



MODULE 1

Who's Paying the Spread

A markout is the signed price change between a fill and the BBO midpoint at a later horizon τ , measured in basis points and interpreted from the liquidity provider's perspective — positive markout means the market moved against the LP after the fill, i.

Section A: Foundations & Concepts – Understanding "Who's Paying the Spread"

Welcome to Module 1! In this section, you will lay the groundwork for understanding the critical dynamics of liquidity provision and transaction costs in financial markets. We'll explore key concepts that reveal how market movements impact different participants and, crucially, "Who's Paying the Spread."

The Concept of Markout: Measuring Price Movement Post-Trade

When you execute a trade, the market doesn't stand still. Prices constantly fluctuate. A fundamental concept for assessing the impact of these movements on a trade is the **markout**. Essentially, a markout measures the change in price between the time a trade is filled and a later point in time.

Key Concept: Markout

A markout is the price change between a fill and the market mid-price at a later horizon τ , signed so that adverse selection (price moving against the liquidity provider) yields a positive value from the taker's perspective.

Understanding markouts is crucial for liquidity providers (LPs), as it directly relates to their profitability. For a bid fill (meaning an LP bought from a taker), the LP's Profit & Loss (PnL) markout can be expressed with a specific formula:

$$\text{PnL_LP}(\tau) = \delta b + (m\tau - m0)$$

Here, δb represents the initial spread capture, and $(m\tau - m0)$ is the markout itself—the change in the market mid-price (m) from the time of the fill ($m0$) to a later horizon ($m\tau$). This formula helps quantify the true profitability or loss after accounting for subsequent market movement.

Exercise 1: Defining Markout in Your Own Words

Based on the definition, how would you explain "markout" to someone new to trading, focusing on its purpose?

Your answer:

Markouts aren't always calculated the same way. The method depends on what you want to analyze. Common ways to calculate markouts include:

MARKOUT CALCULATION BASIS	DESCRIPTION
Mid-to-mid	Measures the change in the market mid-price over a specific period, often used for overall market volatility analysis.
Trade-to-mid (fill-to-mid)	Measures the price change from the actual fill price to the market mid-price at a later horizon τ . This is what's used in the LP PnL formula.
Trade-to-trade (fill-to-cross)	Measures the price change from the fill price to the price of a subsequent opposing trade (e.g., a buy fill to a sell fill).

Exercise 2: Applying the LP PnL Markout Formula

Imagine a liquidity provider fills a bid order. The initial spread capture (δb) was \$0.02. At the time of the fill, the market mid-price (m_0) was \$100.00. One minute later (τ), the mid-price (m_τ) has moved to \$99.95. Calculate the LP's $PnL_{LP}(\tau)$ for this bid fill.

Your answer:

The 'Winner's Curse' and Adverse Selection

The concept of markouts is closely tied to the risk faced by liquidity providers, often referred to as the **market maker's 'winner's curse'**. This happens when a market maker gets their order filled, but the market immediately moves against their position, leading to a loss.

Key Concept: Market Maker's 'Winner's Curse'

This is the risk of being hit at the bid just before the market goes down, or at the ask just before the market goes up. It's a prime example of adverse selection.

This "curse" highlights **adverse selection**—the situation where one party in a transaction has more or better information than the other. If you, as a liquidity provider, are consistently "picked off" right before the market moves unfavorably, it suggests you might be trading against more informed participants.

Exercise 3: Identifying the 'Winner's Curse'

Describe a scenario from your own experience or imagination where a market maker (or any trader) might fall victim to the 'winner's curse'. What was the trade, and how did the market move immediately after?

Your answer:

Uninformed Traders and Market Noise

Not all traders are "informed." Many market participants are simply making trades for reasons unrelated to predicting short-term price movements. These are often called **uninformed (noise) traders**. For such traders, the expectation is that their trades do not, on average, predict future price movements.

Key Concept: Uninformed (Noise) Traders

For uninformed traders, the expected mid-price change after a trade is zero: $E[m_\tau - m_0 \mid \text{uninformed}] = 0$. This means their trades don't carry predictive information about future price direction.

This distinction is crucial for liquidity providers. If you are primarily trading against uninformed flow, your markouts should, on average, be close to zero, allowing you to profit from the spread. If you are constantly experiencing negative markouts (from your perspective), you might be facing informed flow.

Exercise 4: The Impact of Uninformed Trading

You are a market maker. You execute 100 trades against what you believe are purely "noise" traders over a short period. Based on the concept of uninformed traders, what would you expect the average markout ($m\tau - m_0$) to be across these 100 trades, and why?

Your answer:

Effective vs. Realized Spread: Decomposing Execution Cost

To truly understand "Who's Paying the Spread," you need to differentiate between various measures of transaction costs. Two key measures are the **effective spread** and the **realized spread**.

Key Concept: Effective Spread

Effective spread measures the cost of an execution relative to the mid quote at the time of the trade. It tells you the immediate cost of trading.

Key Concept: Realized Spread

Realized spread measures effective spread net of subsequent price movement, separating execution cost from post-fill information/selection effects. It shows the true cost after accounting for adverse selection.

The relationship between these two spreads is highly informative. If you, as a trader, buy an asset and the mid-price falls shortly after, your realized spread will be higher than your effective spread. This difference strongly suggests that you were likely "picked off" or adversely selected—meaning the market moved against you immediately after your trade, implying the counterparty might have had better information.

Exercise 5: Calculating Effective Spread

You place a market buy order and it gets filled at \$50.05. At the exact moment of your trade, the market mid-price was \$50.00. What is your effective spread for this trade?

