

Medium-Horizon Return Predictability in the KOSPI 200 Futures Basis

Timing, Signal Construction, and Out-of-Sample Evidence

Jaehyun Jo

Independent Researcher

sophrosophia@gmail.com

September 2026

Abstract

This paper asks whether the KOSPI 200 futures basis contains predictive information for cash-index returns beyond the short horizons emphasized in prior Korean research. Using daily predictor data from January 2004 through July 2026, I construct a fair-value-adjusted end-of-day basis for the nearby KOSPI 200 futures contract from the CD 91-day rate, the FY1 expected KOSPI 200 dividend yield, and actual time to expiration; three-month predictive regressions end in April 2026 because future returns must be observed. Cash and futures end-of-day observations are not perfectly synchronized. This creates measurement noise in daily basis observations generally and an additional timing concern at the forecast boundary, because the month-end cash close also anchors the subsequent return. The primary monthly signal therefore excludes the final trading day, and a stronger robustness test delays the return start to the next trading-day close. Higher basis values are associated with lower subsequent KOSPI 200 returns, with the clearest full-control evidence at two- to three-month horizons. In the three-month specification, the coefficient is -8.15 (HAC $t = -2.13$, $p = 0.033$) after conditioning on recent returns, realized volatility, VKOSPI, USD/KRW, and the maturity cycle. A two-sided moving-block max- $|t|$ correction across the one-, two-, three-, and six-month horizons gives a three-month family-wise adjusted p -value of 0.010 in the parsimonious model and 0.091 with the full control set. A 120-month rolling forecast based on the monthly average basis produces an out-of-sample R^2 of 3.73%. Alternative monthly and duration summaries sometimes produce larger OOS R^2 values, but those comparisons are exploratory and do not establish a distinct persistence factor. The results are sensitive to the global financial crisis and to contract construction near expiration. Rolling into a matched second-nearby contract before expiration materially weakens the full-control in-sample coefficient, although the parsimonious three-month relation and rolling OOS performance remain positive. Overall, the evidence is consistent with medium-horizon predictive information in the KOSPI 200 futures basis, but not with a causal mechanism or an unconditional trading anomaly.

Keywords: KOSPI 200, index futures, futures-cash basis, return predictability, cost of carry, out-of-sample forecasting

JEL Classification: G12, G13, G14

1. Introduction

The futures-cash basis is one of the most direct observable links between a derivatives market and its underlying cash market. Under the standard cost-of-carry relation, the price of an equity-index futures contract is tied to the spot index through financing costs, expected dividends, and time to maturity. In actual markets, however, the observed futures price can deviate from this benchmark. Transaction costs, funding conditions, short-sale frictions, hedging demand, liquidity demand, investor positioning, and market microstructure can all contribute to the observed basis.

A large literature studies these deviations through the lens of market integration and limits to arbitrage. Roll, Schwartz, and Subrahmanyam (2007) document a dynamic relation between index futures-cash bases and stock-market liquidity. More recently, Hazelkorn, Moskowitz, and Vasudevan (2023) show across global equity-index markets that the basis contains information about demand for equity-market exposure, comoves with futures positions, and negatively predicts subsequent spot and futures returns. Their evidence also shows that basis dynamics can persist beyond very short intraday horizons.

The KOSPI 200 market provides a useful single-market setting because it combines a long futures history with an extensive literature on price discovery and mispricing. KOSPI 200 futures have long played an important price-discovery role. Nam et al. (2006) show that KOSPI 200 futures lead the underlying cash index at intraday frequencies, and Joo, Seon, and Lee (2016) document an important information role for overnight futures trading. Korean research also studies the determinants of KOSPI 200 futures mispricing. Kim (2018) links futures-implied spot-price deviations to prior market-return moments, while Ryu, Ryu, and Yang (2025) show that investor sentiment, volatility, short-sale restrictions, and market reforms influence mispricing and limits to arbitrage.

Most importantly for the present paper, Kim (2020) already documents that KOSPI 200 futures mispricing predicts subsequent cash-market returns. Using a cost-of-carry-based implied spot index, that paper finds a negative relation between mispricing and future KOSPI 200 returns, with the reversal effect lasting up to roughly two weeks. The present paper therefore does **not** claim to discover return predictability from KOSPI 200 futures mispricing for the first time.

Instead, this paper asks three narrower questions. First, does the information in the KOSPI 200 futures basis extend from the short horizons documented in Korea to subsequent returns over roughly two to three months? Second, does the result depend on how noisy daily basis observations are summarized within the month? Third, does any predictive relation remain when forecasts are generated pseudo-real-time using only returns that would already have been observed at each forecast origin?

I construct a daily fair-value-adjusted basis for the nearby KOSPI 200 futures contract. The preferred fair value uses the Korean CD 91-day rate as the financing-rate proxy, the FY1 expected dividend yield of the KOSPI 200, and actual remaining days to contract expiration. A matched second-nearby futures series is available over the main sample and is used to test whether the results depend on remaining in the nearby contract close to expiration. The main monthly predictor is the average of daily abnormal-basis observations excluding the final trading day of each month.

The month-end exclusion addresses a specific forecast-boundary problem rather than the broader fact that cash and futures closing prices are non-synchronous on many trading days. Under current Korea Exchange rules, the KOSPI cash market regular session ends at 15:30, whereas regular trading in KOSPI 200 futures extends to 15:45; trading hours have changed over the sample, so the empirical

design does not assume a fixed historical gap.¹ Non-synchronous closes can add measurement noise to any daily end-of-day basis. The final trading day creates an additional issue, however, because the same month-end cash close also defines the starting level of the subsequent return. If the futures close contains information arriving after that cash close, including the month-end basis would allow part of the predictor to be observed after the return starting point. The baseline therefore excludes that one day. Robustness tests exclude additional month-end observations, and the strongest timing test moves the return start to the following trading-day close. These choices reduce forecast-boundary overlap; they do not claim to synchronize every daily cash and futures observation.

The empirical results can be summarized in four steps. First, higher fair-value-adjusted basis values are associated with lower subsequent KOSPI 200 returns, with the clearest full-control evidence at two- to three-month horizons. The three-month coefficient is -11.34 (HAC $t = -3.01$) in the parsimonious regression and -8.15 ($t = -2.13$, $p = 0.033$) after controlling for recent returns, realized volatility, VKOSPI, USD/KRW changes, and the maturity cycle. A one-standard-deviation increase in the monthly basis corresponds to approximately 1.88 percentage points lower subsequent three-month return. The six-month full-control coefficient is much less precise, so the evidence is not uniform across horizons. A max- $|t|$ adjustment across the one-, two-, three-, and six-month horizons yields a three-month family-wise p -value of 0.010 for the parsimonious model and 0.091 for the full-control model.

