Available Online: 14 September 2026 DOI: 10.54254/2977-5701/2026.36759

Chinese convertible bonds: portfolio diversification and volatility spread trading with a dynamic percentile framework

Ziyan (Kevin) Tian

Central University of Finance and Economics, Beijing, China

13671292116@163.com

Abstract. Chinese convertible bonds are examined in two settings: portfolio construction and relative-value trading. At the portfolio level, a convertible bond index is compared with the HS300 equity index and a broad bond index over 2005–2024 using return, volatility, Sharpe ratio, maximum drawdown, illustrative allocations, and mean–variance frontiers. The convertible bond index earns a return close to equities with lower volatility and a smaller drawdown, while allocations containing convertibles record higher Sharpe ratios than the stock–bond benchmark. At the trading level, mispricing is defined as the gap between GARCH volatility estimated from the underlying stock and the implied volatility embedded in the bond price. A positive gap indicates that the conversion option is priced at a lower volatility than the stock estimate. Following a static delta-hedged Feikai CB example, fixed market-wide thresholds are compared with bond-specific expanding-window percentile rules based only on information available at each date. With a P50 exit, clean-convergence trades increase from 945 to 1,415 and the timeout share falls from 25.0% to 9.4%. The contrast between CITIC and Wencan shows that spread dispersion, rather than spread level alone, is important for determining whether apparent mispricing can be traded.

Keywords: Chinese convertible bonds, portfolio diversification, GARCH–implied volatility spread, dynamic percentile framework, delta hedging

1. Introduction

A convertible bond combines a debt claim with the holder's right to convert into the issuer's shares. Coupon and principal payments support its value when the stock performs poorly, while the conversion feature preserves part of the upside when the stock rises. Because both pieces matter, valuation cannot be reduced to an ordinary bond calculation. Interest rates and credit quality remain relevant, but so do the share price, expected volatility, conversion terms, and the contractual clauses attached to the issue.

This mixed payoff makes the Chinese market interesting for two different reasons. In an asset allocation, convertible bonds may provide equity participation without reproducing the full risk of an equity index. At the security level, the option component can occasionally trade at a volatility that differs materially from an estimate based on the underlying stock. A hedged position may benefit if that difference later narrows, although the hedge and the cost of implementing it are central to the result.

Two questions guide the empirical work. First, does adding Chinese convertible bonds improve the risk-return choices available to a conventional stock-bond investor? Second, can the gap between GARCH and implied volatility serve as a workable signal for systematic convertible bond trades? The first question is addressed with index statistics and frontier analysis; the second is studied with bond-level signals and backtests.

The tests proceed from broad allocation evidence to individual trades. The long-sample section first compares the three asset indices and constructs two- and three-asset efficient frontiers. The next section studies historical, GARCH, and implied volatility at both the bond and market levels. Finally, the trading section moves from one Feikai CB position to two market-wide rules: an absolute-threshold benchmark and a bond-specific percentile specification.

The allocation results place convertible bonds between equities and straight bonds, but not as a simple average of the two. Their index return is 6.74% a year, compared with 7.18% for the HS300, while annualized volatility is 17.45% rather than 25.76%. The two example portfolios that hold convertible bonds achieve Sharpe ratios of 0.476 and 0.499; the corresponding stock-bond benchmark reaches 0.459. The frontier results point in the same direction by widening the set of attainable risk-return combinations.

The backtests reveal a problem with applying one spread level to the whole market. Under the benchmark rule, entry requires a GARCH-minus-implied spread above 40%, and a profitable position is closed after the spread drops below 10%. Those numbers are easy to interpret, but the same 40% observation can be exceptional for one bond and routine for another because their historical distributions differ widely.

The percentile specification responds to this heterogeneity by using each bond as its own reference. P90 and P50 are recalculated at every eligible date from observations available strictly beforehand, so the backtest does not borrow information from the future. Relative to the fixed rule, the P50 version generates 1,415 rather than 945 clean-convergence trades and cuts the timeout proportion from 25.0% to 9.4%. This change in convergence frequency is the paper's main trading result.

CITIC and Wencan make the cross-sectional issue visible. Their median spreads, 17.9% and 20.1%, are close enough that a ranking based on level would treat them similarly. Their standard deviations are not close: 2.9% for CITIC and 27.1% for Wencan. The first bond remains cheap-looking but changes little; the second repeatedly moves between relatively high and low spread states. In this setting, enough variation to create an exit matters as much as the initial appearance of cheapness.

Implementation is deliberately simplified in the present backtest. Delta is set when a trade opens and is not rebalanced afterward. In the Feikai example, 63.22 of the total 84.81 profit per bond comes from the short stock, while the bond contributes 21.59. The position therefore retains meaningful directionality even though it is described as delta hedged. Rebalancing could isolate volatility convergence more closely, but only with explicit assumptions about borrowing, trading costs, liquidity, and execution.

The remainder of the paper is arranged as follows. Section 2 introduces the contract structure, institutional clauses, data, and volatility definitions. Section 3 reports the portfolio evidence, and Section 4 develops the mispricing measure. Section 5 contains the case study and market-wide backtests. Section 6 considers the limits of the current design and possible extensions, before Section 7 summarizes the findings.

