prices to implied interest rates depends on the definition of the reference period and on the underlying day-count convention. The monthly average rate is obtained by weighing each daily rate by the number of calendar days to which it applies. Overnight fixings published on a business day typically apply to that day and, in many cases, to the subsequent weekend days, so the averaging procedure must account for both business and non-business days. For three-month SOFR futures, the reference period is a reference quarter anchored to standard International Monetary Market (IMM) dates. IMM dates are defined as the third Wednesday of March, June, September, and December. Each three-month SOFR futures contract refers to the quarter between two consecutive IMM dates. The contract specification prescribes how these daily fixings are compounded and annualized on a 360-day basis to obtain the reference rate that enters the futures payoff.

Clarifying these conventions is essential, because the empirical strategy reconstructs the forward path of overnight rates at a daily frequency and then aggregates them using exactly the conventions embedded in the contract specifications.

2.4 Monetary policy, FOMC meetings, and target rate changes

In the current Federal Reserve operating framework, the stance of monetary policy is communicated through a target range for the federal funds rate. This range is set by the Federal Open Market Committee (FOMC) and is typically adjusted in increments of 25 basis points. The EFFR usually trades within this target range, and policy decisions that change the range tend to generate discrete jumps in the EFFR.

Market participants form expectations not only about the future level of the target range, but also about the probability of policy moves at future FOMC meetings. These expectations are embedded in the prices of short-term interest rate futures. A contract whose reference month contains an FOMC meeting reflects a mixture of scenarios: one in which the target range remains unchanged, and others in which it is increased or decreased by one or more 25-basis-point steps. The resulting futures price can thus be interpreted as a probability-weighted average of the corresponding monthly average EFFR outcomes.

Understanding how these expectations are encoded in futures prices is crucial for constructing a forward curve for EFFR that is consistent with both the timing of FOMC meetings and the market-implied probabilities of policy changes. This, in turn, is necessary to obtain a coherent forward curve for SOFR and to evaluate potential arbitrage opportunities between three-month and one-month SOFR contracts.

2.5 Extracting expectations from futures prices: intuitive framework

The empirical strategy adopted in this thesis relies on a simple decomposition of each contract's reference period into three segments: (i) a realized segment, covering the days from the start of the contract to the current date; (ii) the current (front) month, comprising the remaining days in the ongoing month; and (iii) subsequent months, consisting of the months that follow the current month within the SOFR 3 months contract's reference period. Each segment is associated with a different information set and is treated differently in the construction of the forward curve.

For a given Fed Funds Futures contract, the futures price summarizes the market-implied monthly average EFFR. When part of the month has already elapsed, this average combines the realized EFFR (taken directly from the Federal Reserve's publications) and the expected EFFR over the remaining days, which must be inferred from the futures price. By isolating the contribution of realized days, one can back out the implied average rate for the not-yet-realized days in the current month. For subsequent months, where no days have yet occurred, the same logic implies that the futures price directly reflects the expected monthly average rate.

When a month includes an FOMC meeting, the expected path of EFFR over that month is not flat: under the possibility of a hike or a cut, the EFFR is expected to exhibit a jump on the meeting date. In such cases, the monthly average implied by the futures price can be decomposed into a pre-meeting rate and a post-meeting rate, weighted by the number of days before and after the meeting and by the probability of a policy change. The formal expressions used to retrieve pre- and post-meeting rates from futures prices and meeting probabilities are presented in the Methodology section.

Meeting probabilities themselves can be extracted from calendar spreads, defined as price differences between adjacent Fed Funds Futures contracts. A FOMC decision affects only the part of each contract month that lies after the meeting date. As a consequence, the meeting has a different impact on the expected monthly averages encoded in the two contracts, depending on how many post-meeting days each month contains. The calendar spread thus isolates the incremental effect of the meeting on expected short-term rates. If a full 25-basis-point hike or cut were fully priced in, the spread would trade at a predictable "theoretical" level determined by the contracts' exposures to the meeting. In practice, the observed spread is typically smaller in absolute value, reflecting the fact that the market assigns a probability strictly between zero and one to a policy move. Comparing the observed spread to its full-move benchmark yields an estimate of the market-implied probability of a hike or cut.

By combining these elements, realized overnight rates from the Federal Reserve, expected monthly averages from FFF and one-month SOFR futures, meeting probabilities inferred from calendar spreads, and the credit component extracted from the basis between secured and unsecured contracts, the analysis constructs a daily forward curve for EFFR and SOFR that is consistent with market prices and institutional details. This forward curve provides the theoretical benchmark against which observed prices of three-month SOFR futures are compared in order to assess the presence of potential arbitrage opportunities.

2.6 Alternative Markets

While CME derivatives effectively serve institutional hedging needs through standardized and capital-intensive instruments, the emerging landscape of prediction markets addresses a fundamentally different segment of participants. Unlike traditional futures, which are structured around linear or convex payoffs suitable for balance sheet management, alternative markets like Kalshi offer binary "event contracts" designed to democratize access to financial risk transfer.

2.7 Historical Context and Emergence

The theoretical foundation of prediction markets lies in the efficient market hypothesis and the concept of information aggregation, famously articulated by F.A. Hayek, which suggests that decentralized markets are superior mechanisms for processing dispersed information. The practical application of this theory materialized in 1988 with the establishment of the Iowa Electronic Markets (IEM) by the University of Iowa, a pedagogical and research-based platform designed to forecast election outcomes. While the IEM demonstrated the predictive efficacy of such mechanisms, the sector largely remained confined to academic experiments or unregulated, offshore entities for several decades.

In this context, Kalshi was founded in 2018 with the specific objective of integrating event contracts into the formal United States financial system. Unlike its predecessors, which often operated in legal gray areas, Kalshi sought explicit regulatory approval from the onset. In late 2020, the platform received designation as a DCM from the CFTC, marking a paradigm shift as the first federally regulated exchange dedicated exclusively to event contracts.

2.8 Operational Mechanism and Contract Structure

Functionally, the platform operates as a centralized limit order book exchange that facilitates the trading of binary event contracts. These instruments are structured as "all-or-nothing" derivatives based on the resolution of unambiguous future events. The pricing mechanism is inherently probabilistic: contracts are traded within a bounded price interval, typically normalized between zero and one unit of currency. The market price at any given moment reflects the aggregate consensus probability of the event occurring. Upon the resolution of the underlying event, settlement is binary; contracts associated with the realized outcome reach their full notional value, while opposing contracts expire worthless. This structure allows market participants to express views on future probabilities through a standardized financial instrument, distinct from traditional equity or fixed-income assets.

2.9 Regulatory Framework and Market Architecture

A defining characteristic of Kalshi is its institutional architecture, which diverges significantly from traditional betting operators and unregulated prediction platforms. As a CFTC-regulated exchange, the platform adheres to the Commodity Exchange Act, ensuring compliance with rigorous standards regarding market manipulation, transparency, and financial integrity. Unlike sportsbooks or betting operations where the operator acts as the counterparty to every transaction, Kalshi functions strictly as an exchange matching buyers and sellers. Furthermore, funds are segregated and protected, and trades are cleared through a regulated clearinghouse structure. This regulatory oversight aligns the platform more closely with traditional futures and commodities markets than with the gambling industry, mitigating counterparty risk and establishing a formalized legal framework for the trading of event risks.