Second, the result is not tied to the actual month-end basis observation. Monthly mean, median, positive-day share, and longest positive streak all produce negative three-month coefficients, whereas the actual month-end basis is weak. Yet this comparison does not identify persistence as a separate economic factor: a penultimate-day observation and short moving averages also contain forecasting information in some exercises. The safer conclusion is that several broader or pre-boundary summaries are more informative than the actual month-end observation in this sample, while persistence, smoothing, and timing noise cannot be cleanly separated.

Third, the main monthly average retains positive pseudo-real-time forecasting performance. A 120-month rolling model produces an OOS R^2 of 3.73% and a one-sided Clark-West p -value of 0.018. Positive-day share and longest streak have larger OOS R^2 values, but those measures were developed during the research process and do not decisively outperform a five-day smoothed basis in direct loss comparisons. Their stronger OOS statistics are therefore treated as exploratory rather than as independent anomaly discoveries.

Fourth, the robustness evidence both strengthens and limits the interpretation. Alternative Korean financing-rate and dividend inputs leave the three-month coefficient negative, and delaying the return start to the next trading-day close yields a coefficient of -9.87 with $t = -2.60$. At the same time, the full-sample in-sample significance is materially amplified by the global financial crisis. Contract construction also matters: rolling from the nearby into the second-nearby contract during the final five trading days reduces the full-control three-month coefficient to -4.70 ($t = -1.67$), although the basis-only t -statistic remains -2.65 and the 120-month rolling OOS R^2 remains 2.82%. Thus the evidence is inconsistent with a simple, invariant trading anomaly.

The contribution is therefore best described as an incremental medium-horizon extension of the KOSPI 200 basis literature with explicit attention to forecast-boundary timing, signal construc-

¹Korea Exchange currently lists the KOSPI regular cash session as 09:00–15:30 and the KOSPI 200 futures regular session as 08:45–15:45. These current hours illustrate the non-synchronicity concern; the paper does not assume identical trading hours throughout 2004–2026.

tion, contract rolling, and pseudo-real-time validation. Relative to Kim (2020), the paper studies substantially longer return horizons. Relative to Hazelkorn, Moskowitz, and Vasudevan (2023), it provides a long-sample Korean implementation using local fair-value inputs and directly documents how the result changes across monthly summaries and nearby-contract roll rules. The paper does not claim a new basis anomaly or identify a causal economic mechanism.

The rest of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data and variable construction. Section 4 presents the empirical design. Section 5 reports baseline predictive results. Section 6 compares alternative basis summaries and pseudo-real-time forecasts. Section 7 presents robustness tests and limitations. Section 8 concludes.

2. Related Literature

2.1 Futures-Cash Basis, Arbitrage, and Liquidity Demand

The no-arbitrage benchmark for equity-index futures is fundamentally simple: the futures contract can be replicated through a financed position in the underlying index portfolio, adjusted for dividends. In a frictionless setting, economically meaningful deviations from the replication value should attract arbitrage capital and disappear quickly.

Real markets are not frictionless. Financing costs, balance-sheet constraints, transaction costs, short-sale restrictions, and liquidity can impede arbitrage. Roll, Schwartz, and Subrahmanyam (2007) show that deviations in the equity-index futures-cash relation are dynamically linked to stock-market liquidity. Their results establish a direct connection between the law of one price and the conditions under which arbitrage can be implemented.

Hazelkorn, Moskowitz, and Vasudevan (2023) broaden the interpretation from pure arbitrage to liquidity demand. Across global equity-index markets, they show that the futures-cash basis co-moves with dealer and investor positions, is positively related to contemporaneous equity returns, and negatively predicts subsequent spot and futures returns. Their results suggest that the basis can summarize demand for equity-market exposure rather than simply record a transitory pricing mistake.

Persistence is already part of this literature. Hazelkorn et al. report persistent supply-demand imbalances reflected in the basis and return predictability beyond very short horizons. Consequently, the present paper does not claim that basis persistence itself is new. Its narrower question is whether simple within-month summaries—such as the monthly average, the share of days with a positive fair-value-adjusted basis, or the longest positive run—are useful for medium-horizon forecasting in the KOSPI 200 market and in pseudo-real-time tests.

2.2 KOSPI 200 Futures and Price Discovery

The Korean index-derivatives literature has long emphasized price discovery. Nam et al. (2006), using minute-level data, find that KOSPI 200 futures lead the cash index and analyze pricing bias across the KOSPI 200 stock-index derivatives markets. Joo, Seon, and Lee (2016) show that overnight trading in KOSPI 200 futures contributes materially to the daily price-discovery process.

These findings establish that the futures market can incorporate information before the cash market. That makes timing particularly important when daily futures and cash observations are combined.

A daily closing basis constructed from non-synchronous prices can mix the economic basis with a short interval of price discovery. The daily end-of-day basis can therefore contain timing noise throughout the sample. The empirical design does not claim to remove that measurement problem. Instead, it focuses on the additional forecast-boundary issue at month-end and uses a next-trading-day return start as the strongest check on direct timing overlap.

2.3 KOSPI 200 Futures Mispricing

Kim (2018) studies the determinants of KOSPI 200 futures mispricing using futures-implied spot prices for nearby and second-nearby contracts. Prior 60-day returns, volatility, and skewness are significant determinants of the measured deviation. This evidence suggests that the spot-futures relation varies with the state of the underlying equity market.

Kim (2020) is closest to the current return-predictability question. The paper constructs futures-implied spot prices using a cost-of-carry framework, studies demand-based explanations for the mispricing, and finds a negative relation between futures mispricing and future cash-index returns. The predictive effect lasts up to approximately two weeks.

Ryu, Ryu, and Yang (2025) examine KOSPI 200 futures mispricing through investor sentiment and limited arbitrage. They find that the sentiment-mispricing relation is stronger when volatility is high and when short selling is constrained, and that regulatory reforms affect the relation.

Together, these papers make clear that neither KOSPI 200 mispricing nor short-horizon predictability is unexplored. The empirical gap addressed here is more specific: medium-horizon predictability, forecast-boundary-aware daily construction, contract-roll sensitivity, and pseudo-real-time forecasting over a sample exceeding two decades.

2.4 Contribution Relative to Prior Work

The paper offers four incremental pieces of evidence.

First, it provides evidence that the KOSPI 200 futures basis contains predictive information at a substantially longer horizon than the short-run reversal emphasized in prior Korean evidence. The full-control relationship is most evident around two to three months and is weak by six months.

Second, it compares multiple ways to aggregate the daily basis. The analysis distinguishes actual month-end observations, short moving averages, monthly mean and median measures, and duration measures. The comparison is descriptive about signal construction; it does not identify persistence as a separate economic factor because smoothing, timing, and duration are difficult to separate with these data.