2. Convertible bond background and data

2.1. Convertible bond mechanics

For valuation purposes, a convertible bond is often separated into a straight-bond component and an equity conversion option. The first component provides coupons and repayment at maturity; the second allows the holder to exchange the claim for a predetermined number of shares. With a face value of 100 yuan and a conversion price of 10 yuan, for example, one bond converts into ten shares [1, 2].

Conversion has little appeal when the share price is well below the conversion price, and the security behaves mainly as debt in that region. Once the share price rises sufficiently, the conversion right becomes valuable and the bond takes on more equity-like behavior. This changing sensitivity across stock-price states is the source of its hybrid payoff and much of its potential diversification benefit.

Volatility affects the conversion option for the same asymmetric reason that it affects an ordinary call. Wider stock-price outcomes raise the chance of a valuable conversion, while losses are partly buffered by the debt claim, subject to credit risk. Other inputs unchanged, a higher expected stock volatility should therefore raise both the option component and the convertible bond's value [3].

2.2. Main clauses in Chinese convertible bonds

Actual Chinese issues contain features that make the straight-bond-plus-call description only an approximation. Three clauses are especially relevant: redemption at the issuer's option, a put held by the investor, and a possible downward revision of the conversion price. Each can alter the bond's effective maturity, the likelihood of conversion, and the cash flows that investors expect to receive.

The issuer's redemption clause is commonly associated with forced conversion. A typical trigger requires the stock to close at no less than 130% of the conversion price on at least 15 days within a 30-trading-day window. Once triggered, the issuer may redeem near par plus accrued interest. If conversion value is already far above that amount, holders have a strong incentive to convert instead.

Calling the bond can serve several purposes for the issuer. Conversion replaces debt with equity, removes a future principal payment, and ends the low-coupon financing period. The burden is shifted away from immediate cash interest and toward dilution of existing shareholders. This trade-off helps explain why redemption policy is economically relevant rather than a minor contractual detail.

The investor put works in the opposite direction by allowing the holder to return the bond under specified weak-price conditions, which provides some downside support. A downward revision reduces the conversion price after poor stock performance. It raises the number of shares received on conversion and improves the holder's option position, but also increases potential dilution for shareholders.

These clauses enter the volatility discussion indirectly. Redemption may shorten the life of a bond that has performed well, whereas a revision changes the effective strike of the embedded option. Consequently, implied volatility reflects not only stock uncertainty but also beliefs about issuer behavior and contractual outcomes. The spread used later should be read with that institutional qualification in mind.

2.3. Data and sample construction

Daily observations on a Chinese convertible bond index, the HS300, and a bond index form the allocation sample. Their common period extends from January 2005 through March 2024. Daily returns and normalized price paths are calculated from these series, followed by annualized return and volatility, Sharpe ratio, skewness, kurtosis, and maximum drawdown.

Bond-level tests require closing prices for the convertible and its underlying stock, conversion prices, delta estimates, implied volatility, and GARCH volatility. The market aggregate is available from 2016 to 2024, while the trade records used in the dynamic backtest begin in 2017 and continue into early 2024.

Feikai CB is retained as a transparent worked example before the rule is expanded to the full universe. The resulting files record entry and exit dates, spreads, holding periods, profit by leg, total profit, and exit type. Separate outputs contain the strategy summaries, bond-level percentiles, and the focused comparison between CITIC and Wencan. A complete inventory of the supplementary files is provided in Appendix C.

2.4. Volatility measures and the spread signal

Historical volatility measures the realized dispersion of past stock returns over a chosen window. A GARCH estimate is also based on past returns, but it treats volatility as conditional and persistent, so recent large shocks can influence forecasts for subsequent periods. The estimate used here is annualized before being compared with the bond's implied volatility. The principal variables used in the empirical analysis are summarized in Appendix Table A1.

Implied volatility is obtained in the reverse direction: the observed convertible price is inserted into the pricing model, and the volatility input is solved for. It is therefore neither a directly observed variable nor automatically the best forecast of future realized volatility. Its role is to express the market price of the embedded option in a common volatility unit.

The trading signal used throughout the empirical sections is defined in Equation (1):

$$Spread_{i,t} = GARCH\ Volatility_{i,t} - Implied\ Volatility_{i,t} \tag{1}$$

When the spread is positive, the stock-based $GARCH$ estimate exceeds the $Volatility$ implied by the convertible price, which is consistent with a relatively cheap embedded option. A negative value indicates the reverse. This language describes relative pricing under the model; it does not establish an error in the market price by itself.

The size of the gap must also be judged against the bond's own history. A 20% observation may be rare for one issue yet completely ordinary for another. That cross-sectional difference motivates the move from absolute cutoffs to bond-specific percentiles in Section 5.

3. Portfolio-level evidence

3.1. Asset-level performance

Table 1 provides the first answer to the allocation question. The convertible bond index returns 6.74% annually, only 0.44 percentage points below the HS300's 7.18%, but its volatility is 17.45% rather than 25.76%. On this simple return-to-risk comparison, the convertible bond index consequently has the higher Sharpe ratio.

Drawdowns tell a similar, though less comfortable, story. A maximum loss of 52.02% shows that convertible bonds are far from a low-risk substitute for straight debt. Even so, the decline is smaller than the 72.30% drawdown of the HS300. The bond index lies at the other end of the range, with much lower risk and a lower return. In the sample, convertibles occupy a genuine middle position without closely tracking either benchmark. The normalized price paths in Figure 1 show the same intermediate position over time.