Chapter 3: Data

3.1 Traditional Market Data

The transition from LIBOR to SOFR began in 2018, but it was not until 2020, when SOFR derivatives reached sufficient liquidity, that SOFR could effectively serve as a benchmark rate. For this reason, the data sample begins on 18 March 2020, corresponding to the first liquid quarterly SOFR futures contract, and extends to 4 December 2025.

Daily data on three-month SOFR futures, as well as one-month Fed Funds and SOFR futures prices, are obtained from Bloomberg. Intraday data are also sourced from Bloomberg, using the highest frequency permitted under our license, namely one observation every one second; due to the limited computational power the sample for the high frequency data starts at the beginning at the last IMM date, 17 September 2025, and ends at the same day of the daily sample, 21 November 2025. The Federal Reserve publishes daily effective Fed Funds and SOFR rates, and the associated historical series are retrieved from the Federal Reserve's database, together with the complete set of FOMC meeting dates.

For daily data, the paper uses settlement prices; for intraday data, the study collects both bid and ask quotes (and the relative number of orders) and compute mid-prices for model estimation. Although all data sources provide observations only on business days, we assume, when weekend values are required for the construction of the model, that Saturday and Sunday observations are equal to the preceding Friday's values.

Three-month SOFR futures contracts start and expire on IMM dates (see Appendix 1), whereas Fed Funds Futures and one-month SOFR futures begin on the first calendar day of each month and expire on the last. Further contract specifications are reported in Table 1.

Table 1. Futures Contracts Specifications

Specification	Fed Funds Futures	1-Month SOFR Futures	3-Month SOFR Futures
CME Ticker	ZQ	SR1	SR3
Underlying Rate	Arithmetic Avg. of daily EFFR	Arithmetic Avg. of daily SOFR	Compounded daily SOFR
Reference Period	1 Calendar Month	1 Calendar Month	1 Quarter (IMM to IMM)[1]
Price Quotation	100 - R	100 - R	100 - R
Contract Unit	Interest on $5,000,000 for 30 days	Interest on $5,000,000 for 30 days	Interest on $1,000,000 for 3 months
DV01	$41.67	$41.67	$25.00
Min. Price Fluctuation	0.25 bps front month 0.50 bps back month	0.25 bps front month 0.50 bps back month	0.25 bps front month 0.50 bps back month
Settlement Method	Cash Settled	Cash Settled	Cash Settled
Listing Cycle	Monthly (up to 60 months)	Monthly (up to 13 months)	Quarterly IMM (up to 10 years)

The dataset comprises 1,491 daily observations, as well as 1,596,000 intraday observations sampled at one second frequency.

For the maintenance margin information are found in CME website[2].

3.2 Proprietary Data

The analysis incorporates proprietary datasets regarding transaction costs, specifically clearing house commissions, initial margin requirements mandated by the CME, and applicable financing rates for capital borrowing.

Note 1. IMM dates are available in Appendix 1.

Note 2. Data retrieved from the following website in date 4/12/2025 (Website)

3.3 Alternative Markets Data

To perform a comparative analysis between non-traditional and traditional market structures, this study utilizes intraday data sourced from the alternative exchange, Kalshi. Specifically, we retrieved the full limit order book data, comprising bid prices, ask prices, and trading volumes, for the trading session of December 4, 2025. The dataset focuses on event contracts referencing the FOMC meeting scheduled for January 2026. For the traditional market benchmark, we collected parallel high-frequency data (bid, ask, and volume) for the CME Fed Funds Futures expiring in January and February 2026.

The Kalshi platform lists five mutually exclusive binary options corresponding to potential policy outcomes: "Maintain Rate", "Cut 25 bps", "Cut > 25 bps", "Hike 25 bps", and "Hike > 25 bps". Market participants may express directional views by entering "Yes" or "No" positions, which correspond operationally to transacting at the bid and ask prices present in the central limit order book.

A critical microstructural feature observed in the dataset is the presence of asymmetric liquidity. Specifically, for the "Hike 25 bps" and "Hike > 25 bps" contracts, liquidity was notably absent on the bid side. This constraint implies that participants were restricted in their ability to execute bilateral trades; specifically, the market structure allowed participants to only "bet on" the hike (establishing long positions via the ask) while precluding the execution of opposing views due to the lack of bids.

Functionally, Kalshi contracts are structured as binary derivatives priced within a bounded interval of $0.00 to $1.00. Settlement is determined by the realization of the underlying event: contracts expire at a value of $1.00 if the event occurs, and $0.00 otherwise. While the exchange facilitates trading in terms of contract quantity or notional dollar amounts (with automatic conversion), this study standardizes all positions in terms of USD notional exposure to ensure consistency in portfolio construction and PnL analysis. To provide a comprehensive overview of the market liquidity at the time of strategy implementation, a snapshot of the order book for all five binary options is provided in Appendix 3.

While Kalshi allows trading in terms of contract quantity, all positions in this study are converted into notional USD exposure to ensure consistency with CME futures analysis. Regarding transaction costs, Kalshi transparently published a fees schedule.

3.4 Historical Distribution of Policy Actions

From Bloomberg we took the historical target EFFR from 1948 to now, we define the policy action at meeting t as Δr_t, representing the change in the upper bound of the Federal Funds Target Range. Table 2 presents the descriptive statistics of these historical changes.

Figure 1. Frequency Distribution of FOMC Rate Changes (in bps)

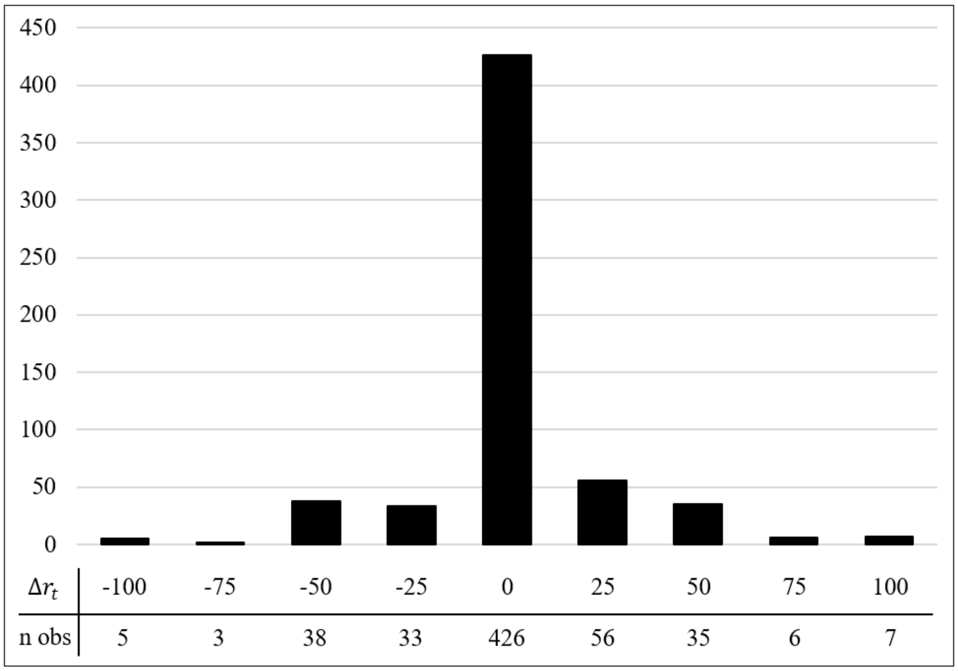

Table 2. Descriptive Statistics of FOMC Rate Changes (in bps)

Metric	Total Observations	Mean	Standard Deviation	Skewness	Kurtosis
Value	608	1.52	25.73	0.18	7.31

Chapter 4: Literature Review

This study draws upon three distinct strands of financial literature: the efficiency of institutional interest rate markets, the behavioral microstructure of prediction markets, and the theoretical limits of arbitrage in segmented venues. By synthesizing these fields, we frame the research question within the broader debate on how market structure influences price discovery.