Third, it evaluates the relation using pseudo-real-time forecasting rules. Predictive regressions can look significant in sample even when they provide little out-of-sample accuracy improvement. Rolling and expanding exercises therefore form a central part of the paper rather than a minor robustness check.

Fourth, it uses a matched second-nearby futures series to quantify contract-roll sensitivity directly. The resulting early-roll tests show that the strongest controlled in-sample coefficient is not invariant to nearby-contract construction, an important limitation that can be measured rather than left as a data caveat.

3. Data and Variable Construction

3.1 Data and Sample

The preferred sample begins in January 2004, when the expected-dividend series required for the main fair-value specification becomes available, and extends through July 2026 for monthly predictor construction. Because a three-month future return must be observed, the latest usable three-month forecast origin is April 2026.

The analysis uses daily observations on:

- the KOSPI 200 cash index;
- the nearby and second-nearby KOSPI 200 futures contracts;
- FY0 and FY1 KOSPI 200 dividend yields;
- nearby-contract expiration dates;
- Korean money-market interest rates;
- nearby- and second-nearby-contract volume and open interest;
- VKOSPI; and
- the USD/KRW exchange rate.

The KOSPI 200 cash index, vendor-defined nearby and second-nearby futures series, dividend yields, contract-expiration dates, futures volume and open interest, and Korean interest-rate series are retrieved through **FnGuide QuantiWise** Excel exports. The workbook identifies the nearby series as KSFA020, the second-nearby series as KSFA021, and the KOSPI 200 index as IKS200; the preferred expected-dividend item is the KOSPI 200 FY1 dividend yield. VKOSPI and USD/KRW are obtained from **Bloomberg L.P.** using VKOSPI Index and KRW Curncy, respectively, with PX_LAST. The analysis uses the historical daily vendor series as supplied for each date and does not independently reconstruct archival vintages of the expected-dividend series. Appendix Table B1 provides the raw identifiers used in the replication pipeline.

The main monthly sample contains 271 month-end-excluded basis observations. The three-month full-control regression contains 267 observations. Table 1 reports the summary statistics used in the paper.

Table 1: Summary Statistics

Variable	N	Mean	Std. Dev.	P5	Median	P95	Min	Max
KOSPI 200 monthly return (%)	270	1.056	6.761	-7.566	0.789	10.746	-23.631	35.344
3M forward KOSPI 200 return (%)	268	3.358	12.168	-11.976	2.098	22.989	-27.739	84.097
Safe monthly mean abnormal basis (%)	271	0.062	0.230	-0.371	0.100	0.374	-1.109	0.739
Safe monthly median abnormal basis (%)	271	0.063	0.256	-0.396	0.100	0.370	-1.456	0.711
Penultimate-day abnormal basis (%)	271	0.122	0.373	-0.411	0.140	0.574	-3.626	1.190
Month-end abnormal basis (%)	271	0.076	0.357	-0.559	0.101	0.516	-1.654	1.395
Positive-basis day share	271	0.655	0.262	0.182	0.700	1.000	0.053	1.000
Longest positive streak share	271	0.424	0.292	0.056	0.333	1.000	0.048	1.000
Realized volatility (%)	271	5.566	3.606	2.709	4.706	12.074	1.781	32.407
Safe monthly mean VKOSPI	271	21.340	10.806	12.015	18.506	38.239	10.754	85.002
Safe month-end USD/KRW change (%)	270	0.120	3.133	-4.221	-0.026	5.184	-14.283	18.112
Mean daily futures volume (contracts)	271	239,938	80,566	134,672	220,664	382,077	102,590	589,760
Mean daily open interest (contracts)	271	169,479	84,494	86,432	115,405	310,489	77,312	350,019

Notes: Monthly basis variables use the fair-value-adjusted end-of-day KOSPI 200 futures basis. The main monthly mean excludes the final trading day. Positive-basis-day share and longest-streak share are fractions between 0 and 1. Volume and open interest are mean daily contract counts.

3.2 Fair-Value-Adjusted Basis

For a futures contract expiring at T , theoretical fair value is

$$F_{t,T}^* = S_t \exp [(r_t - q_t)\tau_{t,T}],$$

where S_t is the KOSPI 200 cash index, r_t is the annualized financing rate, q_t is the annualized expected dividend yield, and $\tau_{t,T}$ is remaining time to expiration expressed in years.

The preferred specification uses the CD 91-day rate for r_t and the FY1 expected KOSPI 200 dividend yield for q_t .

The daily abnormal basis is

$$AB_{t,T} = 100 \frac{F_{t,T} - F_{t,T}^*}{S_t}.$$

I use *abnormal basis* as shorthand for this fair-value-adjusted end-of-day measure. Because the historical cash and futures observations are not synchronized to the exact same intraday timestamp, the variable should not be interpreted as an executable arbitrage profit or as a direct measure of arbitrage profits available to a trader.

For continuous-roll robustness, the same fair-value formula is recomputed using the selected contract's own expiration date. The main alternatives replace the nearby contract with the second-nearby contract during the final five trading days before expiration or when the nearby contract has 7, 10, or 14 calendar days or fewer remaining. The second-nearby series is available throughout the 2004–2026 main sample. In the vendor continuous-series data, second-nearby trading volume rarely exceeds nearby volume before expiration, while open-interest crossing generally occurs only one to two days before expiry; liquidity-cross rules therefore do not provide an economically earlier roll date.

3.3 Forecast-Boundary Timing and the Monthly Signal

Let L_m be the final trading day of month m . The primary predictor is

$$\overline{AB}_m^{safe} = \frac{1}{N_m - 1} \sum_{\substack{t \in m \\ t \neq L_m}} AB_t.$$

I refer to this measure as the timing-safe monthly signal only in the narrow sense of forecast-boundary timing. Cash and futures closes can be non-synchronous on every trading day, so earlier daily basis observations may still contain measurement noise. The final trading day raises an additional concern because the month-end cash close S_m also anchors the subsequent return. Excluding that day's basis ensures that the predictor does not use a potentially later futures close while measuring the future return from the same day's earlier cash close. The conditioning variables in the full regression are measured using information available by month-end. The average fair-value-adjusted basis is 0.062% with a standard deviation of 0.230%; its sample range is -1.109% to 0.739%.

Alternative timing tests exclude the last two, three, and five trading days. A stronger specification moves the future-return starting point to the next trading-day cash close, so the forecast boundary

occurs after all month-end end-of-day observations. Figure 1 plots the month-end-excluded monthly average basis over the full sample.