Table 1. Asset and sample-portfolio performance statistics

Asset / Portfolio	Annualized Return (%)	Annualized Volatility (%)	Sharpe Ratio	Maximum Drawdown (%)
Convertible Bond Index	6.74	17.45	0.386	−52.02
HS300 Index	7.18	25.76	0.279	−72.30
Bond Index	4.91	1.54	3.185	−4.28
Portfolio 1: 50% CB / 30% HS300 / 20% Bond	7.16	15.04	0.476	−46.3
Portfolio 2: 30% CB / 40% HS300 / 30% Bond	7.15	14.34	0.499	−46.4
Portfolio 3: 0% CB / 60% HS300 / 40% Bond	7.09	15.44	0.459	−50.3

Note: Values are rounded from the project calculations. The Sharpe ratio is calculated under a zero risk-free-rate assumption as annualized return divided by annualized volatility. Sample period: January 2005 to March 2024. Source: author's calculations based on the project index data.

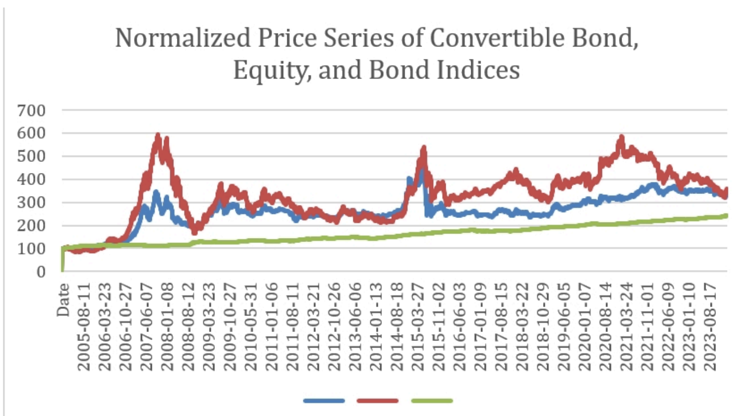

Figure 1. Normalized price series of convertible bond, equity, and bond indices

Note: All three indices are normalized to 100 at the beginning of the common sample. Sample period: January 2005 to March 2024. Source: author's calculations based on the project index data.

3.2. Sample portfolios

Three simple allocations illustrate what this middle position contributes. Portfolio 1 assigns 50% to convertible bonds, 30% to the HS300, and 20% to bonds. Portfolio 2 uses weights of 30%, 40%, and 30%, respectively. Portfolio 3 serves as the comparison portfolio and contains only 60% HS300 and 40% bonds.

Average performance is almost unchanged across the three choices: annualized returns range only from 7.09% to 7.16%. Risk is where the allocations separate. Both portfolios containing convertibles have lower volatility and higher Sharpe ratios than Portfolio 3, with Portfolio 2 reaching the highest ratio, 0.499. The contribution of the new asset is thus better balance rather than a large increase in average return.

This comparison should not be read as unconditional dominance. Straight bonds remain the least volatile asset in the sample, and equities may offer more participation in a strong bull market. What convertibles add is a payoff pattern that the particular stock and bond indices do not reproduce on their own [4].

3.3. Mean–variance frontier

The frontier analysis removes the need to rely on only three selected portfolios. Expected portfolio return is the weighted mean of asset returns, and variance is calculated from the weight vector and the sample covariance matrix. For every target return, the exercise searches for the feasible weights that produce the lowest variance while still summing to one [5]. The calculations are summarized by Equations (2) and (3).

$$Portfolio\ Return = w'\mu \tag{2}$$

$$Portfolio\ Variance = w'\Sigma w \tag{3}$$

Figure 2 compares a frontier built from the HS300 and bond index with one that also permits an allocation to convertible bonds. Adding the third asset expands the feasible region. The difference is most relevant at moderate and higher target returns, where the three-asset set reaches similar expected returns with less volatility. This pattern is consistent with the sample portfolios rather than being driven by one chosen set of weights.

Because the bond index has exceptionally low sample volatility, the global minimum-variance solution is naturally concentrated in that index. The result is mechanically sensible but not very informative about the role of convertibles. More economically useful is the part of the frontier away from that corner. There, convertible bonds offer another combination of equity sensitivity and downside support and improve the menu available to investors who require more return.

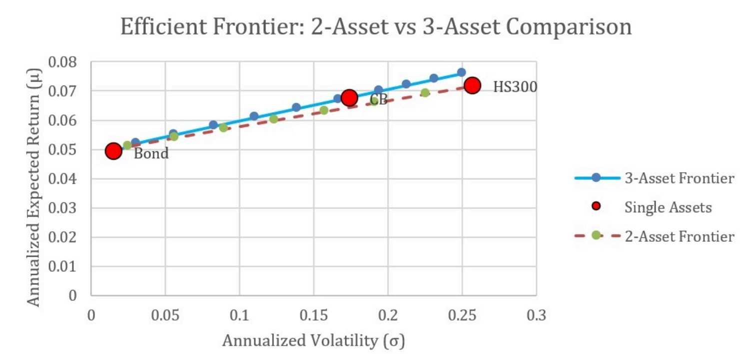

Figure 2. Mean–variance frontier for two-asset and three-asset portfolios

Note: The horizontal axis reports annualized volatility and the vertical axis reports annualized expected return. The two-asset frontier uses the CSI 300 and bond indices, while the three-asset frontier also includes the convertible bond index. Sample period: January 2005 to March 2024. Source: author's calculations.