4.1 Efficiency and Information Aggregation in Institutional Markets

The pricing of monetary policy expectations has been foundational to the study of fixed-income derivatives. In their seminal work, Piazzesi and Swanson (2008) challenged the then-prevailing view that federal funds futures rates were pure expectations of future policy. By decomposing excess returns into risk premia and forecast errors, they established that while these futures are the most accurate predictors available, significantly outperforming survey-based forecasts and simple macro-models, they contain a time-varying risk premium. Crucially, they found this premium to be countercyclical: investors demand higher compensation for holding rate risk during recessions. Their methodology demonstrates that futures prices aggregate information from agents with actual capital at risk ("skin in the game"), thereby incorporating real-time macroeconomic news more efficiently than non-binding survey responses.

Building on this, Gürkaynak (2007) investigated the speed of information incorporation using high-frequency event study methodology. By analyzing asset price reactions within tight windows around FOMC announcements, they confirmed that institutional interest rate markets incorporate new information with high velocity, often completing the price discovery process within minutes of a policy release. Their findings paint a picture of a market dominated by sophisticated, algorithmic participants where statistical anomalies are rapidly eroded. Our research aligns with this consensus, as the fleeting nature of the arbitrage windows we identify within the CME ecosystem, often persisting for mere seconds, confirms the high degree of efficiency characterizing the institutional segment.

4.2 Behavioral Anomalies in Prediction Markets

Parallel to traditional finance, a growing body of literature investigates "prediction markets" as alternative mechanisms for aggregating dispersed information. Wolfers and Zitzewitz (2004) provide a comprehensive overview, positing that decentralized markets can effectively forecast uncertain events, from election outcomes to economic indicators. They argue that the incentive structure of these markets, where payoffs are tied directly to accuracy, encourages participants to reveal their true beliefs.

However, unlike wholesale financial markets dominated by institutional capital, prediction markets are often characterized by a retail-driven participant composition. This demographic shift introduces specific pricing distortions. Snowberg and Wolfers (2010) document the prevalence of the "favorite-longshot bias" in these venues, where participants systematically overvalue low-probability outcomes (longshots) and undervalue high-probability ones (favorites). Their study explicitly tests two competing hypotheses for this bias: "risk-love" (preference for high variance) versus "misperception" (cognitive inability to assess small probabilities). They find strong evidence supporting the misperception hypothesis. In the context of our study, this suggests that retail participants on Kalshi may structurally misprice the "tail events" of monetary policy (e.g., unexpected rate cuts), creating a systemic divergence from the efficient pricing observed in CME futures.

4.3 Market Segmentation and Limits to Arbitrage

The persistence of pricing discrepancies between a highly efficient institutional market (CME) and a biased prediction market (Kalshi) presents a theoretical puzzle. According to the classical "Law of One Price," rational arbitrageurs should rapidly trade against these misalignments until they disappear. The framework of Shleifer and Vishny (1997) provides the theoretical resolution to this paradox through the concept of "Limits to Arbitrage."

They argue that real-world arbitrage is not risk-free but requires specialized capital and is subject to agency risks. Professional arbitrageurs typically manage other people's money and are evaluated on short-term performance ("performance-based arbitrage"). Consequently, they often face binding capital constraints, such as margin calls or risk limits, precisely when price divergences are most acute and the opportunity is greatest. Gromb and Vayanos (2010) extend this analysis to market segmentation, showing how barriers to capital mobility impede the flow of liquidity. When arbitrageurs cannot freely move capital between segmented pools (e.g., due to the regulatory and operational friction between a futures exchange and a designated contract market like Kalshi), pricing gaps can endure for extended periods.

This thesis contributes to this literature by documenting a specific, actionable instance of such segmentation. Unlike previous studies that often rely on daily closing prices or theoretical models, our use of high-frequency intraday data allows us to observe the mechanics of this segmentation at a granular level. We demonstrate that the "Limits to Arbitrage" are not merely theoretical constructs but tangible barriers that allow risk-free alpha to persist in the digital age.

Chapter 5: Methodology

The forward SOFR curves implied by one-month and three-month futures contracts must be internally consistent; otherwise, arbitrage opportunities would arise between instruments referencing the same underlying rate over overlapping horizons. The empirical strategy therefore focuses on identifying such potential arbitrage opportunities by comparing the "front" three-month SOFR futures contract with the set of one-month contracts that span the same reference period.

In doing so, the analysis must account for both realized and forward-looking components of the contract's valuation. The realized component consists of historical data, specifically, the daily EFFR and SOFR fixings published by the Federal Reserve, which enter the pricing formula as known quantities. The forward-looking component is based on market expectations: the market-implied path of the Effective Federal Funds Rate (EFFR), the market-implied credit component (which, together with EFFR, determines expected SOFR), and the market-implied probabilities of policy rate cuts or hikes at future FOMC meetings. These expectations are inferred from the term structure of Fed Funds Futures (FFF) and one-month SOFR futures.

In practice, a three-month SOFR futures contract typically overlaps with four consecutive one-month SOFR futures. Throughout the analysis, it is convenient to decompose the life of the three-month contract into three segments: a realized segment, the current (front) month, and subsequent months. Each segment is associated with a different information set and extraction procedure. Once the EFFR and credit components have been recovered across all segments, they are combined into an implied SOFR forward curve, which is then mapped into a theoretical three-month SOFR futures price using the contract's compounding formula.

5.1 Contract decomposition and information sets

For each observation date and each three-month SOFR futures contract, the reference quarter is partitioned into three segments, as shown in Figure 1:

I. **Realized segment** – from the start of the contract to the current date, where overnight EFFR and SOFR fixings are known. These realized values are taken directly from the Federal Reserve's database. On the first trading day of the contract, this segment is empty and therefore contributes zero.

II. **Current (front) month** – the remaining days in the ongoing month, for which EFFR and SOFR are unknown and must be inferred from market expectations embedded in one-month SOFR and FFF contracts.

III. **Subsequent months** – the one or two months following the current month that still fall within the three-month reference period, whose realizations are again unknown and are extracted from market-implied expectations. If the valuation date lies in the last reference month of the contract, this segment is empty and contributes no additional observations.

Figure 2. Time Segmentation

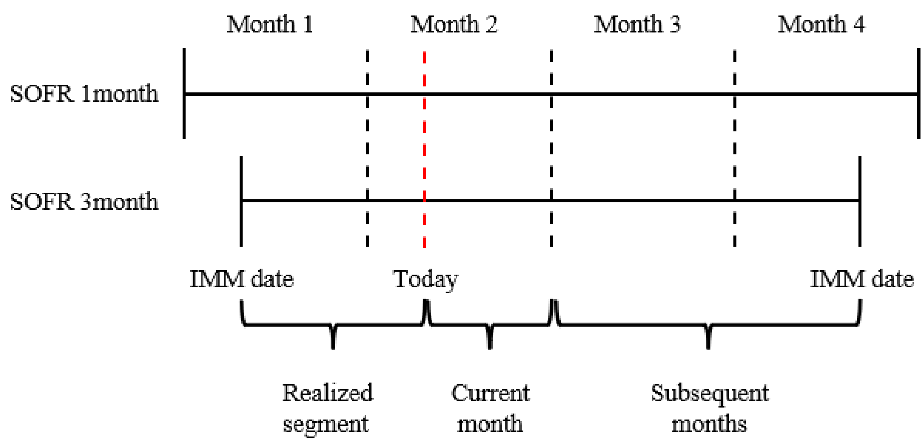

The remainder of this section details how information from FFF and one-month SOFR futures is used to construct the EFFR and credit curves over these segments, and how these curves are then used to price three-month SOFR futures.