Your answer:

Exercise 6: Realized Spread and Adverse Selection

Continuing from Exercise 5: One minute after your buy fill at \$50.05, the market mid-price has fallen to \$49.98. Calculate your realized spread. What does the relationship between your effective spread and realized spread tell you about this trade?

Your answer:

Checklist: Analyzing Your Trade Costs

- ☐ Did I calculate the effective spread to understand the immediate cost?
- ☐ Did I calculate the realized spread to account for post-trade price movements?
- ☐ Did I compare effective and realized spread to identify potential adverse selection?
- ☐ If buying, and $\text{Realized Spread} > \text{Effective Spread}$, did I consider if I was "picked off"?

VPIN (Toxicity) Levels: Gauging Informed Flow

To proactively manage the risk of adverse selection, liquidity providers often monitor indicators of order flow toxicity. One such indicator is **VPIN**. VPIN stands for Volume-Synchronized Probability of

Informed Trading and is a measure of order flow imbalance that can signal the presence of informed trading.

A high VPIN level suggests that a significant portion of the trading volume is driven by informed participants, making it riskier to provide liquidity. There's a critical threshold:

Key Concept: VPIN (Toxicity) Threshold

A VPIN (toxicity) level above 45% indicates dominant informed flow, at which point providing liquidity is inadvisable.

Monitoring VPIN can help you decide when to widen your spreads, reduce your size, or even temporarily withdraw from providing liquidity to avoid substantial losses from adverse selection.

VPIN Toxicity Levels and Liquidity Provision

Below 45%

Advisable

Above 45%

Inadvisable

Exercise 7: Responding to VPIN Signals

You are a liquidity provider monitoring VPIN levels for a particular stock. You observe the VPIN level rise from 30% to 55% over a short period. Based on your understanding, what actions might you consider taking, and why?

Your answer:

Effective Sample Size (n_{eff}) for EWMA: Statistical Power

When analyzing market data, especially time-series data like prices or order flow, you often use statistical estimators that give more weight to recent observations. An **Exponentially Weighted Moving Average (EWMA)** is a common tool for this. To understand the reliability of such an estimator, you need to consider its statistical power.

The **effective sample size (n_{eff}) of an EWMA** quantifies this power. It represents the number of independent observations that would yield the same statistical precision as your EWMA estimator. This quantity is determined solely by the decay parameter (α).

Key Concept: Effective Sample Size of EWMA

The effective sample size of an EWMA with decay parameter α is $n_{\text{eff}} = \frac{2}{\alpha}$, a quantity fixed by α alone that quantifies the statistical power of the estimator independently of fill count.

A larger n_{eff} indicates a more statistically powerful estimator, meaning it's based on more "effective" data points. The decay parameter α controls how quickly older observations lose their weight. A smaller α gives more weight to older data (larger n_{eff}), while a larger α makes the EWMA more responsive to recent changes (smaller n_{eff}).

Exercise 8: Understanding the Decay Parameter (α)

You are setting up an EWMA to track short-term price volatility. You have two choices for the decay parameter α : 0.1 and 0.5. Which α value would give you a higher effective sample size (n_{eff}), and what does that imply about the estimator's reliance on historical data?

Your answer:

Exercise 9: Connecting Concepts – Markout and Realized Spread

Explain how the concept of a "markout" is directly integrated into the calculation and interpretation of the "realized spread."

Your answer:

Exercise 10: Synthesizing Risks for a Liquidity Provider

As a liquidity provider, how do the 'winner's curse' and a high VPIN level both indicate a similar underlying risk? What is that risk, and why is it problematic?

Your answer:

Section B: Principles & Frameworks

To truly understand "Who's Paying the Spread," you need to grasp the fundamental theories and frameworks that govern market dynamics. This section breaks down key concepts, from the impact of informed trading to how market makers manage risk and calculate their profit and loss.

Understanding Adverse Selection

In financial markets, not all participants have the same information. Some traders possess superior insights into future price movements. This imbalance gives rise to a critical concept: **adverse selection**.

Key Concept: Adverse Selection

Adverse selection occurs when a counterparty has information about future price direction. For a market maker, this specifically means trading against participants whose order flow reflects superior information about future price dynamics. You risk being hit at the bid just before the market falls or at the ask just before it rises – this is often called the 'winner's curse'.

When an *informed trader* executes a buy order, their knowledge implies that the market price is likely to move in their favor. The expected change in the mid-price after such an informed buy can be described by a specific relationship:

$$E[m_t - m_0 \mid \text{informed buy}] = \mu(1 - e^{-\lambda t})$$

Here, μ (mu) represents the **long-run price impact**, indicating how much the price is expected to move in the long term due to the informed trade. λ (lambda) is the **signal decay rate**, which tells you

how quickly that new information is incorporated into the market price. Both μ and λ are positive values ($\mu > 0$ and $\lambda > 0$).

Exercise 1: Identifying Adverse Selection

Think of a recent situation where you might have bought something, only for its value to drop shortly after, or sold something just before its value surged. How does this personal experience relate to the 'winner's curse' faced by a market maker?

Your answer:

Exercise 2: Interpreting Informed Buy Impact

An informed trader makes a buy. The market's mid-price is m_0 . Using the formula $E[m_\tau - m_0 \mid \text{informed buy}] = \mu(1 - e^{-\lambda\tau})$, explain in your own words what μ and λ represent in this context.

Your answer:

Market Maker Profit and Loss (PnL)

Market makers aim to profit by capturing the **spread** – the difference between the bid and ask prices. However, adverse selection directly impacts their profitability. Your PnL as a liquidity provider (LP) isn't simply the spread captured; it also depends on how the market moves after your trade, especially if you traded with an informed party.