4. Volatility mispricing signal

4.1. Historical, GARCH, and implied volatility

Portfolio analysis treats the convertible bond as one asset; a trading strategy instead needs to look inside it. Since the conversion right becomes more valuable as stock volatility rises, the observed bond price can be translated into an implied volatility. That quantity can then be compared with a volatility estimate formed from returns on the underlying stock.

Historical volatility summarizes realized variation over a fixed past window. GARCH uses the same broad source of data but allows the conditional variance to change over time and to reflect volatility clustering. Implied volatility comes from the current bond price instead. The spread therefore compares information contained in the stock's return history with the valuation currently placed on the conversion option [6, 7].

Neither side of this comparison is a frictionless measure of fair value. The GARCH estimate changes with model specification and estimation error. The implied number can absorb credit conditions, liquidity, contract clauses, and investor demand in addition to expected stock volatility. Despite these qualifications, using one consistently defined spread makes unusually wide valuation gaps easier to locate and compare [8, 9].

4.2. Individual-bond evidence

Bond-level series confirm that the spread behaves very differently across issues. A large bank convertible such as CITIC may remain at a positive spread for an extended period without much movement. Higher-beta industrial issues can cross repeatedly between relatively cheap and expensive states, while other bonds stay near the middle and offer few clear signals.

These patterns matter before any backtest is run. If one bond stays close to 20% while another ranges from -40% to +50%, an identical entry threshold is measuring two different events. The bond-level evidence therefore suggests that unusualness should be defined relative to an issue's own distribution.

4.3. Market-wide spread evidence

To describe the market as a whole, the GARCH-minus-implied spread is first calculated for every bond observed on a given date and then averaged with equal weights. The resulting series is a broad indicator of option cheapness: a large positive value means that many bonds carry implied volatilities below the corresponding stock-based estimates.

The average remains positive through much of 2016-2024. Its widest gap occurs around 2018-2019, reaching about 29% in June 2018. A 120-trading-day moving average makes the persistence clearer: spreads stay relatively high through much of 2018-2020 and then narrow after 2021. Figure 3 establishes this time variation, but the figure alone cannot identify which economic forces caused each regime.

At the same time, the observed universe grows from only a small number of bonds in 2016 to more than 500 in 2024. Market expansion alongside lower average spreads is consistent with increased attention and maturity, although that explanation is not tested directly. The safer conclusion is that the opportunity set changes through time, which makes a permanent market-wide cutoff difficult to defend.

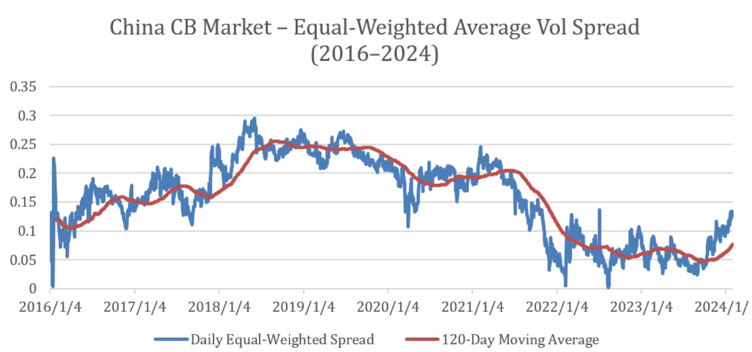

Figure 3. Market-wide average GARCH-minus-implied volatility spread, 2016–2024

Note: The daily series is the equal-weighted average of GARCH volatility minus implied volatility across available convertible bonds. The orange line is the 120-trading-day moving average. Sample period: January 2016 to January 2024. Source: author's calculations based on the project volatility data.

4.4. From the signal to a hedged trade

A sufficiently positive spread suggests buying the convertible because the option appears inexpensive relative to the stock-based volatility estimate. Shorting the underlying stock offsets part of the position's first-order equity sensitivity. The gap may then close through a rise in implied volatility, a fall in the GARCH estimate, or movement in both measures [10].

Without the stock position, a long convertible is also a directional bet on the issuer's shares. The purpose of the hedge is to reduce that directionality so that relative valuation plays a larger role in PnL. The Feikai case in the next section shows how the position is constructed and, just as importantly, why a hedge fixed only at entry remains imperfect.

5. Dynamic percentile-based arbitrage strategy

5.1. Single-bond case study: Feikai CB

Feikai CB (123078.SZ) provides an unusually clear case because its positive spread is both large and easy to connect to the mechanics of a static hedge. On 25 April 2022, GARCH volatility was close to 59% and implied volatility was about 12.5%, leaving a spread of roughly 46.5 percentage points. Under the model, the conversion option therefore appeared inexpensive relative to the stock's conditional volatility estimate.