Figure 1: Month-End-Excluded Monthly Average Fair-Value-Adjusted KOSPI 200 Futures Basis

3.4 Alternative Monthly Basis Summaries

To assess whether the result depends on a single aggregation rule, I construct several monthly summaries of the daily basis.

The monthly mean and median measure the magnitude of the basis across the month.

The positive-day share is

$$PositiveShare_m = \frac{1}{N_m - 1} \sum_{\substack{t \in m \\ t \neq L_m}} \mathbf{1}(AB_t > 0).$$

The longest-streak share is the maximum number of consecutive positive-basis days divided by the number of eligible trading days.

For comparison, I also construct the average basis over the final five and ten eligible trading days of each month. The monthly mean and five-day average are highly correlated (approximately 0.86), and the monthly mean and positive-day share have a correlation of approximately 0.88. Joint-regression evidence must therefore be interpreted with caution.

3.5 Returns and Controls

The future h -month KOSPI 200 return is

$$R_{m,m+h} = 100 \left(\frac{S_{m+h}}{S_m} - 1 \right), \quad h \in \{1, 2, 3, 6\}.$$

The full regressions include the prior one-month KOSPI 200 return, realized volatility, a month-end-excluded monthly VKOSPI average, the corresponding USD/KRW change, and month-of-quarter indicators that control for the recurring futures maturity cycle. Monthly realized volatility is defined as the square root of the sum of squared daily KOSPI 200 percentage returns within the month. The FX control is the percentage change in the last pre-boundary USD/KRW observation of the month relative to the corresponding observation in the preceding month.

4. Empirical Design

4.1 Predictive Regressions

The parsimonious predictive regression is

$$R_{m,m+h} = \alpha_h + \beta_h X_m + \varepsilon_{m,m+h},$$

where X_m is a basis measure.

The full specification is

$$R_{m,m+h} = \alpha_h + \beta_h X_m + \Gamma_h Z_m + \varepsilon_{m,m+h},$$

where Z_m contains the market controls described above.

Because multi-month future returns overlap, inference uses heteroskedasticity- and autocorrelation-consistent standard errors. The implementation uses a minimum HAC lag of three months and increases the lag length with the forecast horizon. The coefficient β_h is interpreted as predictive association, not as a causal effect of the futures basis on subsequent cash returns.

4.2 Pseudo-Real-Time Forecasting

At every forecast origin, the model is estimated using only observations for which the future return would already have been fully realized. The OOS specifications are deliberately parsimonious: each forecast model contains an intercept and one basis predictor, without the full in-sample control set. This avoids estimating a comparatively large control vector in short rolling windows and makes the signal comparison transparent.

The model forecast is compared with a historical-mean return forecast computed from the same information available in the estimation window. Out-of-sample performance is

$$R_{OOS}^2 = 1 - \frac{\sum_m (R_m - \widehat{R}_m)^2}{\sum_m (R_m - \bar{R}_m)^2}.$$

Here $\widehat{R}_m$ is the basis-model forecast and $\bar{R}_m$ is the recursively available historical-mean benchmark. A positive OOS R^2 therefore means that the basis model has a lower sum of squared forecast errors than the benchmark. The paper reports expanding-window, 60-month rolling, and 120-month rolling estimates. Clark-West tests are reported for nested forecast comparisons; the reported

Clark-West p-values are one-sided tests of whether the predictor model improves on the historical-mean benchmark, with HAC(3) inference for the three-month forecast exercise. OOS R^2 measures forecast-error improvement and should not be read as a trading return or investor utility gain.

The forecasting exercises are appropriately described as **pseudo-real-time validation**. The variables themselves were developed during the broader research process, so the later sample is not a pristine untouched discovery sample.

5. Main Results

5.1 Baseline Predictive Regressions

I begin with the month-end-excluded monthly average basis. Its predictive coefficients for future KOSPI 200 returns are negative across the reported horizons.

In the basis-only specifications, the one-, two-, three-, and six-month coefficients are -4.60, -8.15, -11.34, and -10.71. Their HAC t-statistics are -2.62, -2.79, -3.01, and -2.04.

The strongest evidence is at two and three months. The three-month parsimonious coefficient is -11.34 with p-value 0.003.

The full-control results are more conservative. The one-month coefficient is -4.36 ($t = -1.94$), the two-month coefficient is -5.69 ($t = -2.10$), the three-month coefficient is -8.15 ($t = -2.13$, $p = 0.033$), and the six-month coefficient is -5.41 ($t = -0.88$).

The full specification therefore provides conventional-significance evidence around two to three months but not a uniformly significant relation across horizons. In the common-sample family-wise bootstrap reported in Appendix A, the three-month adjusted p-value is 0.010 for the basis-only model and 0.091 with the full control set. Thus the parsimonious result is stronger than the full-control result under the conservative horizon correction. Table 2 reports the complete horizon-by-horizon estimates.

Table 2: Baseline Return Predictability

Horizon	Basis-only β	t	p	N	R^2	Full-control β	t	p	N	R^2
1M	-4.604	-2.618	0.009	270	0.024	-4.364	-1.938	0.053	269	0.041
2M	-8.149	-2.786	0.005	269	0.032	-5.688	-2.104	0.035	268	0.107
3M	-11.345	-3.007	0.003	268	0.045	-8.150	-2.134	0.033	267	0.164
6M	-10.707	-2.037	0.042	265	0.013	-5.409	-0.880	0.379	264	0.118

Notes: The dependent variable is the future KOSPI 200 return over the indicated horizon. Full controls include the prior one-month return, realized volatility, a month-end-excluded mean VKOSPI, the corresponding USD/KRW change, and month-of-quarter indicators. HAC/Newey-West standard errors are used.

A one-standard-deviation increase in the monthly fair-value-adjusted basis is associated with approximately 1.88 percentage points lower subsequent three-month KOSPI 200 return. This is an effect-size calculation from the predictive regression, not evidence of an implementable net return.

5.2 Quintile Evidence

Sorting months into abnormal-basis quintiles reveals an asymmetric descriptive pattern. Average subsequent three-month returns are 6.32%, 5.81%, and 5.84% in the first three quintiles, then fall to 0.36% in the fourth quintile and -1.55% in the highest-basis quintile.

This pattern is consistent with stronger predictability in unusually high-basis states, but it does not establish a statistically distinct threshold or asymmetric model. Threshold specifications explored during the research process do not consistently dominate the linear forecast out of sample. The quintile evidence is therefore descriptive and is not used to define an ex-post trading rule. Figure 2 displays the quintile pattern.




Figure 2: Subsequent Three-Month KOSPI 200 Return by Abnormal-Basis Quintile

6. Alternative Basis Summaries and Pseudo-Real-Time Forecasting

Having established the baseline relation for the monthly average, I next ask whether the result depends on how the daily basis is summarized and whether those alternatives improve forecast accuracy.