The **unified LP PnL formula** accounts for fills across both bid and ask:

$$\text{PnL}_{\text{LP}}(\tau) = \delta - |m_\tau - m_0| \cdot 1[\text{adverse}]$$

In this formula, δ (delta) is the **half-spread captured**. The second part, $|m_\tau - m_0| \cdot 1[\text{adverse}]$, represents the loss due to adverse selection, where $1[\text{adverse}]$ is an indicator function that is 1 if

adverse selection occurs and 0 otherwise. This shows that your captured spread can be eroded, or even turned into a loss, if the market moves against you due to superior information held by your counterparty.

Consider a specific scenario: you, as a market maker, buy one unit at the bid price ($m_0 - \delta^b$). Your PnL at a horizon τ (tau) is calculated differently:

$$\text{PnL}_{\text{LP}}(\tau) = \delta^b + (m_\tau - m_0)$$

Here, δ^b is the half-spread captured on the bid side. Adverse selection, in this specific case, occurs when the price falls after you buy. If m_τ is less than m_0 , then $(m_\tau - m_0)$ is negative, reducing your PnL or causing a loss.

Exercise 3: Calculating PnL with Adverse Selection

You, as a market maker, capture a half-spread (δ) of \$0.05. An informed trader buys from you, and the mid-price subsequently moves by \$0.07 against your position. Assume adverse selection occurred, meaning $|m_\tau - m_0| = \$0.07$. Calculate your PnL using the unified LP PnL formula.

$\text{PnL}_{\text{LP}}(\tau) =$

Exercise 4: Analyzing PnL for a Bid Fill

You buy one unit at the bid price ($m_0 - \delta^b$). Your captured half-spread on the bid is $\delta^b = \$0.03$. After your buy, the mid-price drops by \$0.04 ($m_\tau - m_0 = -\$0.04$). Calculate your PnL for this specific scenario and state whether adverse selection occurred.

Your PnL:

Adverse Selection (Yes/No):

PNL SCENARIO	FORMULA	DESCRIPTION
Unified LP PnL	$\text{PnL}_{\text{LP}}(\tau) = \delta - m_{\tau} - m_0 \cdot 1[\text{adverse}]$	General formula, applies to both bid/ask fills, explicitly shows adverse selection impact on captured half-spread.
LP PnL (Buy at Bid)	$\text{PnL}_{\text{LP}}(\tau) = \delta^b + (m_{\tau} - m_0)$	Specific to buying one unit at the bid; adverse selection occurs if the price falls after the buy.

Optimal Spread and Risk Management Frameworks

Given the risks of adverse selection and inventory imbalances, how do market makers set their spreads optimally? Frameworks like the **Avellaneda-Stoikov model** provide an answer. This model proposes that an optimal spread is not a single, static number but rather a dynamic value that accounts for different types of risk.

Key Concept: Avellaneda-Stoikov Optimal Spread

In the Avellaneda-Stoikov model, the optimal spread decomposes into two primary terms: an **inventory risk premium** and an **adverse selection component**. This framework helps market makers balance the desire to capture spread with the need to manage risks associated with holding inventory and trading against informed participants.

The **inventory risk premium** accounts for the risk associated with holding an unbalanced position (too much buy-side or sell-side inventory). The **adverse selection component** directly addresses the risk of trading with informed counterparties, as discussed earlier. By adjusting these components, a market maker can dynamically set a spread that reflects current market conditions and their own risk appetite.

Exercise 5: Decomposing Your Spread

Imagine you are a market maker. You set your spread based on the Avellaneda-Stoikov model. Describe how you would qualitatively adjust the two components of your optimal spread (inventory risk premium and adverse selection component) if you suddenly observe a significant increase in informed trading activity in the market.

Your answer:

Exercise 6: Real-World Inventory Risk

Think of a retail business that sells seasonal items. How is their risk of holding unsold inventory (analogous to inventory risk premium for a market maker) managed as the season progresses? How might they adjust their "spread" (i.e., pricing strategy) to mitigate this risk?

Your answer:

Measuring and Adapting to Spread Dynamics

Understanding theoretical frameworks is crucial, but practical application requires robust measurement and adaptive strategies. Regulatory bodies often define how to measure the effectiveness of trading venues and the quality of execution.

One such measurement is the **realized spread**. Under **SEC Rule 605**, for sell orders, the realized spread is calculated as double the difference between the midpoint of the consolidated best bid and offer five minutes after execution and the execution price. This provides a snapshot of how much of the quoted spread was actually "realized" by the market maker after short-term price movements settle.

Evolution of Spread Understanding

- **[[DATE-UNKNOWN]]** Avellaneda-Stoikov Model introduces optimal spread components (inventory risk, adverse selection).
- **[[DATE-UNKNOWN]]** SEC Rule 605 establishes standardized calculation for realized spread for sell orders.

When you are building systems that adapt to market dynamics, especially those involving estimates like Exponentially Weighted Moving Averages (EWMA), you need to ensure these estimates are reliable. The **warmup factor** helps prevent premature action based on insufficient data.

The warmup factor is calculated as: $\text{warmup} = 1 - (1 - \alpha)^n$, where 'n' is the number of fills processed so far, and 'α' (alpha) is a smoothing constant. This factor scales a t-statistic toward zero when few fills have been processed. In essence, it prevents your system from acting confidently on immature EWMA estimates, ensuring you have enough data before making significant adjustments.

Exercise 7: Applying SEC Rule 605

You executed a sell order at \$100.00. Five minutes after execution, the midpoint of the consolidated best bid and offer is \$99.90. Calculate the realized spread for this sell order according to SEC Rule 605.