The position buys one convertible bond and shorts shares of the issuer. A conversion price of 15.56 implies a conversion ratio of approximately 6.43 shares for each 100 yuan of face value. Multiplying that ratio by the entry delta of about 0.80 gives a short position of roughly 5.15 shares per bond. The hedge calculation is shown in Equation (4).

$$Shares\ Shorted = (100/ConversionPrice) \times Delta \tag{4}$$

Entry prices were 129.41 for the convertible and 125.32 for the stock. The 5.15-share hedge is left unchanged until exit, which makes this a static rather than dynamically rebalanced position. Profit on the bond leg is its price change; profit on the short leg is the decline in the stock price multiplied by the number of shares sold short. Adding the two gives total PnL.

The path between entry and exit is more informative than the final profit alone. When both securities initially rose, the loss on the short shares exceeded the gain on the bond and produced a substantial interim drawdown. PnL recovered only after the stock declined. In other words, an entry-date delta calculation did not keep the position neutral as market conditions changed. Figure 4 displays this cumulative PnL path from entry to exit.

Exit was determined by convergence rather than by a preset calendar date. On 1 November 2022, the spread had declined to approximately 2.25%, and the position was closed. Total profit reached 84.81 per bond: 21.59 on the convertible leg and 63.22 on the short-stock leg. That decomposition becomes important when the strength of the hedge is assessed in Section 6.

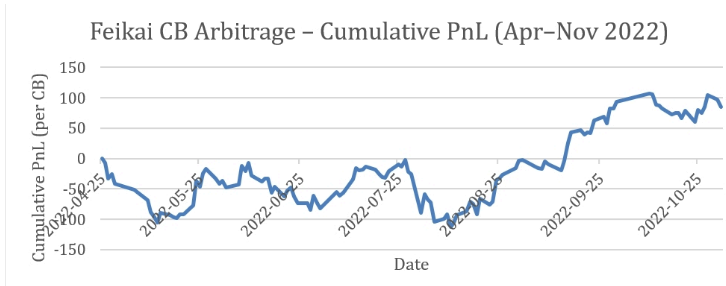

Figure 4. Feikai convertible bond arbitrage: cumulative PnL, April–November 2022

Note: The figure reports cumulative profit and loss for a static delta-hedged position consisting of a long position in Feikai CB and a short position in the underlying stock. The trade was opened on April 25, 2022 and closed on November 1, 2022. PnL is reported in RMB per convertible bond. Source: author's calculations.

5.2. Fixed-threshold full-market backtest

The first market-wide specification deliberately keeps the rule simple and applies identical cutoffs to every issue. A spread above 40% opens a position. The position is closed once the spread is below 10% and cumulative PnL is positive. If neither condition is satisfied within 252 trading days, the observation is classified and closed as a timeout. Appendix B summarizes the full backtest rule details.

This benchmark is useful precisely because its economic logic can be read directly from the thresholds: buy after the embedded option becomes markedly cheap by the chosen measure and sell after much of the gap disappears. It also supplies a common reference against which the percentile rule can be evaluated.

Its simplicity comes at a cost. Forty percent may be an extreme observation for a stable bond that almost never reaches it, but an unremarkable value for an issue whose spread routinely swings through a wide range. The benchmark discards both the usual level and the dispersion of each bond's own series, so equivalent-looking signals need not carry equivalent information.

5.3. Dynamic percentile framework

The dynamic rule replaces the universal cutoff with each bond's expanding historical distribution. Trading is not allowed until at least 250 valid prior spread observations are available. From then onward, P90 and P50 are recomputed on every date using all valid observations strictly earlier than that date. The estimation sample therefore expands as time passes; it is not restricted to the latest year.

Entry occurs when the current spread reaches or exceeds the bond's contemporaneous P90, placing it among roughly the highest 10% of that issue's observed history. In the principal specification, exit requires a spread at or below P50 together with positive total PnL. The same 252-day stop remains in force. The alternative P20 test keeps P90 for entry and changes only the exit condition, thereby demanding a deeper convergence before closing.

Timing of the percentile calculation is essential. Full-sample percentiles would reveal the future shape of the distribution to trades made early in the sample. Instead, a decision dated t uses only records from before *t*. This restriction makes the historical simulation closer to information that could actually have been available and removes the most direct form of look-ahead bias.

The thresholds also have an intuitive reading. P90 marks unusually high cheapness for the bond in question, even if its absolute spread never approaches 40%. P50 represents a return to the middle of its history, while P20 waits for the spread to move into a low historical state. Because the benchmark is bond specific, a routine observation in a high-dispersion issue is less likely to be mistaken for an exceptional opportunity. The entry and exit rules are formalized in Equations (5) and (6).

$$Entry : Spread_{i,t} \geq P90_{i,t} \tag{5}$$

$$Exit : Spread_{i,t} \leq P50_{i,t}; Total\ PnL > 0 \tag{6}$$

5.4. Dynamic P50 versus fixed threshold

Relative to the fixed rule, dynamic P50 produces a substantial change in how trades are completed. As shown in Table 2, clean-convergence exits rise from 945 to 1,415, an increase of about one half. As a share of all positions, signal exits increase from 56.0% to 71.4%.

Timeouts move in the opposite direction, falling from 25.0% to 9.4%. Signal trades also remain open for an average of 34 days instead of 42. Their mean PnL is lower, 65.4 versus 76.2 per bond, which is consistent with accepting median convergence sooner and releasing capital rather than waiting for a larger absolute move.