6.1 Broad versus Point-in-Time Basis Measures

The full-control three-month coefficient for the monthly mean is -8.15 ($t = -2.13$). The median produces -6.54 ($t = -2.11$). Positive-day share produces -6.56 ($t = -2.04$), while longest positive streak produces -5.92 ($t = -2.05$).

The actual month-end basis behaves differently. Its coefficient is 2.08 with t -statistic 0.70. The penultimate-day basis is also weak in the main in-sample regression, with coefficient -1.64 and t -statistic -0.66.

The contrast should not be overstated. Smoothing alone can improve a noisy daily basis measure. Indeed, the final-five-day average has a basis-only HAC t-statistic of -3.13, and the final-ten-day average has a t-statistic of -2.53. Moreover, the penultimate-day basis is weak in sample but has positive rolling OOS performance. The evidence therefore does not show that month-long persistence is necessary. A narrower statement is that the actual month-end basis is a weak summary in these tests, whereas several broader or pre-boundary summaries produce more stable evidence. Table 3 presents the comparisons.

Table 3: Alternative Monthly Basis Summaries

Panel A. Individual predictor specifications							
Predictor	β	t-stat	p-value	N			
Monthly mean	-8.150	-2.134	0.033	267			
Monthly median	-6.539	-2.109	0.035	267			
Penultimate-day point	-1.642	-0.663	0.508	267			
Month-end point	2.080	0.704	0.481	267			
Positive-basis day share	-6.562	-2.045	0.041	267			
Longest positive streak share	-5.916	-2.048	0.041	267			

Panel B. Monthly mean plus an additional measure							
Joint with monthly mean	Mean β	Mean t	Mean p	Other β	Other t	Other p	N
Penultimate-day point	-8.037	-1.957	0.050	-0.246	-0.098	0.922	267
Month-end point	-9.626	-2.330	0.020	3.739	1.181	0.238	267
Positive-basis day share	-6.907	-1.067	0.286	-1.291	-0.242	0.809	267
Longest positive streak share	-5.988	-1.354	0.176	-2.647	-0.814	0.415	267

Notes: All specifications use the three-month future return and the full control set. Persistence variables are correlated with the monthly mean, so incremental in-sample information should be interpreted cautiously.

6.2 Out-of-Sample Results

I next ask whether the in-sample differences across basis summaries translate into out-of-sample forecast accuracy. The answer is partly yes, but the comparison still does not identify persistence as the source of the signal.

For the monthly mean, the OOS R^2 values are 3.22% for a 60-month rolling window, 3.73% for a 120-month rolling window, and 3.39% for the expanding window. The corresponding Clark-West p-values are 0.031, 0.018, and 0.004.

Under the 120-month rolling design:

- monthly median: 3.03%;
- penultimate-day basis: 5.08%;
- positive-day share: 7.72%;
- longest positive streak: 7.69%;
- actual month-end basis: -0.39%.

The actual month-end basis fails to improve on the historical-mean benchmark. However, the penultimate-day basis generates a positive OOS R^2 , which cautions against interpreting the OOS comparison as proof that duration itself is the relevant economic feature.

A five-day average, motivated by the short-window smoothing used in the international basis literature, generates an OOS R^2 of 1.96%; the ten-day average produces 1.79%. Monthly and duration

measures are numerically stronger. However, a direct forecast-loss comparison against the five-day benchmark yields one-sided p-values of 0.125 for the monthly mean, 0.083 for positive-day share, and 0.064 for the longest streak.

Accordingly, the paper does not claim statistically decisive dominance of the duration measures over a conventional smoothed basis. The duration variables and several alternative summaries were developed during the research process, and no family-wise correction is applied across this full set of predictor definitions. Their OOS statistics should therefore be read as exploratory validation. In this sample, the duration summaries have larger OOS R^2 values than the five- and ten-day averages. Table 4 reports the OOS statistics; Figures 3 and 4 show the cross-model OOS R^2 comparison and cumulative forecast gains, and Figure 5 benchmarks monthly summaries against short-window smoothing.

More generally, the persistence and short-window comparisons should be treated as secondary evidence. The monthly mean is the paper’s main signal; median, positive-day share, longest streak, penultimate-day, month-end, and short-window averages were all examined during the research process. Reported OOS tests preserve the forecasting information set, but they do not undo the ex-post choice of which signal definitions to inspect.

Table 4: Pseudo-Real-Time Out-of-Sample Forecasting

Predictor	Window	OOS R^2 (%)	CW t	CW p	N	Start	End
Monthly mean	60	3.22	1.867	0.031	206	2009-03-31	2026-04-30
Monthly mean	120	3.73	2.098	0.018	146	2014-03-31	2026-04-30
Monthly mean	expanding	3.39	2.639	0.004	206	2009-03-31	2026-04-30
Monthly median	60	2.22	2.088	0.018	206	2009-03-31	2026-04-30
Monthly median	120	3.03	2.324	0.010	146	2014-03-31	2026-04-30
Monthly median	expanding	2.57	2.801	0.003	206	2009-03-31	2026-04-30
Penultimate-day point	60	2.67	1.598	0.055	206	2009-03-31	2026-04-30
Penultimate-day point	120	5.08	1.813	0.035	146	2014-03-31	2026-04-30
Penultimate-day point	expanding	2.23	2.410	0.008	206	2009-03-31	2026-04-30
Positive-basis day share	60	5.25	1.616	0.053	206	2009-03-31	2026-04-30
Positive-basis day share	120	7.72	1.800	0.036	146	2014-03-31	2026-04-30
Positive-basis day share	expanding	4.99	2.413	0.008	206	2009-03-31	2026-04-30
Longest positive streak share	60	5.62	2.040	0.021	206	2009-03-31	2026-04-30
Longest positive streak share	120	7.69	2.181	0.015	146	2014-03-31	2026-04-30
Longest positive streak share	expanding	5.57	2.982	0.001	206	2009-03-31	2026-04-30

Notes: The benchmark is the recursively available historical-mean return forecast. Positive OOS R-squared indicates lower mean squared forecast error than the benchmark. Clark-West p-values are one-sided. These are pseudo-real-time validation exercises, not pristine ex-ante discovery tests; OOS R-squared is a forecast-accuracy metric, not a trading return.

7. Robustness and Limitations

The robustness analysis separates five questions: whether the result depends on fair-value inputs, on the forecast boundary, on crisis episodes or calendar subsamples, on nearby-contract construction, and on selection across forecast horizons. The purpose is not to show that every specification produces the same significance level; several tests instead identify where the baseline relation is fragile.