5.2 Fed Funds Futures

Fed Funds Futures are one-month reference period derivatives that encode the market's expectation of the Effective Federal Funds Rate over the corresponding calendar month. At expiry, the contract settles at:

$$P = 100 - R \qquad (1)$$

where R denotes the arithmetic average of the daily EFFR over the entire reference month, including weekends. Under the assumptions discussed in the theoretical background, the contract price can be interpreted, up to a small risk premium, as the market-implied average EFFR over that month.

Within the framework of the three-segment decomposition, the information contained in FFF is used as follows.

I. **Realized segment.**

For days that have already occurred, the overnight EFFR is taken directly from the Federal Reserve's website. Since these are actual realizations, they are used without further adjustment in the construction of the realized segment.

II. **Subsequent months.**

For months following the current (front) month, there is no distinction between realized and future days, since none of the daily observations have yet occurred. In this case, the rate R implied by the FFF price is interpreted directly as the market-implied average EFFR for that future month.

III. **Current month.**

For the not-yet-realized days of the current month, the structure of the contract payoff is exploited to infer the expected EFFR over the remaining days. Let p denote the futures price and r the implied monthly average rate obtained by inverting the pricing relation $p = 100 - r$. This monthly average can be written as the sum of all daily rates divided by the total number of calendar days in the month. Equivalently, the sum of all daily rates is $r \times n$, where n is the number of calendar days in the month.

Subtracting from this total the sum of realized daily EFFR observations already observed in the current month yields the sum of the expected daily EFFR over the remaining days. Dividing this residual by the number of remaining days gives the market-implied average EFFR for the not-yet-realized portion of the month. Rearranging these relationships leads to equation (1), which expresses the futures price as a function of realized and missing daily EFFR contributions for the current month.

Applying this procedure to all FFF contracts in the sample yields a panel of expected average EFFR levels for the current and subsequent months, which will then be refined to account for FOMC meeting dates and implied policy moves.

5.3 Meetings Probabilities

Knowing the probability of a change in the target rate range is crucial for constructing the forward curve of the EFFR. To extract the implied probabilities of cuts and hikes, calendar spreads between adjacent FFF contracts are used. The key idea is to select spreads such that at least one of the two underlying monthly contracts is not affected by the FOMC meeting under consideration, irrespective of whether this is the first or the second month in the spread.

The procedure is as follows. First, for each individual monthly contract the **exposure** to a given FOMC meeting is computed as:

$$Exposure = \frac{Number\ of\ days\ after\ the\ FOMC}{Number\ of\ days\ of\ the\ month} \quad (2)$$

This exposure measures the fraction of the contract month that is affected by the new target range set at that meeting.

Next, the exposure of a calendar spread is defined as the difference between the exposure of the second month and that of the first month. Once the calendar's net exposure to a given meeting has been determined, it is multiplied by 25 basis points (or −25 basis points, corresponding to the usual step size in the target range) to obtain the level at which the calendar spread should trade if a full hike or cut of that size were fully priced in.

Let ΔP^{obs} denote the observed calendar spread and ΔP^{full} the corresponding theoretical full-move value. The market-implied probability p of an hike or a cut of 25 basis points is then obtained as the ratio:

$$p = \frac{\Delta P^{obs}}{\Delta P^{full}} \quad (3)$$

If the extracted probability exceeds 100%, this indicates that the market is pricing in a change in the target range larger than a single 25-basis-point move.

When more than one calendar spread isolates only the same FOMC meeting, the spread with the largest absolute exposure to that meeting is retained, as it provides a more precise estimate of the implied probability. Moreover, in configurations where three FOMC meetings occur in three consecutive months (as in December 2020, December 2021, June 2022, June 2023, December 2023, December 2024, and June 2025), it is necessary first to isolate one of the two "outer" meetings before identifying cleanly the exposure to the central meeting. A historical overview of the resulting implied probabilities and the effective policy decisions, when different from "unchanged", is reported in Appendix 2.

5.4 EFFR distribution: pre- and post-meeting rates

Fed Funds Futures provide the market-implied monthly average of the EFFR. When a given month contains an FOMC meeting, this average necessarily embeds the possibility of a discrete jump in the rate on the meeting date in the event of a hike or a cut. In such cases, it is essential to construct a forward curve for the EFFR that is properly distributed over the month, distinguishing between pre-meeting and post-meeting levels. By contrast, for months without an FOMC meeting there is no need to differentiate between pre- and post-meeting rates, and a single constant rate for the entire month suffices.

For months that include an FOMC meeting, both a pre-meeting and a post-meeting rate are recovered. Let R_{Exp} denote the expected monthly average EFFR implied by the FFF price, m the number of remaining days until the contract's expiration, N the number of days after the FOMC meeting within the contract's reference period, and p the model-implied probability of a rate cut or hike (in 25-basis-point units). The pre-meeting rate is given by:

$$R_{Pre} = R_{Exp} - \left(\frac{N * 0.25 * p}{m}\right) \tag{4}$$

while the post-meeting rate, equal to the expected pre-meeting rate plus the expected policy action at the FOMC, is given by:

$$R_{Post} = R_{Exp} - \left(\frac{N * 0.25 * p}{m}\right) + 0.25 * p \tag{5}$$

These two rates are computed for each contract and each observation date in the sample so that, at every point in time, an implied EFFR curve is obtained that features a discrete jump at each meeting date.

This implied EFFR curve then serves as the basis for constructing the corresponding implied SOFR curve, once the credit component has been added, and ultimately for computing the theoretical price of the three-month SOFR futures contract.

5.5 Credit curve from the SOFR–Fed Funds basis

The only difference between Fed Funds Futures and one-month SOFR futures lies in the underlying reference rate: instead of the EFFR, the one-month SOFR contract is written on the SOFR, which, unlike the former, is a secured overnight rate collateralized by U.S. Treasuries. The basis spread between FFF and one-month SOFR futures therefore reflects the market-implied credit component, as it represents the difference between an unsecured and a secured overnight funding rate.

The payoff structure of the one-month SOFR futures contract is otherwise identical to that of FFF. Consequently, the expected credit component for the remaining days of the current For months (front) month and for the subsequent months is extracted in exactly the same way as the expected EFFR for the corresponding data segments. For the realized segment, the observed SOFR fixings are used directly, and it is not necessary to recover any historical credit spread.

is written on the SOFR, which, unlike the former, is a secured overnight rate collateralized by U.S. Treasuries. The basis spread between FFF and one-month SOFR futures therefore reflects the

More specifically:

I. **Subsequent months**

For months following the current month, there is no distinction between realized and future days, since none of the daily observations have yet occurred. The basis spread is defined as:

$$BASIS = SOFR\ 1month_{Price} - FFF_{Price} \quad (6)$$

and the rate obtained by applying this relation is interpreted directly as the market-implied average credit component for that future month.