The improvement is therefore about the reliability and timing of convergence, not a larger gain on every completed trade. More positions reach a valid exit, fewer capital allocations become stale, and the average signal holding period is shorter. These statistics support better turnover under bond-specific cutoffs; they should not be described as evidence of a change in market liquidity. Figure 5 visualizes the increase in clean-convergence trades and the decline in the timeout share.

Table 2. Fixed-threshold and dynamic P50 strategy comparison

Metric	Fixed Threshold	Dynamic P50
Clean-convergence trades	945	1,415
Signal-exit share (%)	56.0	71.4
Average signal PnL (RMB per bond)	76.2	65.4
Average signal holding days	42	34
Timeout share (%)	25.0	9.4

Note: The fixed strategy uses +40% entry and +10% exit thresholds. Dynamic P50 uses each bond's expanding-window P90 entry and P50 exit. Source: author's backtest outputs.

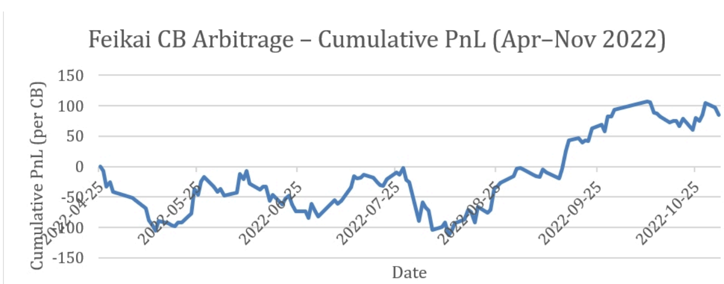

Figure 5. Fixed-threshold vs dynamic P50 strategy performance

Note: Bars report clean-convergence trades. The line reports timeout share. Source: author's backtest outputs.

5.5. Exit rule comparison: P50 versus P20

Table 3 compares the P50 and P20 exit rules. Because both versions continue to enter at P90, the comparison isolates the choice of exit. Waiting for P20 allows the spread to travel further and raises average signal-trade

PnL from 65.4 to 78.1. The additional convergence is slow, however: mean holding time increases from 34 to 47 trading days.

The cost of waiting also appears in unsuccessful exits. The timeout rate rises from 9.4% under P50 to 16.8% under P20, and the real-trade win rate falls from 90.7% to 84.5%. Once signal and timeout trades are pooled, average PnL barely differs between the two rules: 35.5 for P50 and 35.2 for P20.

P50 is consequently retained as the main specification. It gives up some profit conditional on a successful signal exit, but achieves quicker turnover, fewer timeouts, and a higher win rate. P20 remains a meaningful alternative for an investor prepared to commit capital longer in pursuit of deeper convergence.

Table 3. P50 and P20 exit-rule comparison

Metric	P50 Exit	P20 Exit
Average signal PnL (RMB per bond)	65.4	78.1
Average signal holding days	34	47
Timeout share (%)	9.4	16.8
Real-trade win rate (%)	90.7	84.5
Signal + timeout average PnL (RMB per bond)	35.5	35.2

Note: P50 is selected as the main exit rule because it gives faster turnover, fewer timeouts, and a higher win rate. Source: author's dynamic-backtest outputs.

5.6. Cross-sectional validation: CITIC versus Wencan

Table 4 reports the CITIC-Wencan comparison and puts the dispersion argument into a concrete pair. Median spreads are 17.9% and 20.1%, respectively. A cross-sectional screen based only on median cheapness would therefore place the two bonds fairly close together.

Their time-series behavior is almost the opposite. CITIC's spread standard deviation is only 2.9%, near the low end of the sample, whereas Wencan's is 27.1%, near the high end. CITIC stays cheap-looking but changes little. Wencan moves across a much broader range, creating repeated high- and low-spread states in which entries and exits can occur. Figure 6 makes the contrast between similar median spreads and very different dispersion visible.

This pair clarifies why a high spread level cannot be treated as a complete arbitrage signal. Convergence requires movement after entry; a spread that remains permanently high may tie up capital without supplying an exit. Bond-specific percentiles make use of both position within the historical distribution and the distribution's width, which a single fixed threshold cannot do.

Table 4. CITIC and Wencan spread distribution

Bond	Median Spread (%)	Spread Standard Deviation (%)	Interpretation
CITIC	17.9	2.9	Persistently cheap but low movement
Wencan	20.1	27.1	Large swings and repeated trade windows

Note: Median spread levels are similar, but dispersion is very different. Source: author's percentile and comparison files.

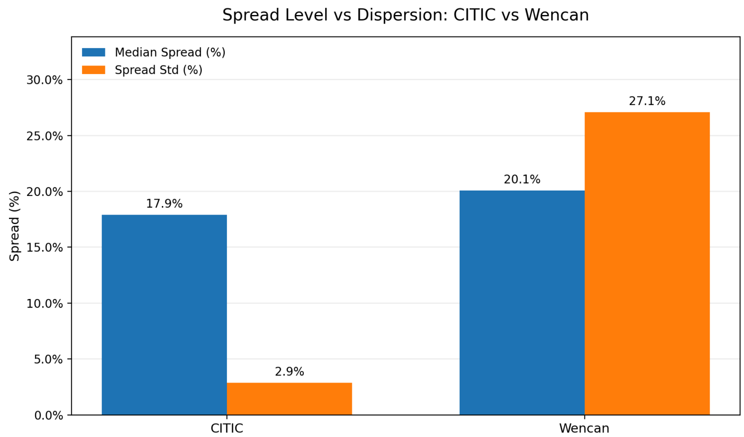

Figure 6. Spread level and dispersion: CITIC vs Wencan

Note: Median spread and spread standard deviation are calculated from each bond's historical spread series. Source: author's calculations.