Realized Spread =

Exercise 8: The Purpose of the Warmup Factor

Imagine you've just started a new trading algorithm. Why would it be detrimental to act on EWMA estimates with very few fills? How does the warmup factor ($\text{warmup} = 1 - (1 - \alpha)^n$) address this problem?

Your answer:

Exercise 9: Scenario with Warmup Factor

You are using a system that applies a warmup factor. If 'n' (number of fills) is very low (e.g., 1 or 2), will the warmup factor be closer to 0 or 1? What does this imply about the confidence your system places in its estimates?

Your answer:

Checklist: Assessing Information Risk

- ☐ Do you have a clear understanding of what constitutes an "informed trade" in your market?
- ☐ Can you identify potential 'winner's curse' scenarios in your trading activity?
- ☐ Are you able to distinguish between adverse selection losses and other types of trading losses?
- ☐ Do your spread-setting mechanisms implicitly or explicitly account for an adverse selection component?
- ☐ Do you use metrics like realized spread to evaluate execution quality?

Exercise 10: Reflecting on Principles

Based on what you've learned in this section, briefly explain how the concept of adverse selection fundamentally influences how a market maker sets their prices (their spread).

Your answer:

Exercise 11: Connecting the Dots

How does the 'signal decay rate' (λ) from the informed trader formula relate to the 'adverse selection component' in the Avellaneda-Stoikov model? Briefly describe the conceptual link.

Your answer:

Section C: Practical Application

You've explored the foundational concepts of "Who's Paying the Spread." Now, let's put that knowledge into action. This section focuses on the practical methodologies for calculating, interpreting, and assessing the reliability of markouts in real-world trading scenarios.

Key Concept: Markout

A markout is the realized PnL on a trade, measured at a specific time horizon (τ) after execution. Alternatively, it can be defined as the price change within an arbitrary time interval before or after an event, such as a fill or a lit trading signal. It helps you understand how the market moved relative to your trade.

Understanding Liquidity Provider (LP) PnL

As a liquidity provider (LP), understanding your profit and loss (PnL) on a trade is critical. The markout concept provides a clear way to measure this. Let's look at the core formula.

Key Concept: LP PnL Formula

For an ask fill (selling), your LP PnL at horizon τ is calculated as: $\text{PnL}_{\text{LP}}(\tau) = \delta a - (m_\tau - m_0)$. Here, δa represents the half-spread captured, and $(m_\tau - m_0)$ is the change in the mid-price after your fill (m_τ is the mid-price at horizon τ , and m_0 is the mid-price at the time of the fill). For a market maker selling one unit, δa is the spread captured and $(m_\tau - m_0)$ is the adverse price move.

Exercise 1: Deconstructing the LP PnL Formula

Imagine you are a liquidity provider. You've just filled an ask order. Identify what each component in the formula $PnL_{LP}(\tau) = \delta a - (m\tau - m0)$ represents in the context of your trade.

- δa :

- $m\tau$:

- $m0$:

- $(m\tau - m0)$:

Your answer:

Exercise 2: Calculating LP PnL for an Ask Fill

You filled an ask order. The half-spread captured (δa) was \$0.02. At the time of the fill (m_0), the mid-price was \$100.00. Five minutes after the fill (m_τ), the mid-price was \$100.01.

Calculate your $PnL_{LP}(5 \text{ minutes})$ using the formula $PnL_{LP}(\tau) = \delta a - (m_\tau - m_0)$.

Your answer:

The result of your PnL_{LP} calculation tells you whether your trade was profitable in hindsight. This is where the concept of a positive or negative markout comes in.

Key Concept: Interpreting Markouts

A **positive markout** means your trade was profitable in hindsight. For buy orders, this means the mid-price rose after the fill. For sell orders, it means the mid-price fell. Conversely, a **negative markout** indicates the market moved against you, suggesting possible trading against informed flow.

Exercise 3: Interpreting LP PnL Outcomes

Based on the LP PnL you calculated in Exercise 2, was your trade profitable or unprofitable in hindsight? What does this imply about the market's movement after your fill?

Your answer:

While you, as an LP, focus on your PnL, it's also important to understand the perspective of the "taker" – the party who initiated the trade with you.

Key Concept: Taker Markout

Markouts are often reported from the taker's perspective as $MO(\tau) = -PnL_{LP}(\tau)$. A positive taker markout means the LP lost money, and the price moved in the taker's favor after the trade.

Exercise 4: Shifting Perspective: From LP to Taker

Using the LP PnL result from Exercise 2, calculate the markout from the taker's perspective. What does this positive or negative taker markout tell you about the market's movement relative to the taker's trade?

Your answer:

MARKOUT TYPE	LP PNL	TAKER MARKOUT	MID-PRICE MOVEMENT (BUYS)	MID-PRICE MOVEMENT (SELLS)	INDICATION
Positive Markout	Profitable	Negative	Rose after fill	Fell after fill	Successful trade / Market moved in LP's favor
Negative Markout	Unprofitable	Positive	Fell after fill	Rose after fill	Trading against informed flow / Market moved against LP

Real-World Application: SEC Rule 605

The concept of measuring price movement after a trade isn't just theoretical; it's codified in regulations. SEC Rule 605 provides a specific methodology for calculating "realized spread" for certain orders, which is a direct application of markout principles.

Key Concept: SEC Rule 605 Realized Spread

Under SEC Rule 605, for buy orders, the **realized spread** is calculated as double the difference between the execution price and the midpoint of the consolidated best bid and offer five minutes after the time of order execution. This effectively measures how much the market moved in the buyer's favor (or against them) after the trade.