II. **Current month**

For the not-yet-realized days of the current month, the expected credit component is recovered using an expression analogous to the one derived for the EFFR. Let n denote the number of days in the month. The futures price can be written in terms of the average of realized and missing daily BASIS contributions; rearranging this relation yields equation (6), which expresses the futures price as a function of the sum of realized BASIS over the current month and the average missing BASIS over the remaining days.

Applying this procedure for all contracts and dates produces a credit curve defined over the entire reference period of each three-month SOFR futures contract.

5.6 SOFR 3 months: forward curve and futures valuation

Three-month SOFR futures expire as shown in Table 1, with settlement price:

$$P = 100 - R \tag{1}$$

where R is the business-day compounded SOFR rate per annum over the contract's reference quarter and is computed by:

$$R = \left\{ \prod_{i=1}^{n} \left[1 + \left(\frac{d_i}{360} \right) * \left(\frac{r_i}{100} \right) \right] - 1 \right\} * \left(\frac{360}{D} \right) * 100 \tag{7}$$

let n denote the number of U.S. government securities market business days in the reference quarter and let i be a running index over these business days, $i = 1,2, \dots, n$. The operator $\prod_{i=1}^{n}\{\ \}$ indicates the product over all values indexed by i from 1 to n. We denote by r_i the value of SOFR on the i-th U.S. government securities market business day, and by d_i the number of calendar days for which the rate r_i applies and D represents the total number of calendar days in the reference quarter.

To synthesize the theoretical price of the three-month SOFR futures contract, the SOFR forward curve implied by the market is first constructed using the components derived above. In particular, for each day in the reference quarter, the implied SOFR rate is obtained as the sum of the implied EFFR rate and the implied credit component for that day, with jump dynamics aligned with FOMC meeting dates as determined in the EFFR distribution step.

Given this daily SOFR forward curve, the sequence $\{r_i, d_i\}_{i=1}^{n}$is generated for each three-month contract, and equation (7) is applied to compute the model-implied compounded SOFR rate R. The corresponding theoretical futures price is then:

$$P^{model} = 100 - R^{model} \tag{8}$$

As a final step, the model output is rounded to the nearest quarter of a basis point, in order to align with the minimum price increment at which the contract is actually traded. In the empirical analysis, deviations between the observed market price and $P^{\mathrm{model}}$are interpreted as potential mispricings; persistent and economically large deviations, after accounting for transaction costs and margin requirements, are viewed as indicative of possible arbitrage opportunities between three-month and one-month SOFR contracts.

5.7 SOFR 3 months: exposure to the meetings

To implement the trading strategy and design appropriate hedges, it is useful to quantify directly the exposure of each three-month SOFR futures contract to individual FOMC meetings that fall within its reference quarter. The simplest way to obtain this measure is to run the pricing model under a set of counterfactual scenarios.

For a given three-month SOFR futures contract and for each FOMC meeting inside its reference quarter, two scenarios are considered:

1. **Baseline scenario (no policy moves).**

   All FOMC meetings are assumed to leave the target range unchanged. The model delivers a corresponding theoretical futures price, $P^{\text{no-move}}$.

2. **Single-meeting move scenario.**

   Exactly one FOMC meeting, the meeting under consideration, is assumed to result in a 25-basis-point change in the target range (hike or cut), while all other meetings are kept unchanged. The model yields a new theoretical price, P^{move}.

The meeting exposure of that three-month contract to the selected FOMC meeting is then defined as the difference between these two prices,

$$\text{Exposure}^{(j)} = P^{\text{move},(j)} - P^{\text{no-move}}, \tag{9}$$

where the superscript (j) indexes the meeting. Intuitively, this exposure measures how sensitive the three-month SOFR futures price is to a full 25-basis-point policy move at that specific meeting, holding all else equal. These exposures, expressed later in dollar terms using DV01s, form the basis for the hedge ratios used in the trading strategy.

5.8 Cross-Market Arbitrage Strategy

The final component of the methodology addresses the structural arbitrage between the institutional futures market (CME) and the prediction market (Kalshi). The objective is to construct a hybrid portfolio that neutralizes the directional risk of a speculative futures position by utilizing the binary event contracts available on the alternative exchange.

The strategy is predicated on a "static replication" approach. While the CME Fed Funds Futures exhibit a linear payoff profile, generating profits or losses proportional to the magnitude of the interest rate change, the prediction market offers contracts with a binary structure, paying a fixed amount only if a specific event occurs. To bridge this structural difference, the methodology first establishes a directional position in the traditional market. For the purpose of this study, we simulate a short position in the Fed Funds Futures calendar spread, which expresses a view anticipating a rate cut. We then map the Profit and Loss (PnL) profile of this position across the discrete set of scenarios defined in Section 3.1 (ranging from -100 bps to +100 bps).

The hedging process identifies the specific subset of scenarios where the futures position generates a negative return, specifically, outcomes where the Federal Reserve maintains rates or hikes them. To immunize the portfolio against these adverse movements, the strategy allocates capital to the corresponding mutually exclusive contracts on Kalshi. The allocation logic is straightforward: for every dollar of potential loss projected in the futures market for a given scenario, the strategy acquires a sufficient number of binary contracts covering that specific outcome. Since each winning binary contract provides a fixed payout (normalized to $1.00), the number of contracts purchased is calculated to ensure that the net payout fully offsets the futures liability. This effectively creates a "floor" for the portfolio performance.

Consistent with the microstructural features detailed in the Data section, the execution strategy incorporates specific liquidity constraints. Most notably, due to the lack of bid-side liquidity for "Hike" scenarios, the model operates under a strictly long-only constraint. This means the strategy cannot short the probability of a rate cut directly; instead, it synthesizes the hedge by purchasing the complementary "No Cut" and "Hike" options at the prevailing ask price. This approach ensures that the simulation accounts for the full transaction costs and reflects the realistic execution conditions faced by market participants.

Furthermore, consistent with the official fee schedule, transaction costs on the Kalshi platform are calculated according to the following formula:

$$Fees = roundup\ [0.07 * C * P * (1 - P)] \tag{10}$$

Where P is the price of a contract in USD, C is the contract quantity and roundup indicates rounding up to the next whole cent.

5.9 Scenario Bounds

To rigorously define the hedging range for the strategy, we analyzed the empirical distribution of historical Federal Open Market Committee (FOMC) target rate decisions. The objective is to determine whether the distribution of rate changes follows a Gaussian process, which would allow for parametric risk modeling, or whether a non-parametric approach is required.

Preliminary analysis indicates significant deviations from normality. Specifically, the distribution exhibits high excess kurtosis (indicating "fat tails" relative to a normal distribution) and non-zero skewness, reflecting the asymmetric nature of monetary policy cycles.

To formally assess the distributional properties of the data, we conduct a Jarque-Bera test for normality. We define the hypotheses as follows:

- H_0 (Null Hypothesis): The data is distributed according to a Normal distribution
- H_1 (Alternative Hypothesis): The data does not follow a Normal distribution.

The Jarque-Bera test statistic is calculated as:

$$Jarque\ Bera = n * \left(\frac{\widehat{Skew}^2}{6} + \frac{\left(\widehat{Kurt} - 3\right)^2}{24} \right) \quad (11)$$

where n is the sample size, $\widehat{Skew}$ is the sample skewness, and $\widehat{Kurt}$ is the sample kurtosis.