6. Limitations and future research

6.1. Static hedge and residual stock exposure

The backtest fixes delta at entry and keeps the stock position unchanged until exit. This makes the signal easy to study but does not maintain continuous neutrality. Delta changes with the share price, moneyness, estimated volatility, time to maturity, and sometimes the conversion terms, so the initial hedge can drift into an under- or over-hedged position.

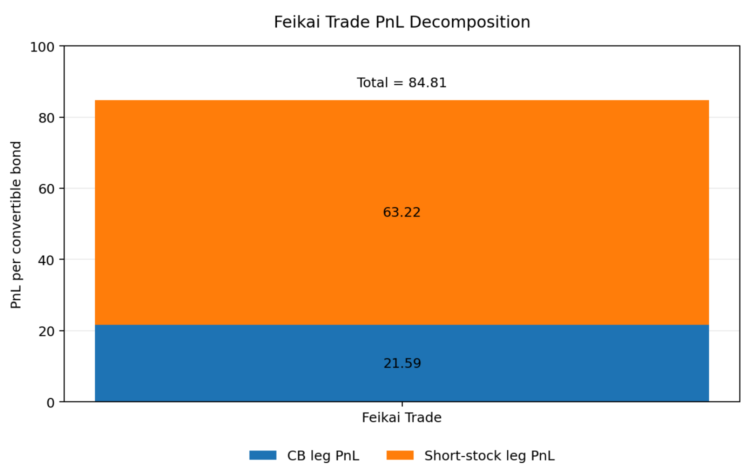

Figure 7. PnL decomposition of the Feikai convertible bond trade

Note: Total PnL of 84.81 consists of 21.59 from the convertible bond leg and 63.22 from the short-stock leg. Source: author's calculations.

Feikai's decomposition makes the residual exposure visible. Of the 84.81 earned per bond, 21.59 comes from the long convertible and 63.22 from the short shares. Most of the realized gain therefore accompanies the later stock decline, so the case should not be presented as a pure estimate of volatility-convergence alpha. The decomposition is shown in Figure 7.

Its more useful role is diagnostic: the spread identifies a relative-value opportunity, while the decomposition shows that the chosen implementation still depends on equity direction. Results using the same hedge design are therefore evidence about signal behavior under a stylized execution rule, not a fully isolated volatility-arbitrage return series.

6.2. Transaction costs and market frictions

Reported PnL does not yet deduct the full set of implementation costs: commissions, financing, stock-borrow fees, bid-ask spreads, slippage, and market impact. Their importance is likely to vary by issue and date, and they would be especially consequential for a strategy that rebalances often. A dynamic hedge may improve exposure control while simultaneously giving up part of that improvement through higher turnover.

Borrow availability is a separate constraint rather than just another fee. In the Chinese market, some underlying stocks cannot be borrowed consistently, and the available quantity and rate can change during a position's life. Thin trading in either leg can also make the closing price an optimistic execution assumption. Signals that are valid in the dataset may therefore be infeasible or much less profitable in practice.

There is model uncertainty on both sides of the spread as well. GARCH volatility depends on the return equation, estimation window, and distributional assumptions. Implied volatility inherits assumptions about interest rates, credit risk, liquidity, and the treatment of contractual clauses in the convertible pricing model. For this reason, the spread is used as an empirical ranking signal, not claimed to be an exact measure of mispricing.

6.3. Future research: dynamic delta hedging and trade filters

A natural next test would rebalance the stock hedge each day. Doing so should reduce residual equity exposure and provide a cleaner assessment of whether spread convergence itself generates returns. Such a backtest would need more than a sequence of deltas: daily borrow availability and fees, trading-cost schedules, rebalancing rules, and treatment of illiquid observations would all have to be specified jointly. Extending those assumptions to the full universe is therefore a separate empirical exercise rather than a small edit to the current model.

Tradeability filters are a second extension. Volume, bid-ask spread, borrow feasibility, and turnover capacity could be imposed before a signal is accepted. The CITIC-Wencan evidence also suggests using spread dispersion itself: an issue with a high but nearly fixed spread may be less attractive than one that repeatedly crosses its own entry and exit percentiles.

Finally, robustness could be assessed across alternative volatility forecasts and convertible pricing specifications. Historical volatility, different GARCH models, or option-implied measures derived from the equity market would provide competing benchmarks. The present paper keeps the completed empirical design fixed and treats these comparisons as future work.

7. Conclusion

Taken together, the results show two distinct uses for Chinese convertible bonds. In the allocation sample, the index earns nearly the HS300 return with lower volatility and a smaller maximum drawdown. Portfolios that include convertibles improve on the benchmark Sharpe ratio, and allowing the asset in the mean-variance problem expands the feasible risk-return set.

At the security level, the GARCH-minus-implied volatility spread offers a consistent, though model-dependent, way to rank apparent option cheapness. The market average is positive for long periods but not stable across time. Its changing level, together with strong differences among individual issues, argues against treating one absolute number as a permanent definition of opportunity.