Figure 3: 120-Month Rolling Out-of-Sample Forecast Performance




Figure 4: Cumulative Three-Month OOS Forecast Gain versus the Historical Mean




Figure 5: Short-Window versus Broader Monthly Basis Summaries

7.1 Alternative Fair-Value Inputs

The preferred construction uses CD 91-day and FY1 expected dividends. The full-control three-month coefficient is -8.15 ($t = -2.13$).

Using FY0 dividends produces -7.65 ($t = -2.18$), while the FY0/FY1 average produces -8.06 ($t = -2.17$).

Alternative rate specifications are also similar. The MSB 91-day construction produces -10.84 ($t = -2.34$); the Korean money-broker call-rate specification produces -12.78 ($t = -2.86$). Within the available Korean rate and dividend proxies, the sign and conventional statistical significance do not depend on a single carry assumption. Table 5 reports the alternative fair-value constructions.

Table 5: Fair-Value Construction Robustness

Construction	β	t-stat	p-value	N
CD91 + FY1	-8.150	-2.134	0.033	267
CD91 + FY0	-7.652	-2.179	0.029	267
CD91 + avg(FY0,FY1)	-8.062	-2.174	0.030	267
MSB91 + FY1	-10.839	-2.336	0.020	236
Call(KMB) + FY1	-12.778	-2.860	0.004	267
Call(Broker) + FY1	-9.780	-2.360	0.018	267
Call(Seoul) + FY1	-9.784	-2.362	0.018	267

Notes: Each row repeats the full-control three-month regression with an alternative financing-rate and/or dividend-yield construction. The MSB91 sample is shorter because of data availability.

7.2 Timing Robustness

Excluding additional month-end observations leaves the coefficient negative.

When the last two trading days are excluded, the coefficient is -7.81 with $t = -2.10$. Excluding three days gives -7.71 with $t = -2.04$, and excluding five gives -7.89 with $t = -2.01$.

The most conservative timing check moves the forecast return start to the following trading-day close. In that specification, the basis coefficient is -9.87, with $t = -2.60$ and $p = 0.009$.

These tests address the forecast-boundary concern rather than the generic non-synchronicity of every daily cash and futures close. Excluding additional month-end days leaves the estimate similar, and moving the return start to the next trading-day close removes the direct overlap between a potentially later month-end futures close and the cash close used to start the future return. Daily basis measurement noise from non-synchronous closes can nevertheless remain elsewhere in the month.

7.3 Crisis Dependence

The global financial crisis matters materially for statistical strength. I define the GFC interval as September 2008 through June 2009 and the COVID interval as March through December 2020. The conservative crisis test removes any forecast origin for which either the predictor month or any month in the subsequent three-month return window intersects the relevant crisis interval.

The full-sample three-month coefficient is -8.15. Under the conservative GFC-overlap exclusion, it falls to -2.56, with $t = -0.90$ and $p = 0.369$. Applying the analogous COVID-overlap exclusion produces -9.66 with $t = -2.47$. Excluding windows that overlap either episode gives -4.17, with $t = -1.45$ and $p = 0.148$.

This is a meaningful limitation: the full-sample in-sample significance is materially amplified by the GFC. The later OOS evidence is not mechanically inherited from GFC observations. From September 2019 onward, the 120-month rolling estimation window has completely rolled beyond the June 2009 end of the GFC interval; over this later evaluation period, OOS R^2 is 3.54% for the monthly mean, 8.07% for positive-day share, and 8.69% for the longest-streak measure. This later interval nevertheless includes COVID and was examined after the predictor definitions had been developed, so it is not a clean crisis-free holdout. The defensible conclusion is that in-sample significance is GFC-sensitive while later forecast performance remains positive after GFC observations leave the estimation window.

7.4 Subperiod Stability

The 2004–2010 coefficient is -9.07 ($t = -1.28$), the 2011–2017 coefficient is -8.74 ($t = -3.21$), and the 2018–2026 coefficient is -5.68 ($t = -0.87$).

All three point estimates are negative, but statistical precision varies substantially. These calendar splits are descriptive rather than estimated structural breaks, and the evidence should be described as time-varying rather than structurally constant.

7.5 Near-Expiry and Continuous-Roll Robustness

The nearby-contract construction is an important source of sensitivity.

A first diagnostic simply removes daily observations close to expiration. Removing observations with five or fewer days to expiration yields a coefficient of -6.23 with $t = -1.95$. Requiring more than seven days gives -4.83 with $t = -1.72$. Requiring more than ten days produces -4.04 with $t = -1.55$, while a fourteen-day exclusion yields -3.82 with $t = -1.49$.

The matched second-nearby contract allows a more direct test that preserves those dates but changes the contract used to compute the basis. I therefore construct continuous-roll signals that

switch from the nearby to the second-nearby futures price before expiration and recompute fair value using the second contract’s own maturity.

The result is clear but not favorable to a strong invariance claim. Rolling during the final five trading days produces a full-control three-month coefficient of -4.70 with $t = -1.67$ and $p = 0.094$. Rolling when nearby time to expiration falls to seven calendar days or less gives -4.25 with $t = -1.60$; the ten- and fourteen-day rules give -3.43 ($t = -1.41$) and -3.29 ($t = -1.36$), respectively.

At the same time, the signal is not eliminated in the simpler or forecasting specifications. Under the final-five-trading-days roll, the three-month basis-only t -statistic is -2.65 and the 120-month rolling OOS R^2 is 2.82% with a one-sided Clark-West p -value of 0.019. The seven-, ten-, and fourteen-day roll rules produce OOS R^2 values of 2.53%, 2.17%, and 2.21%, respectively. These rolled models do not statistically improve upon the front-only forecast; if anything, the front-only model is numerically stronger in this sample.

Using the second-nearby contract for every day provides an additional diagnostic. The monthly second-nearby and nearby basis measures have a correlation of about 0.51. The second-only full-control three-month coefficient is close to zero ($t = -0.17$), although its rolling OOS R^2 is 3.67%. This mixed result reinforces the conclusion that contract maturity and roll mechanics matter for the in-sample relation.

The liquidity data do not supply a useful earlier endogenous roll trigger. Second-nearby volume exceeds nearby volume before or at expiration in only two front-expiry cycles in the main sample, both on the expiration date, while the median open-interest crossing occurs one calendar day before expiry. Accordingly, the paper treats the fixed pre-expiration roll rules as the informative robustness exercise.

The new second-nearby evidence therefore **confirms the near-expiry sensitivity rather than resolving it**. The vendor-defined nearby series remains the paper’s baseline because it is the contract construction used throughout the main analysis, but the weaker controlled coefficients under early roll materially constrain interpretation. Timing, crisis, subperiod, and near-expiry exclusion results are collected in Table 6; explicit continuous-roll results appear in Appendix Table A2.