Exercise 5: Understanding SEC Rule 605 Components

You are analyzing a buy order for compliance with SEC Rule 605. What key pieces of information do you need to calculate the realized spread as defined by the rule?

Your answer:

Exercise 6: Calculating Realized Spread under SEC Rule 605

You executed a buy order at \$50.00. Five minutes after execution, the midpoint of the consolidated best bid and offer was \$50.05.

Calculate the realized spread for this buy order according to SEC Rule 605.

Your answer:

Checklist: Analyzing Trade Profitability

- ☐ Have you calculated the LP PnL for your trade using the formula?
- ☐ Have you interpreted the PnL (positive for profitable, negative for unprofitable)?
- ☐ Did you consider the mid-price movement ($m_t - m_0$) to understand the underlying market behavior?
- ☐ Have you understood the markout from the taker's perspective and what it implies for them?
- ☐ For buy orders, have you applied SEC Rule 605 to calculate the realized spread where applicable?

Measuring Markout Reliability: EWMA and t-statistic

Markouts provide valuable insights, but how reliable are these insights over time? To get a clearer picture, especially with noisy market data, you often need to smooth your observations and assess the statistical significance of your findings.

Key Concept: Exponentially Weighted Moving Average (EWMA)

The **EWMA** is a type of moving average that places greater weight on more recent observations. It's calculated using the formula: $EWMA_{n+1} = x_{n+1} \cdot \alpha + EWMA_n \cdot (1-\alpha)$, where x_{n+1} is the new observation, α is the smoothing parameter (between 0 and 1), and $EWMA_n$ is the previous EWMA value.

Exercise 7: Step-by-Step EWMA Calculation

Suppose your initial EWMA for markouts was 0.005. Your next markout observation (x_{n+1}) is 0.008. You've chosen a smoothing parameter (α) of 0.2.

Calculate the new EWMA ($EWMA_{n+1}$).

Your answer:

Exercise 8: Why Use EWMA?

In the context of analyzing markouts, explain why using an EWMA might be more beneficial than a simple moving average, especially when market conditions can change rapidly.

Your answer:

Key Concept: Markout Reliability (t-statistic)

The reliability of a markout series can be assessed using a t-statistic, analogous to a one-sample t-test. The formula is: $t = n_{\text{eff}} \times c \times \text{warmup}$. Here, **n_eff** provides degrees of freedom (representing the effective number of independent observations), **c** provides the standardized effect size (how strong the signal is), and **warmup** corrects for any bias introduced by the EWMA's initial values.

Exercise 9: Deconstructing the Markout t-statistic

When assessing the reliability of your calculated markouts using the t-statistic ($t = n_{\text{eff}} \times c \times \text{warmup}$), briefly explain the role of each component:

- **n_{eff} :**

- **c :**

- **warmup:**

Your answer:

Exercise 10: Interpreting Markout Reliability

You've calculated a high positive t-statistic for a series of markouts from your trading strategy. What does this high t-statistic indicate about the profitability and reliability of your strategy's markouts?

Your answer:

Checklist: Evaluating Markout Data Reliability

- ☐ Have you smoothed your markout data using an Exponentially Weighted Moving Average (EWMA) to account for recent market dynamics?
- ☐ Have you identified the effective number of observations (n_{eff}) for your markout series?
- ☐ Have you determined the standardized effect size (c) of your markouts?
- ☐ Have you accounted for any EWMA initialization bias using the 'warmup' factor?
- ☐ Have you calculated the t-statistic for your markout data to assess its statistical reliability?
- ☐ Are you able to interpret the t-statistic to determine the confidence level in your markout findings?

Evolution of Markout Definitions

Initial Definition Realized PnL on a trade measured at some horizon τ after execution (Fact 4)

Broader Definition Price change within some arbitrary time interval before or after an event, such as a fill or a lit trading signal (Fact 8)

Welcome to Section D of Module 1! In "Who's Paying the Spread," understanding the mechanics of trading is just the beginning. Now, it's time to learn how we analyze and measure the impact of the spread, particularly focusing on how to detect and quantify **adverse selection**. This section will walk

you through the key metrics and analytical tools used to evaluate execution quality and understand the true cost of trading.

Understanding Markouts: The Core of Execution Analysis

To truly understand who's paying the spread, you need a way to measure the market's reaction to your trades. This is where **markouts** come in. A markout is a fundamental metric used in quantitative finance.

Key Concept: Markout

A markout is defined as the price change within some arbitrary time interval before or after an event, such as a fill or a lit trading signal. Markouts are commonly used as a measurement of execution quality and adverse selection.

In simpler terms, you execute a trade (or observe a signal), and then you look at how the price moves shortly after. If the price moves against your position, that's a sign of adverse selection – meaning you might have traded with someone who had better information.

Exercise 1: Identifying a Markout Scenario

Think of a recent situation where you bought or sold something (even outside of finance, like a house or a car). Describe how you might define a "markout" for that transaction to see if the price moved favorably or unfavorably for you shortly after the deal was done.

Your answer:

The Liquidity Provider's Markout Perspective

For a **liquidity provider (LP)**, the markout is critical for calculating their profit and loss (PnL). The general markout formula from the LP's perspective helps them understand if they captured the spread successfully or if adverse selection eroded their gains.

The formula is: $PnL_{LP}(\tau) = \delta - |m_{\tau} - m_0| \cdot 1[\text{adverse}]$

- δ represents the **half-spread captured**, which is the immediate profit the LP makes from providing liquidity.
- The second term, $|m_{\tau} - m_0| \cdot 1[\text{adverse}]$, is the adverse selection component. It is nonzero only when the price moves against the LP's position. Here, m_{τ} is the mid-price at a later time τ , and m_0 is the mid-price at the time of the trade. If m_{τ} moves unfavorably compared to m_0 , this term quantifies the loss.