Table 3. Jarque Bera Test Results

Metric	JB Test	p-value
Value	473.10	<0.001

Given that the p-value is significantly lower than the standard significance level ($\propto = 0.01$), we reject the Null Hypothesis. This confirms that modelling rate changes using a Gaussian distribution would severely underestimate the tail risk.

Following the rejection of the normality hypothesis, we cannot rely on parametric method, we adopt a non-parametric historical simulation approach to define the hedging bounds.

Given the bidirectional nature of interest rate risk, where the portfolio is exposed to both aggressive easing and tightening cycles, we construct a bilateral confidence interval. We adopt the standard statistical convention of an equal-tailed interval, with $\alpha = 0.0001$.

The 99.99% confidence interval is defined as:

$$CI_{99.99\%} = [Q_{0.005\%}\ , Q_{99.995}]$$

By applying this methodology to the historical dataset, we observe that the interval $[-100\ bps, +100\ bps]$ contains the central probability mass required. Specifically, historical moves larger than a 100 bps in both directions fall into the rejection regions, allowing us to treat them as statistical outliers for hedging purposes.

By applying this methodology to the historical dataset, the interval $[-100\ bps, +100\ bps]$ encompasses more than 99.99% of historical outcomes. Specifically, historical moves larger than a 100 bps in both directions fall into the rejection regions. Thus, the probability of a rate move outside these bounds is statistically indistinguishable from zero. Therefore, the hedging strategy is calibrated to immunize the portfolio strictly within the set of scenarios included in the range:

$$\Omega = [-100\ bps, +100\ bps]$$

Chapter 6: Results

Using daily settlement data, the model yields a maximum absolute difference of 1.25 basis points between the synthetic three-month SOFR futures price and the corresponding market price. Additional summary statistics for the distribution of pricing errors are reported in Table 4. Given that the synthetic price is constructed from a total of four instruments and then compared to a fifth one, for each of them, the daily quote may correspond to either the bid or the ask (which is not observable in the settlement/daily data), such a discrepancy is not entirely surprising. In fact, as shown below, these large differences almost completely disappear once we move to one-second frequency data, and the maximum difference of 1.25 basis points can be attributed to this settlement effect. By analyzing one-second observations by using mid prices, we find that, on average, there are only 43 seconds per trading day during which the model–market discrepancy reaches 0.50 bps, the maximum discrepancy observed in the entire period.

Table 4. Descriptive Statistics of model's mispricing

Metric	Value (bps)	
	Daily	Intraday
Mean	0.23	0.26
Median	0.25	0.25
Max	1.25	0.50
Min	0.00	0.00
Standard Deviation	0.23	0.12

This implies that the market is priced in line with the model for approximately 99.85% of the trading time and is potentially inefficient only in the remaining 0.15% of the time.

To assess whether these short-lived discrepancies are economically meaningful, we construct a hypothetical trading strategy and evaluate its performance.

6.1 Trading strategy and hedging ratios

The construction of the strategy relies on two key ingredients provided by the model:

1. Meeting exposures

   For each three-month contract and each FOMC meeting in its reference quarter, the scenario-based procedure above yields a measure of the contract's exposure to a 25-basis-point policy move at that meeting.

2. Dollar sensitivities

   The DV01 (dollar value of one basis point) of a three-month futures contract differs from that of a one-month contract. As a result, it is not sufficient to compare contracts in terms of raw basis-point changes; hedging must be carried out in dollar space, using DV01s.

In practice, for a given valuation date and a given three-month contract, we first compute the vector of meeting exposures for that contract and for all overlapping one-month futures. These exposures are then converted into dollar terms by multiplying them by the corresponding DV01s. The hedge ratios are chosen so as to make the net tail exposure of the combined position as close as possible to zero, subject to the additional requirement that the resulting total number of contracts remains sufficiently small (eg. maximum 100 contracts) to avoid creating excessive execution risk.

As presented in Table 5, the optimal hedge allocation satisfying the imposed size constraints while minimizing tail exposure involves a long position of 34 three-month SOFR contracts and short positions of 15 October, 8 November and 2 December one-month contracts

Table 5. Hedging Table

		Exposure to meetings			
		Per contract (bps)		Total ($)	
Contract	N of contract	October	December	October	December
Sep 3m	34	13.28	1.66	11,291	1,412
Oct 1m	-15	1.61		(996)	
Nov 1m	-8	25.00		(8,233)	
Dec 1m	-2	25.00	16.94	(2,058)	(1,394)
Tail				4	17

The resulting strategy exhibits a net tail exposure of $21 over the reference period. This aggregate figure decomposes into a negative tail sensitivity to the October meeting, implies that a rate cut in that month would generate a marginal loss of $4, and a positive sensitivity to the December meeting, implying that a rate cut in the last month of the year would yield a profit of $17.

To mitigate execution risk, the strategy is implemented using market orders rather than limit orders. When entering the position, buy orders are assumed to be executed by lifting the best available ask, and sell orders by hitting the best available bid. Before submitting the orders, we verify that there is sufficient depth at those prices to execute the full size implied by the hedge ratios in all relevant contracts. The same procedure is followed when closing the position: all legs are unwound by hitting the market, again conditional on a prior liquidity check. Operationally, when the three-month futures contract appears overpriced relative to the strip, the strategy consists of selling the three-month contract and buying the appropriate combination of one-month SOFR and Fed Funds Futures implied by the hedge ratios. When the three-month contract appears underpriced, the direction of the trade is reverse.

6.2 Intraday Implementation

To illustrate the real-time execution mechanics of the strategy, the paper analyzes the trading session of October 10th, selected as a representative sample of average market conditions. On this date, the model identifies two distinct operational windows characterized by a significant pricing discrepancy. Notably, the mispricing reaches its historical maximum delta of 0.50 bps for extremely short intervals, persisting for only 6 and 4 seconds, respectively. Despite this transient nature, sufficient liquidity is available during these peaks to initiate the full relative-value trade. At the opening of each window, 10:29:00 and 15:53:15, we initiate the relative-value trade by hitting the market in the three-month SOFR futures and in all overlapping one-month contracts, with quantities determined by the hedge ratios described above. The positions are subsequently held until the spread reverts to full alignment, allowing the strategy to exit at 10:44:48 and 15:55:45, respectively.

Table 6. Trades summary

		SOFR Futures Traded Positions				SOFR Futures Prices			
Time	Delta	1m October	1m November	1m December	3m September	1m October	1m November	1m December	3m September
10:29:00	0.50	-15	-8	-2	34	95.8575	96.0850	96.2150	95.9600
10:44:48	0.00	15	8	2	-34	95.8600	96.0800	96.2050	95.9550
13:55:15	0.50	-15	-8	-2	34	95.8550	96.0800	96.2150	95.9575
15:55:45	0.00	15	8	2	-34	95.8600	96.0900	96.2250	95.9600

Using these execution prices and quantities, we compute the realized profit and loss of each round-trip trade, both in basis points and in dollars.

Table 7. Trading PnLs

Change of Prices (bps)				PnL per position ($)				
1m October	1m November	1m December	3m September	1m October	1m November	1m December	3m September	PnL Trade ($)
0.25	-0.50	-1.00	-0.50	-156.26	166.68	83.34	-425.00	-331.24
0.50	1.00	1.00	0.25	-312.52	-333.36	-83.34	212.50	-516.73

Execution simulations incorporate all applicable exchange fees and clearing house commissions. Based on CME membership status, these costs are estimated at $0.19 and $0.05 per contract and trade, respectively. The total transaction cost for the round-trip execution amounts to $28.32.