7.6 Multiple Testing

The paper reports one-, two-, three-, and six-month horizons jointly rather than selecting only the best horizon.

To account for horizon selection, I implement a reproducible two-sided moving-block $\max\{|t|$ procedure across the one-, two-, three-, and six-month horizons. The bootstrap uses a common sample across horizons, residuals from the restricted model under the joint null that the basis coefficient is zero, circular six-month blocks, and 19,999 bootstrap replications. The same resampled residual blocks are used across horizons to preserve their cross-horizon dependence.

In the basis-only family, the common-sample three-month HAC t -statistic is -3.15 and the family-wise adjusted p -value is 0.010. With the full control set, the common-sample three-month t -statistic is -2.30 and the adjusted p -value is 0.091. Table A1 reports all four horizons. This correction addresses selection across the four forecast horizons only; it is not a global correction for every predictor definition, control set, timing rule, or exploratory specification considered during the research process. The parsimonious three-month evidence therefore remains strong within the horizon

Table 6: Timing, Crisis, Subperiod, and Expiry Robustness

Panel A. Timing safeguards

Specification	β	t-stat	p-value	N
Exclude last 1 trading day(s)	-8.150	-2.134	0.033	267
Exclude last 2 trading day(s)	-7.814	-2.095	0.036	267
Exclude last 3 trading day(s)	-7.710	-2.044	0.041	267
Exclude last 5 trading day(s)	-7.886	-2.013	0.044	267
Forecast return starts next trading day (~63d)	-9.870	-2.601	0.009	267

Panel B. Crisis-window exclusions

Specification	β	t-stat	p-value	N
Full	-8.150	-2.134	0.033	267
Ex GFC-overlap window	-2.556	-0.898	0.369	254
Ex COVID-overlap window	-9.664	-2.467	0.014	254
Ex GFC & COVID overlap	-4.165	-1.446	0.148	241

Panel C. Calendar subperiods

Specification	β	t-stat	p-value	N
2004-2010	-9.066	-1.279	0.201	83
2011-2017	-8.739	-3.211	0.001	84
2018-2026	-5.675	-0.872	0.383	100

Panel D. Near-expiry exclusions

Specification	β	t-stat	p-value	N
> 5 days	-6.228	-1.953	0.051	267
> 7 days	-4.830	-1.718	0.086	267
> 10 days	-4.037	-1.552	0.121	267
> 14 days	-3.824	-1.493	0.135	267

Notes: Timing safeguards address non-synchronous cash/futures closes. Crisis tests remove any forecast origin whose predictor month or three-month return window intersects the stated crisis interval. GFC and near-expiry exclusions materially weaken in-sample significance and are reported as limitations rather than omitted.

family, while the full-control result is better characterized as borderline under that conservative family-wise inference.

7.7 Scope of Interpretation

Finally, three qualifications limit what can be inferred from the predictive results. First, the analysis is predictive rather than causal. Without investor-level flow data or another source of identification, it cannot distinguish among liquidity demand, time-varying risk premia, limits to arbitrage, or other mechanisms. Second, the preferred FY1 expected-dividend series is used as supplied in the historical vendor data; the study does not independently reconstruct archival forecast vintages, so any vendor revision or backfill risk cannot be ruled out. Third, the forecasting exercises evaluate squared-error accuracy rather than an implementable portfolio after transaction costs, margins, roll costs, taxes, or risk constraints.

8. Conclusion

This paper asks whether information in the KOSPI 200 futures basis extends beyond the short-horizon return predictability already documented in Korea. Using daily data from 2004 through 2026, I construct a fair-value-adjusted end-of-day basis from Korean financing rates, expected KOSPI 200 dividends, and actual contract maturity, and I evaluate subsequent KOSPI 200 returns over one to six months.

The baseline evidence is concentrated around two to three months. Higher monthly basis values are associated with lower subsequent realized returns, and the three-month coefficient remains negative after controlling for recent returns, realized volatility, VKOSPI, exchange-rate movements, and the maturity cycle. The full-control three-month p-value is 0.033 under conventional HAC inference and 0.091 after a conservative family-wise adjustment across the four reported forecast horizons. The result should therefore be viewed as meaningful but not uniformly strong across specifications or horizons.

The paper also clarifies what the monthly signal does and does not capture. The baseline excludes the final trading day because that day's potentially later futures close could otherwise be combined with the same day's earlier cash close that anchors the future return. This removes a direct forecast-boundary overlap; it does not eliminate non-synchronous-close measurement noise from every daily basis observation. Broader monthly summaries often perform better than the actual month-end observation, but short moving averages and a penultimate-day observation also contain information in some tests. The evidence therefore does not identify persistence or duration as a separate economic factor.

Pseudo-real-time forecasting provides complementary evidence. The monthly average produces positive OOS R^2 under 60-month rolling, 120-month rolling, and expanding estimation windows. Some duration measures are numerically stronger, but they were developed during the research process and do not decisively outperform conventional short-window smoothing in direct loss comparisons. The OOS exercises support forecastability relative to a historical-mean benchmark, not an implementable trading strategy.

The strongest qualifications come from crisis and contract-maturity tests. Excluding GFC-overlap windows materially weakens the full-control in-sample coefficient. Matched second-nearby data

likewise show that rolling out of the nearby contract several days before expiration reduces the magnitude and precision of the controlled coefficient. The parsimonious three-month relation and rolling OOS statistics remain positive under those roll rules, so the signal is not confined entirely to the final days before expiration; nevertheless, the strength of the baseline relation is not invariant to contract construction.

Taken together, the results are best interpreted as incremental evidence that the KOSPI 200 futures basis contains medium-horizon predictive information for subsequent realized returns in this sample. The contribution is not the discovery of a new anomaly and not the identification of a causal mechanism. Rather, it is to document the horizon, timing, signal-construction, forecasting, crisis, and contract-roll conditions under which the predictive relation is stronger or weaker. Whether the underlying economic source is liquidity demand, time-varying risk premia, limits to arbitrage, investor positioning, or another mechanism remains an open question.

Data Availability and Reproducibility

The raw QuantiWise and Bloomberg observations used in this study are proprietary subscription data and cannot be redistributed publicly by the author. The empirical pipeline records the raw series identifiers, transformations, sample rules, and random seed used for bootstrap inference. Given corresponding authorized vendor exports, the code reproduces the transformations and reported results; exact replication may also depend on continued access to the same historical vendor series and vintages. Generated tables, figures, and code contain no vendor credentials.

OpenAI’s ChatGPT was used as an assistant in developing and checking analysis code and in implementing selected robustness procedures. All empirical specifications, data construction, statistical results, and interpretations were reviewed and validated by the author.