Exercise 2: Interpreting LP Markout Components

Imagine you are a liquidity provider. You just sold a share for \$100. The half-spread captured (δ) was \$0.02. Shortly after, the mid-price moved to \$100.05. Using the formula, explain what each part of the adverse selection term ($|m_{\tau} - m_0| \cdot 1[\text{adverse}]$) represents in this specific scenario and what it means for your PnL.

Your answer:

Temporal Evolution: The Markout Curve

Adverse selection isn't a static event; it evolves over time. To capture this dynamic nature, markouts are not just computed at one point but at multiple time horizons.

Key Concept: Markout Curve

Markouts are computed at multiple time horizons (e.g., every 5 seconds for 30 seconds after each trade, or at 1s, 5s, 30s, 1min, 5min) to build a markout curve showing the temporal evolution of adverse selection.

By measuring markouts at different intervals – from very short (1 second) to longer (5 minutes) – you can see how the price impact of a trade unfolds and how persistent any adverse selection might be.

Typical Markout Horizons for Analysis

Horizon Description

1s Captures immediate, very short-term price impact

5s Shows initial price discovery post-event

30s Reveals short-to-medium term market reaction

1min Indicates a more consolidated market impact

5min Reflects longer-term price trends and sustained adverse selection

Exercise 3: Why Multiple Horizons Matter

You are analyzing your trading performance. You observe a positive markout at the 1-second horizon, but a negative markout at the 5-minute horizon for the same set of trades. What might this suggest about the nature of the adverse selection you are experiencing, and why is it important to look at both short and long horizons?

Your answer:

Quantity-Weighted Markouts and VPIN

While a simple markout looks at price change, a more sophisticated approach considers the volume associated with those changes. This is where **quantity-weighted markout** becomes useful.

Quantity-weighted markout uses two Exponentially Weighted Moving Average (EWMA) accumulators:

- $S \leftarrow (1-\alpha)S + \alpha \cdot (m \cdot q)$
- $Q \leftarrow (1-\alpha)Q + \alpha \cdot q$

With $Mq = S/Q$ giving the volume-weighted exponential moving average markout in basis points. Here, 'm' is the mid-price and 'q' is the quantity traded. This method gives more weight to larger trades, which often have a greater impact on the market.

Exercise 4: When to Use Quantity-Weighted Markout

Consider a market where a few very large trades occasionally occur, interspersed with many small trades. Explain why using a quantity-weighted markout might give you a more accurate picture of adverse selection than a simple markout average, especially if you are concerned about the impact of significant order flow.

Your answer:

Another important metric for understanding market imbalance and potential adverse selection is **VPIN** (Volume-Synchronized Probability of Informed Trading).

Key Concept: VPIN

VPIN is calculated as: $| \text{Buy Volume} - \text{Sell Volume} | / \text{Total Volume}$, averaged over N volume buckets. A high VPIN value suggests a significant imbalance between buy and sell volumes, which can indicate increased order flow toxicity and the presence of informed trading.

VPIN helps you identify periods where there's a strong directional pressure in the market, often associated with informed trading that can lead to adverse selection for liquidity providers.

Exercise 5: Markouts vs. VPIN

You are analyzing a trading strategy. You notice that your markouts are consistently negative, indicating adverse selection. Separately, you observe that VPIN for the asset you trade is frequently high. How do these two observations complement each other in painting a picture of the market conditions and your strategy's performance?

Your answer:

Adverse Selection in Optimal Spread Formulas

The concept of adverse selection isn't just for after-the-fact analysis; it's also built into theoretical models for determining optimal trading strategies, such as setting your bid-ask spread. For example, in the Avellaneda-Stoikov optimal spread formula, there's a specific term dedicated to adverse selection.

The adverse selection component of the Avellaneda-Stoikov optimal spread is $(2/\gamma)\ln(1 + \gamma/\kappa)$. Here, γ (**gamma**) is the **risk-aversion parameter**, and κ (**kappa**) is an **order arrival parameter**. This term helps an automated market maker widen their spread to account for the risk of trading with informed participants.

PARAMETER	DESCRIPTION	ROLE IN OPTIMAL SPREAD
γ (gamma)	Risk-aversion parameter	Reflects your willingness to take on risk; higher γ leads to wider spreads to mitigate risk from adverse selection.
κ (kappa)	Order arrival parameter	Relates to the rate at which orders arrive; influences how much the spread needs to widen to account for order flow.

Exercise 6: Risk Aversion and Optimal Spread

If you, as a market maker, become significantly more risk-averse (meaning your γ parameter increases), how would you expect your optimal bid-ask spread to change according to the Avellaneda-Stoikov model, specifically considering the adverse selection term?

Your answer:

Evaluating and Summarizing Adverse Selection

Once you've computed markouts and other metrics across various horizons, you need a way to interpret and summarize this data. Two key concepts help you do this: consistency and a composite score.

Key Concept: Consistency Metric

The consistency metric $c = \text{mean} / (\text{std} + 1\text{e-}8)$ measures the signal-to-noise ratio of the quantity-weighted bps series. The 1e-8 epsilon prevents division by zero in degenerate cases. A higher consistency metric indicates a clearer, more reliable signal of adverse selection.

This metric helps you understand how reliable your markout measurements are. If the mean adverse selection is high but the standard deviation (std) is also very high, the signal might be noisy and less actionable.