The strongest backtest result comes from replacing that absolute number with bond-specific history. Dynamic P50 raises clean-convergence trades from 945 to 1,415 and lowers the timeout share from 25.0% to 9.4%. CITIC and Wencan explain the mechanism: nearly equal median spreads can coexist with radically different dispersion and, therefore, very different chances of completing a convergence trade.

These findings remain conditional on a simplified implementation. A hedge fixed at entry leaves residual stock exposure, and the Feikai decomposition shows that the short-stock leg accounts for most of the observed profit. Full trading and financing costs, reliable access to stock borrow, and liquidity constraints are also absent from the current return calculations. A cost-aware dynamic hedge is the clearest next step for separating volatility convergence from equity direction.

The central conclusion is therefore narrower than a claim that every large spread is arbitrageable. A signal becomes more informative when its level is judged against the same bond's past and when the distribution contains enough movement to support both entry and exit. In this market, relative history provides more useful trading information than a single cutoff applied to all issues.

References

[1] Brennan, M. J., & Schwartz, E. S. (1977). Convertible bonds: Valuation and optimal strategies for call and conversion. *The Journal of Finance, 32*(5), 1699–1715. https://doi.org/10.1111/j.1540-6261.1977.tb03364.x

[2] Ingersoll, J. E., Jr. (1977). A contingent-claims valuation of convertible securities. *Journal of Financial Economics, 4*(3), 289–321. https://doi.org/10.1016/0304-405X(77)90004-6

[3] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy, 81*(3), 637–654. https://doi.org/10.1086/260062

[4] Lummer, S. L., & Riepe, M. W. (1993). Convertible bonds as an asset class: 1957–1992. *The Journal of Fixed Income, 3*(2), 47–56. https://doi.org/10.3905/jfi.1993.408078

[5] Markowitz, H. (1952). Portfolio selection. *The Journal of Finance, 7*(1), 77–91. https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

[6] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica, 50*(4), 987–1007. https://doi.org/10.2307/1912773

[7] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31*(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

[8] Ammann, M., Kind, A., & Wilde, C. (2003). Are convertible bonds underpriced? An analysis of the French market. *Journal of Banking & Finance, 27*(4), 635–653. https://doi.org/10.1016/S0378-4266(01)00256-4

[9] Luo, X., & Zhang, J. (2019). Pricing Chinese convertible bonds with default intensity by Monte Carlo method. *Discrete Dynamics in Nature and Society, 2019*, 1–8. https://doi.org/10.1155/2019/8610126

[10] Agarwal, V., Fung, W. H., Loon, Y. C., & Naik, N. Y. (2011). Risk and return in convertible arbitrage: Evidence from the convertible bond market. *Journal of Empirical Finance, 18*(2), 175–194. https://doi.org/10.1016/j.jempfin.2010.11.008

Appendix

Appendix A. Variable definitions

Table A1. Main variable definitions

Variable	Definition
Historical volatility	Annualized standard deviation computed from the underlying stock's past realized returns.
GARCH volatility	Time-varying annualized conditional volatility estimated from the underlying stock-return process.
Implied volatility	Volatility input that makes the convertible pricing model reproduce the observed market price.
Volatility spread	Difference between GARCH volatility and implied volatility.
P90 / P50 / P20	Issue-level percentiles updated from an expanding sample of previously observed spreads.
Signal trade	Position closed after reaching the spread condition while total PnL is positive.
Timeout trade	Position closed at the maximum holding period because the signal-exit condition was not met.
CB leg PnL	Price change on the long convertible bond holding.
Stock leg PnL	Gain or loss generated by the short position in the underlying shares.
Total PnL	Combined PnL from the convertible bond and stock legs.

Note: Definitions used in the empirical analysis.

Appendix B. Backtest rule details

Fixed-threshold rule: open a position once the spread exceeds 40%. Close it when the spread is below 10% and total PnL is positive; otherwise, apply a timeout after 252 trading days.

Dynamic P50 rule: after 250 valid prior observations, estimate expanding-window P90 and P50 for each bond using only past data. Enter at or above P90, exit at or below P50 when total PnL is positive, and retain the 252-trading-day timeout.

Dynamic P20 comparison: keep P90 as the entry percentile but require the spread to reach P20 for a signal exit. The rule seeks deeper convergence at the cost of longer average holding periods and a greater timeout rate.

Static hedge: The number of shares shorted is calculated according to Equation (4) at entry and held unchanged until the position closes.

Appendix C. Supplementary files

- CB_Stock_Bond_index.xlsx – index-level statistics and mean–variance frontier inputs.
- CB_Task4_Data_v2.xlsx – individual-bond volatility comparison data.
- market_spread_timeseries.xlsx / .csv – market-wide average spread series.
- Feikai_Trade.xlsx – Feikai trade setup, daily data, and PnL calculations.

- task5d_backtest_v2.py and task5d_trades.csv – fixed-threshold full-market backtest.
- task5e_dynamic_backtest.py – dynamic percentile backtest.
- task5e_dynamic_trades_P50.csv and task5e_dynamic_trades_P20.csv – trade-level dynamic outputs.
- task5e_dynamic_summary.csv – P50 and P20 summary metrics.
- cb_percentile_table.csv – bond-specific spread percentiles and dispersion statistics.
- citic_wencan_compare.csv – CITIC and Wencan comparison.