Furthermore, holding derivative positions necessitates the posting of initial margin. According to public CME data, the maintenance margin for the portfolio is $5,947, with an estimated initial margin of $6,541[3]. Accounting for the cost of carry, assuming a financing rate of 6.85% per annum[4] for a medium-sized firm, results in a daily financing cost of $1.23. Consequently, the net loss widens to -$360.79 for the first trade and -$546.28 for the second trade, as shown in Table 8.

Table 8. Total PnL

PnL Trade	Commissions	Financing Costs	Final PnL
-331.24	-28.32	-1.23	-360.79
-516.73	-28.32	-1.23	-546.28

Note 3: Proprietary Data
Note 4: Proprietary Data

6.3 Alternative Markets vs Traditional Markets

This analysis simulates a trading scenario executed on December 4, 2025, at hour 15:30:00, focusing on the outcome of the January 2026 FOMC meeting. The study considers a range of policy scenarios, specifically rate adjustments between -100 bps and +100 bps, as extreme deviations beyond this range are statistically negligible.

The strategy initiates a short position in the Jan/Feb Fed Funds Futures (FFF) calendar contract, predicated on a divergence in implied rate-cut probabilities. Specifically, the traditional market discounts a probability of -28.89%, contrasted with the -30% implied by the Kalshi prediction market. This approach effectively exploits the valuation spread by shorting the overvalued instrument (characterized by the higher implied probability) while taking a long position in the undervalued counterpart. This position, established via a market order at -6.5, effectively speculates on a rate cut of at least 25 bps. According to CME requirements, holding this position necessitates an initial margin of $1,130[5].

As illustrated in point 1 of Table 9, an unhedged FFF position yields a positive PnL in the event of a rate cut. Specifically, the calendar spread price is expected to decrease by 22.58 points for every 25 bps cut, generating profit. However, this exposure carries significant downside risk i.e. if the Federal Reserve maintains the current rate or hikes it, the portfolio incurs losses. For instance, a "no change" scenario results in a loss of approximately $271, with losses increasing further in hiking scenarios.

To mitigate these risks, the strategy incorporates options from the alternative market, Kalshi, to hedge against unfavorable outcomes, a summary of all the positions, the relative costs and fees is shown in Table 9.

1. **Hedging the Maintain Rate Scenario:** A long position is taken in Kalshi contracts betting against a 25 bps cut (effectively a "No Cut" prediction). Purchasing 793 contracts at $0.70 (total cost: $555) provides a payoff that offsets the FFF loss if rates remain unchanged. Conversely, if a cut occurs, these contracts expire worthless, but the loss is absorbed by the profits from the FFF position.
2. **Hedging the "Hike" Scenarios:** To fully immunize the portfolio against rate hikes, additional long positions are established in Kalshi binary options corresponding to "Hike 25 bps" and "Hike > 25 bps" outcomes.

Note 5: Proprietary Data

Table 9. Hedging Strategy Summary

	Instrument	Number of Contracts	Price per Contract	Total Cost	Fees
1.	FFF	1 Calendar	$1,130	$1,130.00	$0.48
2.	Cut 25 bps	793	$0.70	$555.00	$11.66
3.	Hike 25 bps	767	$0.03	$23.00	$1.57
4.	Hike > 25 bps	5000	$0.01	$50.00	$3.47
			Total	$1,758.00	$17.18

As shown in the final summation in Table 10, this hybrid strategy ensures a positive PnL across all considered scenarios, effectively neutralizing the directionality risk associated with the standalone traditional futures position.

Table 10. Scenarios Analysis ($)

	Instrument	Scenarios -100bps	-75bps	-50bps	-25bps	0	25bps	50bps	75bps	100bps
1.	FFF	3,492.78	2,551.87	1,610.96	670.05	(270.86)	(1,211.76)	(2,152.67)	(3,093.58)	(4,034.49)
2.	Cut 25 bps	793.10	793.10	793.10	(555.00)	793.10	793.10	793.10	793.10	793.10
	Subtotal	4,285.87	3,344.97	2,404.06	115.05	522.24	(418.67)	(1,359.58)	(2,300.49)	(3,241.39)
3.	Hike 25 bps	(23.00)	(23.00)	(23.00)	(23.00)	(23.00)	766.68	(23.00)	(23.00)	(23.00)
	Subtotal	4,262.87	3,321.97	2,381.06	92.05	499.24	348.01	(1,382.58)	(2,323.49)	(3,264.39)
4.	Hike > 25 bps	(50.00)	(50.00)	(50.00)	(50.00)	(50.00)	(50.00)	5,000.00	5,000.00	5,000.00
	Total	4,212.87	3,271.97	2,331.06	42.05	449.24	298.01	3,617.42	2,676.51	1,735.61

The inclusion of cumulative fees and borrowing costs, calculated at $17.41 of fees (derived from $0.48 for the FFF calendar and $16.93 for Kalshi) and $18.18 as borrowing cost for $1,758 for 55 days, establishes the strategy's net performance boundaries. Consequently, the adjusted model demonstrates a minimum positive return of $6.69 assuming rates remain unchanged, contrasting with a peak return of $4,177.51 contingent on a 100 bps reduction in rates, as shown in Table 11.

Table 11. Net PnLs ($)

	Scenarios -100bps	-75bps	-50bps	-25bps	0	25bps	50bps	75bps	100bps
Pnl Strategy	4,212.87	3,271.97	2,331.06	42.05	449.24	298.01	3,617.42	2,676.51	1,735.61
Costs and Fees	(35.36)	(35.36)	(35.36)	(35.36)	(35.36)	(35.36)	(35.36)	(35.36)	(35.36)
Net PnL	4,177.51	3,236.61	2,295.70	6.69	413.88	262.65	3,582.06	2,641.15	1,700.25

Chapter 7: Limitations and Future Research

This study validates the existence of arbitrage opportunities between institutional and predictive markets; however, several limitations regarding data scope and practical implementation must be acknowledged.

7.1 Data and Methodological Constraints

The primary limitation pertains to the data sample employed. The analysis of intraday data is constrained to the temporal scope ranging from September 18, 2025, to December 4, 2025. To enhance the precision and generalizability of the findings, future research should expand this timeframe to investigate mispricing across 3-month futures contracts more broadly, rather than restricting the analysis to the specific front-month contract active at the time of writing. Furthermore, the dataset concerning prediction markets, specifically Kalshi, could be significantly enlarged to study the frequency and persistence of pricing misalignments over multiple monetary policy cycles. Additionally, the synthetic pricing model for the 3-month SOFR future is currently calibrated exclusively for the front end of the curve; future studies could extend this framework to test the model's robustness across the broader term structure.

7.2 Execution Risk and Latency

A critical practical limitation concerns the assumption of frictionless execution. The simulation assumes that trades can be executed instantaneously at the displayed limit order book prices. However, the empirical analysis identifies arbitrage windows that are often ephemeral, persisting for as little as 4 to 6 seconds. In a live high-frequency trading environment, technological latency and slippage could significantly erode the theoretical risk-free profits documented here. Consequently, the results likely represent an upper bound on profitability achievable in a real-world setting.