Exercise 7: Interpreting Consistency

You calculate the consistency metric for two different assets. Asset A has a consistency metric of 0.8, while Asset B has a consistency metric of 0.2. Both assets show negative average markouts. Which asset's markout signal would you trust more for making trading decisions, and why?

Your answer:

Finally, to get a single, digestible summary of adverse selection across all timeframes, a **composite score** is often used.

Key Concept: Composite Score

The final composite score is a weighted average of per-horizon tanh scores s_h across all markout horizons h , producing a single value in $(-1, +1)$ summarizing adverse selection consistency across time scales.

This composite score provides a holistic view, combining the insights from all your markout horizons into one easy-to-understand number. A score closer to -1 indicates strong, consistent adverse selection, while a score closer to +1 would suggest consistent favorable price movement (though this is less common for markouts designed to capture adverse selection).

Exercise 8: The Value of a Composite Score

You have markout data for 1s, 5s, 30s, 1min, and 5min horizons, each with its own mean and consistency. Why is it beneficial to condense all this information into a single composite score, especially if you need to quickly compare the adverse selection profile of many different trading strategies or assets?

Your answer:

Checklist: Key Metrics for Adverse Selection Analysis

- ☐ Define markouts for relevant events
- ☐ Understand the LP PnL markout formula
- ☐ Compute markouts at multiple time horizons
- ☐ Apply quantity-weighted markouts where appropriate
- ☐ Monitor VPIN for market imbalance signals
- ☐ Interpret the adverse selection component in optimal spread models
- ☐ Calculate the consistency metric for markout signals
- ☐ Generate a composite score for overall adverse selection assessment

Section E: Synthesis & Mastery – Connecting the Dots on the Spread

You've explored the fundamental elements of "Who's Paying the Spread." Now, let's bring these concepts together. In this section, you'll synthesize your understanding of markouts, adverse selection, and the underlying forces shaping the bid-ask spread. We'll also touch upon how statistical tools and information leakage play a role in a market maker's world.

Key Concept: Markout as a Performance Metric A **markout** is the realized PnL on a trade measured at some horizon τ after execution. It represents the realized price change between a fill price and the market mid-price at a later horizon τ , specifically measuring how the market moved after your execution.

Understanding Markout: Your Post-Trade View

Every trade you make as a liquidity provider (LP) has an outcome that unfolds over time. The markout helps you quantify this outcome.

Exercise 1: Defining Your Markout

In your own words, based on the definition provided, what is a "markout" and what does it aim to measure?

Your answer:

To accurately calculate a markout, you need a consistent reference point. The market's **BBO midpoint** ($(\text{best_bid} + \text{best_ask}) / 2$) serves this purpose. It's used both at the time of your trade's fill (mid_0) and at the chosen horizon τ (mid_τ).

Exercise 2: The Standard Reference Price

Imagine you're calculating the markout for a recent trade. Which specific price reference would you use to measure the market's position both at the moment of the trade and at your chosen analysis horizon?

Your answer:

Key Concept: Markout & Adverse Selection Markout is computed as $(\text{mid}_\tau - \text{mid}_0)$ signed from the LP perspective. This means if the market moves against the LP after a trade, that's a positive markout from the taker's side, indicating **adverse selection**.

Let's illustrate how markout reflects market movement.

SCENARIO	MID-PRICE AT FILL (MID_0)	MID-PRICE AT HORIZON T (MID_T)	MARKOUT (MID_T - MID_0)	LP IMPLICATION (FROM A SELL AT ASK)
Market moves up	100.00	100.05	+0.05	Negative for LP (adverse selection)
Market moves down	100.00	99.95	-0.05	Positive for LP
Market unchanged	100.00	100.00	0.00	Neutral for LP

Exercise 3: Interpreting Markout for a Buy Trade

Using the table above as a guide, consider an LP who *bought* one unit at the bid price. If the markout $(\text{mid}_\tau - \text{mid}_0)$ was -0.03, what does this imply for the LP's position, and does it suggest adverse selection?

Your answer:

Key Concept: Uninformed Traders For **uninformed (noise) traders**, the expected mid-price change after a fill is zero: $E[\text{mid}_\tau - \text{mid}_0 \mid \text{uninformed}] = 0$. This is a baseline for understanding when information might be present.

This concept is crucial for understanding the presence of information.

Exercise 4: The Baseline of Uninformed Trading

If you observe a consistent positive markout (from the taker's perspective) after your trades, what does the principle of uninformed traders tell you about the nature of the flow you are facing?

Your answer:

LP PnL and the Spread Components

The markout is a direct component of your realized profit and loss (PnL) as a liquidity provider. When you sell at the ask price, your PnL at horizon τ is clearly defined.

Key Concept: LP PnL Formula When a market maker sells one unit at ask price ($m_0 + \delta^a$), the LP PnL at horizon τ is: $PnL_{LP}(\tau) = \delta^a - (m_\tau - m_0)$. The second term, $(m_\tau - m_0)$, represents adverse selection if positive from the taker's side.

Exercise 5: Calculating LP PnL and Identifying Adverse Selection

You sell one unit at an ask price of 100.10. At the time of the fill, the BBO midpoint (m_0) was 100.05. One minute later (your horizon τ), the BBO midpoint (m_τ) is 100.12.

1. Calculate the markout ($m_\tau - m_0$).
2. Calculate your LP PnL(τ).
3. Was there adverse selection in this trade, and how do you know?

Your answer:

The **bid-ask spread** isn't just a single number; it's a reflection of multiple economic forces that a liquidity provider needs to cover.