Appendix A. Multiple-Horizon Family-Wise Inference

Table A1. Moving-Block Max- $|t|$ Family-Wise Adjustment

The table reports common-sample HAC t-statistics and two-sided family-wise p-values from a circular moving-block residual bootstrap across the 1-, 2-, 3-, and 6-month return horizons. Blocks contain six monthly observations; 19,999 bootstrap samples are drawn with seed 20260824. The restricted model imposes a zero basis coefficient while retaining the control variables in the full-control family. This procedure controls family-wise error across the four forecast horizons within each displayed model family. It does not correct for the broader research process involving alternative predictor definitions, timing rules, or robustness specifications.

Model	Horizon	Common-sample HAC t	FWER-adjusted p
Basis-only	1M	-2.42	0.066
Basis-only	2M	-2.90	0.021
Basis-only	3M	-3.15	0.010
Basis-only	6M	-2.04	0.143
Full controls	1M	-1.55	0.345
Full controls	2M	-2.14	0.128
Full controls	3M	-2.30	0.091
Full controls	6M	-0.88	0.755

Table A2. Second-Nearby and Continuous-Roll Robustness

The table compares the preferred nearby-contract signal with explicit pre-expiration roll rules using the matched second-nearby KOSPI 200 futures series. Fair value is recomputed using the expiration date of the selected contract. OOS statistics use the same parsimonious 120-month rolling forecast design as Table 4.

Contract rule	3M full β	Full t	Full p	3M basis-only t	120M OOS R^2	CW p
Front only	-8.15	-2.13	0.033	-3.01	3.73%	0.018
Roll final	-4.70	-1.67	0.094	-2.65	2.82%	0.019
5 trading days						
Roll at	-4.25	-1.60	0.109	-2.58	2.53%	0.020
DTE ≤ 7						
Roll at	-3.43	-1.41	0.158	-2.41	2.17%	0.025
DTE ≤ 10						
Roll at	-3.29	-1.36	0.173	-2.43	2.21%	0.026
DTE ≤ 14						
Second-nearby only	-0.27	-0.17	0.868	-1.54	3.67%	0.015

Appendix B. Variable Definitions and Data Sources

Table B1. Raw Data Identifiers and Sources

Variable	Raw identifier	Source	Use
KOSPI 200 cash index	IKS200 / close	FnGuide QuantiWise	Cash index and returns
Nearby KOSPI 200 futures	KSFA020, item P100400	FnGuide QuantiWise	Observed nearby futures price
Second-nearby KOSPI 200 futures	KSFA021, item P100400	FnGuide QuantiWise	Continuous-roll robustness
Nearby / second-nearby volume	KSFA020 / KSFA021, item P100800	FnGuide QuantiWise	Liquidity-roll diagnostic
Nearby / second-nearby open interest	KSFA020 / KSFA021, item P101000	FnGuide QuantiWise	Liquidity-roll diagnostic
KOSPI 200 dividend yield FY0	IKS200, item I501000	FnGuide QuantiWise	Alternative dividend input
KOSPI 200 dividend yield FY1	IKS200, item I502100	FnGuide QuantiWise	Main expected-dividend input
Nearby-contract expiration	KSFA020, item P102500	FnGuide QuantiWise	Time to expiration
CD 91-day rate	xarK0IRKSDACD	FnGuide QuantiWise Economy	Main financing rate
MSB 91-day rate	xarK0IRKSDAMSB91	FnGuide QuantiWise Economy	Alternative financing rate
Call rates	xarK0IRCR*	FnGuide QuantiWise Economy	Alternative financing rates
V-KOSPI 200	VKOSPI Index, PX_LAST	Bloomberg L.P.	Volatility control
USD/KRW	KRW Curncy, PX_LAST	Bloomberg L.P.	Exchange-rate control

Declaration of Generative AI and AI-Assisted Technologies

During the preparation of this work, the author used OpenAI’s ChatGPT to assist with code development and checking, manuscript organization, and language editing. The author reviewed and validated the analyses and all substantive content and takes full responsibility for the final manuscript.

References

- Clark, T. E., & West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1), 291–311.
- Hazelkorn, T. M., Moskowitz, T. J., & Vasudevan, K. (2023). Beyond Basis Basics: Liquidity Demand and Deviations from the Law of One Price. *The Journal of Finance*, 78(1), 301–345. <https://doi.org/10.1111/jofi.13198>
- Joo, S. L., Seon, J., & Lee, J. S. (2016). Overnight Trading and Price Discovery over the Course of a Trading Day: Evidence from Stock Index Futures in Korea. *Asia-Pacific Journal of Financial Studies*, 45(3), 463–491. <https://doi.org/10.1111/ajfs.12136>
- Kim, S. (2018). KOSPI200 Past Return, Volatility, Skewness and Mispricing of KOSPI200 Futures: Using KOSPI200 Index Implied by KOSPI200 Futures. *Korean Journal of Financial Studies*, 47(6), 925–945. <https://doi.org/10.26845/KJFS.2018.12.47.6.925>
- Kim, S. (2020). KOSPI200 Index Futures’ Mispricing: Demand-Based Explanation and Its Predictive Power. *Korean Journal of Business & Management*, 34(1), 89–113. <https://doi.org/10.34274/krabe.2020.34.1.005>
- Nam, S. O., Oh, S. Y., Kim, H. K., & Kim, B. C. (2006). An Empirical Analysis of the Price Discovery and the Pricing Bias in the KOSPI 200 Stock Index Derivatives Markets. *International Review of Financial Analysis*, 15(4–5), 398–414. <https://doi.org/10.1016/j.irfa.2006.02.003>
- Newey, W. K., & West, K. D. (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3), 703–708.
- Roll, R., Schwartz, E., & Subrahmanyam, A. (2007). Liquidity and the Law of One Price: The Case of the Futures-Cash Basis. *The Journal of Finance*, 62(5), 2201–2234. <https://doi.org/10.1111/j.1540-6261.2007.01273.x>
- Ryu, D., Ryu, D., & Yang, H. (2025). Investor Sentiment, Mispricing, and Limited Arbitrage in the Futures Market. *Journal of Futures Markets*, 45(8), 879–895. <https://doi.org/10.1002/fut.22594>

Data and Market Documentation

- FnGuide. (2026). *QuantiWise: Corporate and Equity Investment Analysis System*. Product documentation, accessed August 24, 2026.
- Bloomberg L.P. (2026). *Bloomberg Terminal*. Market data for V-KOSPI 200 and USD/KRW, accessed August 24, 2026.
- Korea Exchange. (2026a). *KOSPI Market Trading Hours*. Market regulation documentation, accessed August 24, 2026.

Korea Exchange. (2026b). *KOSPI 200 Futures Product Specifications*. Derivatives market documentation, accessed August 24, 2026.