7.3 Scalability and Institutional Implementation

Finally, while this study acknowledges liquidity constraints by imposing a theoretical cap on position size, a comprehensive analysis of the market's absorption capacity was beyond the scope of this research. Future studies should rigorously investigate the depth of the prediction market's limit order book to quantify the maximum executable volume available without causing significant price impact. In particular, further research is needed to determine whether such cross-market arbitrage strategies can be effectively scaled to accommodate the capital deployment requirements of large hedge funds and HFT firms, or if they remain structurally confined to smaller proprietary trading scopes.

Chapter 8: Conclusion

This thesis sets out to investigate whether market segmentation creates actionable arbitrage opportunities between traditional and alternative interest rate markets. Our empirical findings strongly validate the theoretical framework established in the literature.

First, consistent with Piazzesi and Swanson (2008), we confirm that the traditional CME market operates with high efficiency; theoretical anomalies are ephemeral and largely inaccessible due to transaction costs. Second, our analysis of Kalshi supports the behavioral biases documented by Snowberg and Wolfers (2010), revealing persistent pricing inefficiencies driven by retail liquidity constraints.

Crucially, the successful execution of the hedged strategy provides empirical evidence for the Limits to Arbitrage (Shleifer and Vishny, 1997). We conclude that despite the mathematical equivalence of the underlying risks, the structural barriers preventing institutional capital from flowing into prediction markets allow these price discrepancies to endure. Thus, we answer our central research question in the affirmative: market segmentation does indeed generate a distinct source of alpha, inaccessible to participants confined to a single market ecosystem.

Chapter 9: Appendix

1. SOFR 3months contracts with beginning and ending IMM dates.

Contract	Beginning Date	Ending Date
Mar2020	3/18/2020	6/17/2020
Jun2020	6/17/2020	9/16/2020
Sep2020	9/16/2020	12/16/2020
Dec2020	12/16/2020	3/17/2021
Mar2021	3/17/2021	6/16/2021
Jun2021	6/16/2021	9/15/2021
Sep2021	9/15/2021	12/15/2021
Dec2021	12/15/2021	3/16/2022
Mar2022	3/16/2022	6/15/2022
Jun2022	6/15/2022	9/21/2022
Sep2022	9/21/2022	12/21/2022
Dec2022	12/21/2022	3/15/2023
Mar2023	3/15/2023	6/21/2023
Jun2023	6/21/2023	9/20/2023
Sep2023	9/20/2023	12/20/2023
Dec2023	12/20/2023	3/20/2024
Mar2024	3/20/2024	6/19/2024
Jun2024	6/19/2024	9/18/2024
Sep2024	9/18/2024	12/18/2024
Dec2024	12/18/2024	3/19/2025
Mar2025	3/19/2025	6/18/2025
Jun2025	6/18/2025	9/17/2025
Sep2025	9/17/2025	12/17/2025

2. FOMC meeting, implied probability from the model and actual decision

FOMC Meeting	Probability implied in the market	Action	Old target rate	New target rate
3/16/2022	101.26%	+25 bps	0.00-0.25	0.25-0.50
5/4/2022	202.07%	+50 bps	0.25-0.50	0.75-1.00
6/15/2022	302.44%	+75 bps	0.75-1.00	1.50-1.75
7/27/2022	301.96%	+75 bps	1.50-1.75	2.25-2.50
9/21/2022	302.85%	+75 bps	2.25-2.50	3.00-3.25
11/2/2022	301.07%	+75 bps	3.00-3.25	3.75-4.00
12/14/2022	201.49%	+50 bps	3.75-4.00	4.25-4.50
2/1/2023	100.59%	+25 bps	4.25-4.50	4.50-4.75
3/22/2023	94.40%	+25 bps	4.50-4.75	4.75-5.00
5/3/2023	96.32%	+25 bps	4.75-5.00	5.00-5.25
7/26/2023	97.76%	+25 bps	5.00-5.25	5.25-5.50
9/18/2024	-203.33%	-50 bps	5.25-5.50	4.75-5.00
11/7/2024	-101.73%	-25 bps	4.75-5.00	4.50-4.75
12/18/2024	-100.89%	-25 bps	4.50-4.75	4.25-4.50
9/17/2025	-99.23%	-25 bps	4.25-4.50	4.00-4.25
10/29/2025	-87.65%	-25 bps	4.00-4.25	3.75-4.00

3. Kalshi Binary Options in Order Book, 12/4/2025 15:30:00

Cut 25 bps				
Yes		Price	No	
Volume ($)	Number of contracts		Number of contracts	Volume ($)
		$0.79	3,096	$2,446
		$0.75	62,000	$46,500
		$0.73	25,209	$18,403
		$0.72	31,234	$22,488
$11,117	15,882	$0.70		
$35,942	52,090	$0.69		
$7,582	11,316	$0.67		
$18,200	28,000	$0.65		

Cut > 25 bps				
Bid		Price	Ask	
Volume ($)	Number of contracts		Number of contracts	Volume ($)
		$0.10	1,100	$110
		$0.08	120	$10
		$0.07	100	$7
		$0.05	2,700	$135
		$0.04	50,000	$2,000
		$0.03	34,367	$1,031
		$0.02	180,948	$3,619
$8,000	800,001	$0.01		

Hike 25 bps				
Bid		Price	Ask	
Volume ($)	Number of contracts		Number of contracts	Volume ($)
		$0.16	1,379	$221
		$0.15	1,291	$194
		$0.14	1,262	$177
		$0.13	910	$118
		$0.12	910	$109
		$0.11	734	$81
		$0.04	52,000	$2,080
		$0.03	640,881	$19,226

Hike > 25 bps				
Bid		Price	Ask	
Volume ($)	Number of contracts		Number of contracts	Volume ($)
		$0.32	1	$0
		$0.16	2,184	$349
		$0.15	1,627	$244
		$0.14	1,604	$225
		$0.13	8,892	$1,156
		$0.12	869	$104
		$0.11	290	$32
		$0.04	2,000	$80
		$0.01	541,751	$5,418

Fed Maintain Rates				
Bid		Price	Ask	
Volume ($)	Number of contracts		Number of contracts	Volume ($)
		$0.63	3,224	2,031
		$0.61	64,100	39,101
		$0.59	31,131	18,367
		$0.58	32,745	18,992
$40,643	71,304	$0.57		
$17,086	30,511	$0.56		
$451	820	$0.55		
$1,064	2,008	$0.53		

Chapter 10: References

Piazzesi, Monika, and Eric Swanson. “Futures Prices as Risk-adjusted Forecasts of Monetary Policy.” *National Bureau of Economic Research*, June 1, 2004. https://doi.org/10.3386/w10547.

Gürkaynak, Refet S, Brian P Sack, and Eric T Swanson. “Market-Based Measures of Monetary Policy Expectations.” *Journal of Business and Economic Statistics* 25, no. 2 (March 24, 2007): 201–12. https://doi.org/10.1198/073500106000000387.

Wolfers, Justin, and Eric Zitzewitz. “Prediction Markets,” May 1, 2004. https://doi.org/10.3386/w10504.

Snowberg, Erik, and Justin Wolfers. “Explaining the Favorite-Longshot Bias: Is it Risk-Love or Misperceptions?” *National Bureau of Economic Research*, April 1, 2010. https://doi.org/10.3386/w15923.

Vishny, Andrei Shleifer Robert W. “The Limits of Arbitrage.” *NBER*, July 1, 1995. https://doi.org/10.3386/w5167.

Gromb, Denis, and Dimitri Vayanos. “Limits of Arbitrage: The State of the Theory,” March 1, 2010. https://doi.org/10.3386/w15821.