ECONOMIC FORCE	DESCRIPTION
Inventory Risk	The risk that the value of your inventory (the assets you hold) changes unfavorably before you can offset your position.
Adverse Selection	The risk that you trade with someone who has better information than you, leading to losses as the market moves against your position shortly after your trade.
Order Processing Costs	The operational costs associated with handling and executing orders, including technology, connectivity, and personnel.

Exercise 6: Understanding Spread Components

Think about a time you bought or sold something with a noticeable difference between the buy and sell price (e.g., currency exchange, second-hand goods). Which of the three economic forces listed above do you think was most significant in determining that "spread" for the counterparty, and why?

Your answer:

Information Leakage: The Silent Threat

One significant contributor to adverse selection is **price reading (information leakage)**. This occurs when your own quotes, perhaps unintentionally, reveal information about your inventory or its sign.

Exercise 7: Identifying Price Reading

You notice that immediately after you adjust your bid and ask prices, certain sophisticated traders, informally called '**skew sniffers**', seem to anticipate market movements and trade profitably against you. How would you describe what's happening here, using the terminology you just learned?

Your answer:

Exercise 8: The Link Between Price Reading and Adverse Selection

Explain how price reading directly contributes to the adverse selection component of the bid-ask spread from a liquidity provider's perspective.

Your answer:

Leveraging Statistical Tools for Insight

To manage the risks associated with markouts and adverse selection, liquidity providers often employ statistical methods. Two such tools mentioned are the Exponentially Weighted Moving Average (EWMA) and the tanh-compressed score.

Key Concept: EWMA for Mean Estimation The Exponentially Weighted Moving Average (EWMA) is a special case of the Exponentially Weighted Moving Model (EWMM). It's obtained when fitting data with a constant model using a **square (quadratic) loss function** to estimate the mean. This makes it useful for tracking the average of a series that might change over time, giving more weight to recent observations.

Exercise 9: Applying EWMA Conceptually

You are an LP analyzing the markouts from your trades over the last few days. You suspect that the average markup (and thus, adverse selection) might be slowly increasing. How might using an EWMA help you detect and monitor this trend more effectively than a simple moving average?

Your answer:

When analyzing statistical signals, such as the significance of an observed markup or adverse selection, a **t-statistic** is often used. To make these signals more interpretable and comparable, they can be transformed.

Key Concept: Tanh-Compressed Score The **tanh-compressed score** $s = \tanh(t / 3.5)$ maps the unbounded **t-statistic** to $(-1, +1)$. The divisor 3.5 calibrates the saturation point so that $|t| \approx 3.5$ (a strong statistical signal) maps to $|s| \approx 0.998$. This provides a normalized measure of signal strength.

Exercise 10: Interpreting the Tanh-Compressed Score

You've calculated a t-statistic of 5.0 for the significance of adverse selection on your most recent batch of trades. Using the tanh-compressed score, how would you describe the strength and direction of this signal compared to a t-statistic of 1.0? What does the resulting 's' value tell you?

Your answer:

Capstone Exercise: Managing Your Spread

Exercise 11: The Informed Trader Scenario

You are an LP trading a liquid stock. Over the past hour, you've noticed a pattern:

- You consistently sell at your ask price ($m_0 + \delta^a$).
- Immediately after these fills, the BBO midpoint (m_τ) tends to rise significantly, resulting in a positive ($m_\tau - m_0$) markout from the taker's perspective.
- Your calculated LP PnL(τ) on these trades is consistently negative.

Based on what you've learned:

1. What phenomenon is likely occurring here, and which component of your bid-ask spread is most impacted?
2. What type of trader might you be interacting with, and what might they be doing?
3. How might you use EWMA or the tanh-compressed score to monitor and confirm this issue?

Your answer:

Checklist: Assessing Your Spread Management

☐

Do I regularly calculate markouts for my trades using the BBO midpoint?

☐

Can I identify the adverse selection component within my LP PnL?

☐

Am I aware of the three economic forces contributing to my bid-ask spread?

☐

Do I monitor for signs of price reading or information leakage in my trading activity?

☐

Do I have methods (like EWMA or t-statistic analysis) to statistically evaluate my markouts and adverse selection?

WHAT THE RESEARCH SAYS

For a bid fill, the LP PnL markout formula is: $PnL_{LP}(\tau) = \delta b + (m\tau - m_0)$.

For informed traders, the expected mid-price change after an informed buy follows: $E[m_\tau - m_0 \mid \text{informed buy}] = \mu(1 - e^{-\lambda\tau})$, where μ is the long-run price impact and λ is the signal decay rate.

For an ask fill, the LP PnL markout formula is: $PnL_{LP}(\tau) = \delta a - (m\tau - m_0)$, where δa is the half-spread captured and $(m\tau - m_0)$ is the mid-price change after the fill.

REFERENCES FOR THIS MODULE

1. A markout is the realized PnL on a trade measured at some horizon τ after execution.
2. A markout is defined as the realized PnL on a trade measured at some horizon τ after execution.
3. The warmup factor $warmup = 1 - (1-\alpha)^n$ (where n is fills so far) scales the t-statistic toward zero when few fills have been processed, preventing the system from acting on immature EWMA estimates.
4. In market making, adverse selection arises when the market maker trades against participants whose order flow reflects superior information about future price dynamics — specifically, the risk of being hit at the bid just before the market falls or at the ask just before it rises (the 'winner's curse').

— DEEP RESEARCH AUDIO —

References

1. A markout is the realized PnL on a trade measured at some horizon τ after execution.
2. A markout is defined as the realized PnL on a trade measured at some horizon τ after execution.
3. The warmup factor $\text{warmup} = 1 - (1-\alpha)^n$ (where n is fills so far) scales the t-statistic toward zero when few fills have been processed, preventing the system from acting on immature EWMA estimates.